ANSI X3H2-96-151
ISO/IEC JTC1/SC21/WG3 DBL ?

I S O
INTERNATIONAL ORGANIZATION FOR STANDARDIZATION
ORGANISATION INTERNATIONALE DE NORMALISATION

April 11, 1996

Subject:	SQL/Temporal
Status:	Change Proposal
Title:	Adding Valid Time to SQL/Temporal
Source:	ANSI Expert's Contribution
Authors:	Richard T. Snodgrass, Michael H. Böhlen, Christian S. Jensen and Andreas Steiner

Abstract: This change proposal specifies the addition of tables with valid-time support into SQL/Temporal, and explains how to use these facilities to migrate smoothly from a conventional relational system to a temporal system. Initially, important requirements to a temporal system that may facilitate such a transition are motivated and discussed. The proposal then describes the language additions necessary to add valid-time support to SQL3 while fulfilling these requirements. The constructs of the language are divided into four levels, with each level adding increased temporal functionality to its predecessor. The proposal formally defines the semantics of the query language by providing a denotational semantics mapping to well-defined algebraic expressions. Several alternatives for implementing the language constructs are listed. A prototype system implementing these constructs on top of a conventional DBMS is publicly available.

References

[1] Böhlen, M. H. and Marti R. *On the Completeness of Temporal Database Query Languages*, in *Proceedings of the First International Conference on Temporal Logic*. Ed. D. M. Gabbay, Ohlbach H. J.. Lecture Notes in Artificial Intelligence 827. Springer-Verlag, July 1994, pp. 283-300.

[2] Böhlen, M. H. *Valid-Time Integrity Constraints*, Aalborg University, October, 1995, 21 pages.

[3] Böhlen, M. H., C. S. Jensen and R. T. Snodgrass. *Evaluating the Completeness of TSQL2*, in *Proceedings of the VLDB International Workshop on Temporal Databases*. Ed. J. Clifford and A. Tuzhilin. VLDB. Springer Verlag, Sep. 1995.

[4] Clifford, J., A. Croker and A. Tuzhilin. *On Completeness of Historical Relational Query Languages. ACM Transactions on Database Systems*, 19, No. 1, Mar. 1994, pp. 64–116.

[5] Jackson, M. A. *System Development*. Prentice-Hall International Series in Computer Science. Prentice-Hall International, Inc., 1983.

[6] Jensen, C. S. and R. Snodgrass. *Temporal Specialization and Generalization. IEEE Transactions on Knowledge and Data Engineering*, 6, No. 6 (1994), pp. 954–974.

[7] Melton, J. (ed.) *SQL/Foundation*. October, 1995. (ISO/IEC JTC 1/SC 21/WG 3 DBL-LHR-002.)

[8] Melton, J. (ed.) *SQL/Temporal*. October, 1995. (ISO/IEC JTC 1/SC 21/WG 3 DBL-LHR-009.)

[9] Schueler, B. *Update Reconsidered*, in *Architecture and Models in Data Base Management Systems*. Ed. G. M. Nijssen. North Holland Publishing Co., 1977.

[10] Snodgrass, R. T., S. Gomez and E. McKenzie. *Aggregates in the Temporal Query Language TQuel. IEEE Transactions on Knowledge and Data Engineering*, 5, Oct. 1993, pp. 826–842.

[11] Snodgrass, R. T. and H. Kucera. *Rationale for Temporal Support in SQL3*. 1994. (ISO/IEC JTC1/SC21/WG3 DBL SOU-177, SQL/MM SOU-02.)

[12] Snodgrass, R. T., K. Kulkarni, H. Kucera and N. Mattos. *Proposal for a new SQL Part—Temporal*. 1994. (ISO/IEC JTC1/SC21 WG3 DBL RIO-75, X3H2-94-481.)

[13] Snodgrass, R. T. (editor) *The Temporal Query Language TSQL2*. Kluwer Academic Pub., 1995.

[14] Steiner, A. and M. H. Böhlen. The TimeDB Temporal Database Prototype, September, 1995. Available at `ftp://www.iesd.auc.dk/general/DBS/tdb/TimeCenter` or at `ftp://ftp.cs.arizona.edu/tsql/timecenter/TimeDB.tar.gz`.

[15] Tsichritzis, D.C. and F.H. Lochovsky. *Data Models*. Software Series. Prentice-Hall, 1982.

[16] Yourdon, E. *Managing the System Life Cycle*. Yourdon Press, 1982.

1 Introduction

This change proposal (a modification and extension to LHR-096, improving the proposed language extensions) introduces additions to SQL/Temporal to add valid-time support to SQL3. We outline a four-level approach for the integration of time. We motivate and discuss each level in turn, and we define the syntactic extensions that correspond to each level. We will see that the extensions are fairly minimal. Each level is described via a quick tour consisting of a set of examples. These examples have been tested in a prototype which is publicly available [14].

The proposed language constructs ensure temporal upward compatibility, sequenced valid semantics, and non-sequenced semantics, important properties that will be discussed in detail in Section 5.

2 The Problem

Most databases store time-varying information. For such databases, SQL is often the language of choice for developing applications that utilize the information in these databases. However, users also realize that SQL does not provide adequate support for temporal applications. To illustrate this, the reader is invited to attempt to formulate the following straightforward, realistic statements in SQL3. An intermediate SQL programmer can express all of them in SQL for a non-time-varying database in perhaps five minutes. However, even SQL experts find these same queries challenging to do in several *hours* when time-varying data is taken into account.

- An Employee table has three columns: Name, Manager and Dept. We then store historical information by adding a fourth column, When, of data type `PERIOD`. Manager is a foreign key for Employee.Name. This means that at each point in time, the character string value in the Manager column also occurs in the Name column (probably in a different row) at the same time. This cannot be expressed via SQL's foreign key constraint, which doesn't take time into account. Formulate this constraint instead as an assertion.

- Consider the query "List those employees who are not managers." This can easily be expressed in SQL, using `EXCEPT` or `NOT EXISTS`, on the original, three-column table. Things are just a little harder with the When column; a where predicate is required to extract the current employees. Now formulate the query "List those employees who were not managers, and indicate when." `EXCEPT` and `NOT EXISTS` won't work, because they don't consider time. This simple temporal query is challenging even to SQL experts.

- Consider the query "Give the number of employees in each department." Again, this is a simple query in SQL. Formulate the query "Give *the history of* the number of employees in each department." This query is extremely difficult without temporal support in the language.

- Now formulate the modification "Change the manager of the tools department for 1994 to Bob." This modification is difficult in SQL because only a portion of many validity periods needs be changed, with the information outside of 1994 retained.

Most users know only too well that while SQL is an extremely powerful language for writing queries on the current state, the language provides much less help when writing temporal queries, modifications, and constraints.

3 Outline of the Solution

The problem with formulating these SQL statements is due to the extreme difficulty of specifying in SQL the correct values of the timestamp column(s) of the result. The solution is to allow the DBMS to compute these values, moving the complexity from the application code into the DBMS. With the language extensions proposed in this change proposal, the above queries can all be easily written by an intermediate SQL programmer in about five minutes.

Referential integrity can be specified using sequenced valid semantics (which will be defined, exemplified, and provided a formal definition later in this document):

```
CREATE TABLE Employee(
   Name    VARCHAR(30),
   Manager VARCHAR(30) VALID REFERENCES Employee (Name),
   Dept    VARCHAR(20)) AS VALID DAY
```

Here we indicate that the table has valid-time support through "`AS VALID DAY`" and that the referential integrity is to hold for each point in time through "`VALID REFERENCES`".

For the query "List those employees who are not managers," we are interested only in the current employees. We use temporal upward compatibility to extract this information from the historical information stored in the Employee table.

```
SELECT Name FROM Employee EXCEPT SELECT Manager FROM Employee
```

This results in a conventional table, with one column.

We use sequenced valid semantics in the query "List those employees who were not managers, and when."

```
VALID SELECT Name FROM Employee EXCEPT SELECT Manager FROM Employee
```

The added "`VALID`" reserved word specifies that the query is to be evaluated at each point in time. At some times, an employee may not be a manager, whereas at other times, the employee is a manager. A one-column table results, but this time with valid-time support (i.e., the periods of time when each was not a manager is included).

The query "Give the number of employees in each department" is easy given temporal upward compatibility.

```
SELECT Dept, COUNT(*)
FROM Employee
GROUP BY Dept
```

Again, we just get the current count for each department. To extract the "*the history of* the number of employees in each department", only a simple change is required.

```
VALID SELECT Dept, COUNT(*)
FROM Employee
GROUP BY Dept
```

For each department, a time-varying count will be returned.

Modifications work in similar ways. The modification "Change the manager of the tools department for 1994 to Bob" can be expressed by following `VALID` with a period expression.

```
VALID PERIOD '[1994-01-01 - 1994-12-31]' UPDATE Employee
SET Manager = 'Bob'
WHERE Dept = 'Tools'
```

Here again, we exploit our knowledge of SQL to first write the update ignoring time, then change it in minor ways to take account of time.

These statements are reminiscent of the kinds of SQL statements that application programmers are called to write all the time. The potential for increased productivity is dramatic. Statements that previously took hours to write, or were simply too difficult to express, can take only minutes to write with the extensions discussed here.

4 Scope

Research on temporal databases has identified several properties crucial to temporal database systems, including support for valid-time, transaction time, temporal aggregates, indeterminacy, time granularity, user-defined calendars, vacuuming, and schema versioning. This document is the second in a series that will propose constructs for SQL/Temporal drawn from the consensus temporal query language TSQL2 [13]. The first [12], which was accepted in July, 1995, concerned the `PERIOD` data type.

The present change proposal addresses support for valid-time, specifically temporal upward compatibility, sequenced valid, and nonsequenced valid support. The next proposal will add support for transaction time. Future proposals will concern time granularities, temporal indeterminacy, and other features relevant to SQL3 that are fully supported in TSQL2. However, it is important that each proposal be comprehensive in its motivation of the additions, its presentation of the syntactic changes, and its specification of the semantics of the new constructs. For this reason, each change proposal should be separately considered and evaluated by the SQL3 standards committees.

While the language additions proposed here are modest, the productivity gains made available to the application programmer are significant. In particular, we will show how adding a single reserved word will convert any conventional (termed *snapshot*) query into a temporal query that extracts the history of the aspect being queried. This permits users to express rather complex temporal queries easily, by first formulating them as snapshot queries, then adding the reserved word. This parallel will be exploited in the semantics, permitting *any* SQL3 query to be rendered temporal. Moreover the syntactic modification not only holds for queries but also for view definitions, insert statements, delete statements, update statements, cursor declarations, table constraint definitions, column constraint definitions, and the definition of assertions.

We now return to the important question of migrating legacy databases. In the next section, we formulate several requirements of SQL/Temporal to allow graceful migration of applications from conventional to temporal databases.

5 Migration

The potential users of temporal database technology are enterprises with applications[1] that need to manage potentially large amounts of time-varying information. These include financial applications such as portfolio management, accounting, and banking; record-keeping applications, including personnel, medical records, and inventory; and travel applications such as airline, train, and hotel reservations and schedule management. It is most realistic to assume that these enterprises are already managing time-varying data and that the temporal applications are already in place and working. Indeed, the uninterrupted functioning of applications is likely to be of vital importance.

For example, companies usually have applications that manage the personnel records of their employees. These applications manage large quantities of time-varying data, and they may benefit substantially from built-in temporal support in the DBMS [11]. Temporal queries that are shorter and more easily formulated are among the potential benefits. This leads to improved productivity, correctness, and maintainability.

This section explores the problems that may occur when migrating database applications from an existing to a new DBMS, and it formulates a number of requirements to the new DBMS that must be satisfied in order to avoid different potential problems when migrating. Formal definitions of these requirements may be found in Appendix A.

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. As the existing model is given, the focus is on formulating requirements to the new data model.

Much of the section is 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. Further, we generally assume that the transition is from the SQL3 standard to SQL/Temporal.

[1] We use "database application" nonrestrictively, for denoting any software system that uses a DBMS as a standard component.

5.1 Upward Compatibility

Perhaps the most important aspect of 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 with the same functionality as with the existing system. The next two definitions are intended to capture what is needed for that to be possible.

We 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 [15]. For example, the central data structure of the relational model is the table, and the central, user-level query language is SQL.

We define a data model to be *syntactically upward compatible* with another data model if all the data structures and legal query expressions of the latter model are contained in the former model. For a data model to be *upward compatible* with another data model, we add the requirement that all queries expressible in the existing model must evaluate to the same results in both models. Syntactic upward compatibility implies that all existing databases and query expressions in the old system are also legal in the new system. The second condition guarantees that all existing queries 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.

To explore the relationship between SQL/Temporal and SQL3, we employ a series of figures that demonstrate increasing query and update functionality. In Figure 1, a conventional table 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 modification that took the table from one state to the next state.

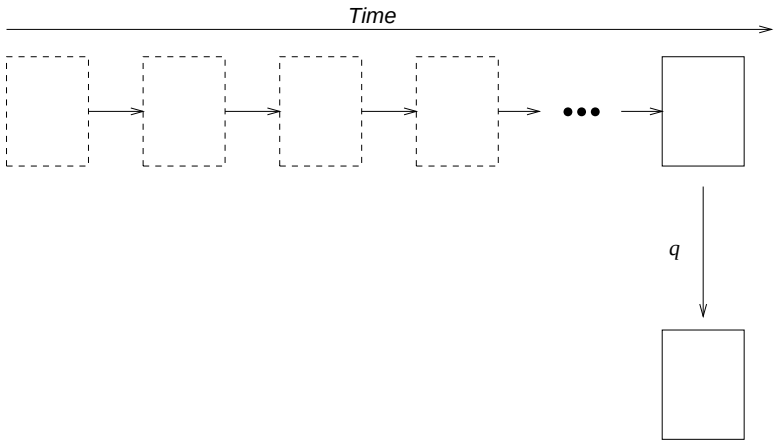

Figure 1: Level 1 evaluates an SQL3 query over a table without temporal support and returns a table also without temporal support

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 SQL3 are instances of tables in SQL/Temporal, (2) all SQL3 modifications to tables in SQL3 result in the same tables when the modifications are evaluated according to SQL/Temporal semantics, and (3) all SQL3 queries result in the same tables when the queries are evaluated according to SQL/Temporal.

By requiring that SQL/Temporal is a strict superset (i.e., only *adding* constructs and semantics), it is relatively easy to ensure that SQL/Temporal is upward compatible with SQL3.

Throughout, we provide examples of the various levels. In Section 6, we show these examples expressed in SQL/Temporal.

EXAMPLE 1: A company wishes to computerize its personnel records, so it creates two tables, an employee

table and a monthly salary table. Every employee must have a salary. These tables are populated. A view identifies those employees with a monthly salary greater than $3500. Then employee Therese is given a 10% raise. Since the salary table has no temporal support, Therese's previous salary is lost. These schema changes and queries can be easily expressed in SQL3. □

5.2 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 replaced with a new temporal model. No application code has been modified, and all tables have no temporal support. 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 table to add valid-time support (e.g., by using the `ALTER` statement of SQL/Temporal).

Note that it is undesirable to be forced to change the application code that accesses the table without temporal support that is replaced by a table with valid-time support. We formulate a requirement that states that the existing applications on tables without temporal support will continue to work with no changes in functionality when the tables they access are altered to add valid-time support. 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 tables with valid-time support.

Temporal upward compatibility is illustrated in Figure 2. When valid-time support is added to a table, the history is preserved, and modifications over time are retained. In this figure, the state to the far left 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 table without valid-time support.

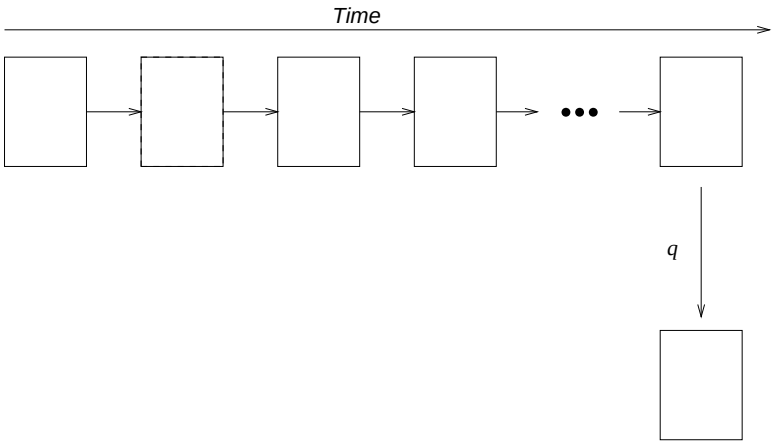

Figure 2: Level 2 evaluates an SQL3 query over a table with valid-time support and returns a table with similar support

The query q is an SQL3 query. Due to temporal upward compatibility the semantics of this query must not change if it is applied to a table with valid-time support. Hence, the query only applies to the current state, and a table without temporal support results.

EXAMPLE 2: 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 SQL3 `INSERT` statements, respects temporal upward compatibility. That means it is valid from now on. Queries and views on these tables with newly-added valid-time support work exactly as before. The

SQL3 query to list where high-salaried employees live returns the current information. Constraints and assertions also work exactly as before, applying to the current state and checked on database modification. □

It is instructive to consider 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 [16]); in the second, the design is based on the *structure* of the reality that the system is about (the "Jackson" approach [5]). It has been argued that the latter approach is superior because structure may remain stable when the function changes while the opposite is generally 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 attributes 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 data with temporal. The typical situation, when replacing a non-temporal system with a temporal system, 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 temporal data, which is encoded using tables without temporal support, in an ad hoc manner. When adopting the new system, upward compatibility guarantees that it is not necessary to change the database schema or application programs. However, without changes, the benefits of the added valid-time 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 table with valid-time support: the semantics of queries and modification are retained from tables without temporal support; the only difference is that intermediate states are also retained. Second, in those cases where the original 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, conventional tables that do contain temporal information and for which temporal support has been added can still be queried and modified by conventional SQL3 statements in a consistent manner.

5.3 Syntactically Similar Sequenced Temporal Extensions

The requirements covered so far have been aimed at protecting investments in legacy code and at ensuring uninterrupted operation of existing applications when achieving substantially increased temporal support by migrating to a temporal data model. Upward compatibility guarantees that (non-historical) legacy application code will continue to work without change when migrating, and temporal upward compatibility in addition allows legacy code to coexist with new temporal applications following the migration.

The requirement in this section aims at protecting the investments in programmer training and at ensuring continued efficient, cost-effective application development upon migration to a temporal model. This is achieved by exploiting the fact that programmers are likely to be comfortable with the non-temporal query language, e.g., SQL3.

Sequenced valid semantics states that SQL/Temporal must offer, for each query in SQL3, a temporal query that "naturally" generalizes this query, in a specific technical sense. In addition, we require that the SQL/Temporal query be syntactically similar to the SQL3 query that it generalizes.

With this requirement satisfied, SQL3-like SQL/Temporal queries on tables with temporal support have

semantics that are easily ("naturally") understood in terms of the semantics of the SQL3 queries on tables without temporal support. The familiarity of the similar syntax and the corresponding, naturally extended semantics makes it possible for programmers to immediately and easily write a wide range of temporal queries, with little need for expensive training.

Figure 3 illustrates this property. We have already seen that an SQL3 query q on a table with valid-time support applies the standard SQL3 semantics on the current state of that table, resulting in a table without temporal support. This figure illustrates a new query, q', which is an SQL/Temporal query. Query q' is applied to the table with valid-time support (the sequence of states across the top of the figure), and results in a table also with valid-time support, which is the sequence of states across the bottom.

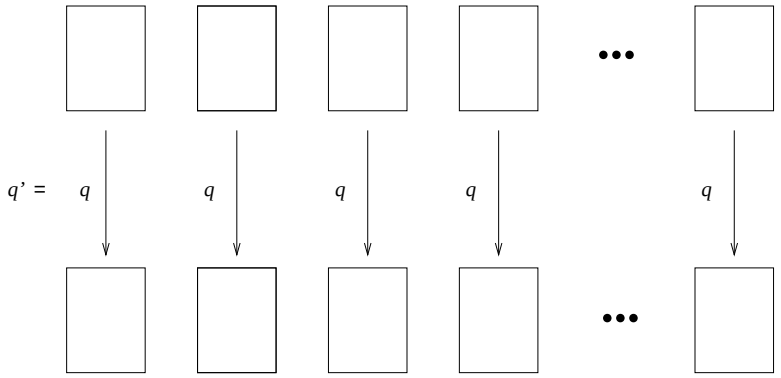

Figure 3: Level 3 evaluates an SQL/Temporal query over a table with valid-time support and returns a table with similar support

We would like the semantics of q' to be easily understood by the SQL3 programmer. Satisfying sequenced semantics along with the syntactical similarity requirement makes this possible. Specifically, the meaning of q' is precisely that of applying SQL3 query q on each state of the input table (which must have temporal support), producing a state of the output table for each such application. And when q' also closely resembles q syntactically, temporal queries are easily formulated and understood. To generate query q', one needs only prepend the reserved word `VALID` to query q.

This requirement ensures that the temporal model is a user-friendly (i.e., minor) extension of the snapshot model.

Example 3: We ask for the history of the monthly salaries paid to employees. Asking that question for the current state (i.e., what is the salary for each employee) is easy in SQL3; let us call this query q. To ask for the history, we simply prepend the keyword `VALID` to q to generate the SQL/Temporal query. Sequenced semantics allows us to do this for all SQL3 queries. So let us try a harder one: list *the history of* those employees for which no one makes a higher salary and lives in a different city. Again the problem reduces to expressing the SQL3 query for the current state. We then prepend `VALID` to get the history. Sequenced semantics also works for views, integrity constraints and assertions. □

These concepts also apply to sequenced *modifications*, illustrated in Figure 4. A valid-time modification destructively modifies states as illustrated by the curved arrows. As with queries, the modification is applied on a state-by-state basis. Hence, the semantics of the SQL/Temporal modification is a natural extension of the SQL modification statement that it generalizes.

Example 4: It turns out that a particular employee never worked for the company. That employee is deleted from the database. Note that if we use an SQL3 `DELETE` statement, temporal upward compatibility requires deleting the information only from the current (and future) states. By prepending the reserved word `VALID` to the `DELETE` statement, we can remove that employee from every state of the table.

Many people misspell the town Tucson as "Tuscon", perhaps because the name derives from an American Indian word in a language no longer spoken. To modify the current state to correct this spelling requires

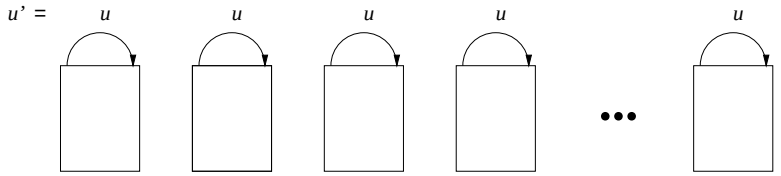

Figure 4: Level 3 also evaluates an SQL/Temporal modification on a table with valid-time support

a simple SQL `UPDATE` statement; let's call this statement u. To correct the spelling in all states, both past and possibly future, we simply prepend the reserved word `VALID` to u. □

5.4 Non-Sequenced Queries and Modifications

In a sequenced query, the information in a particular state of the resulting table with valid-time support is derived solely from information in the state at that same time of the source table(s). However, there are many reasonable queries that require other states to be examined. Such queries are illustrated in Figure 5, in which each state of the resulting table requires information from possibly all states of the source table.

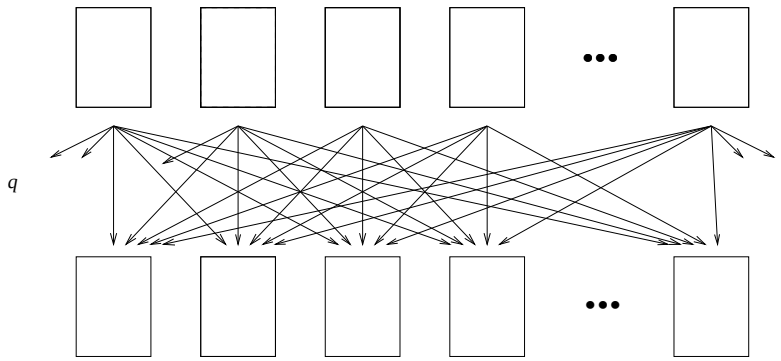

Figure 5: Level 4 evaluates a non-sequenced SQL/Temporal query over a table with valid-time support and returns a table with similar support

In this figure, two tables with valid-time support are shown, one consisting of the states across the top of the figure, and the other, the result of the query, consisting of the states across the bottom of the figure. A single query q performs the possibly complex computation, with the information usage illustrated by the downward pointing arrows. Whenever the computation of a single state of the result table may utilize information from a state at a different time, that query is non-sequenced. Such queries are more complex than sequenced queries, and they require new constructs in the query language.

Example 5: The query, "Who was given salary raises?", requires locating two consecutive times, in which the salary of the latter time was greater than the salary of the former time, for the same employee. Hence, it is a non-sequenced query. □

The concept of non-sequenced queries naturally generalizes to modifications. *Non-sequenced modifications* destructively change states, with information retrieved from possibly all states of the original table. In Figure 6, each state of the table with valid-time support is possibly modified, using information from possibly all states of the table before the modification. Non-sequenced modifications include future modifications.

Example 6: We wish to give employees a 5% raise if they have never had a raise before. This is not a temporally upward compatible modification, because the modification of the current state uses information

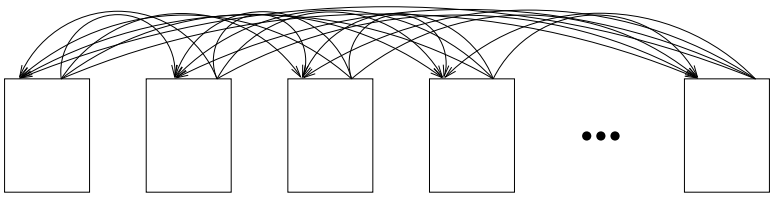

Figure 6: Level 4 also evaluates a non-sequenced SQL/Temporal modification on a table with valid-time support

in the past. For the same reason, it is not a sequenced update. So we must use a slightly more involved SQL/Temporal `UPDATE` statement. In fact, only the predicate "if they never had a raise" need be nonsequenced; the rest of the update can be temporally upward compatible. □

Views and cursors can also be nonsequenced.

EXAMPLE 7: We wish to define a snapshot view of the `salary` table in which the row's timestamp period appears as an explicit column. We can also define a valid-time view on this snapshot view that uses the explicit period column as an implicit timestamp. □

5.5 Summary

In this section, we have formulated three important requirements that SQL/Temporal should satisfy to ensure a smooth transition of legacy application code from SQL3's non-temporal data model to SQL/Temporal's data model. We review each in turn.

Upward compatibility and temporal upward compatibility guarantee that legacy application code needs no modification when migrating and that new temporal applications may coexist with existing applications. They are thus aimed at protecting investments in legacy application code.

The requirement that the temporal data model be a sequenced extension of the existing data model guarantees that a core subset of the temporal data model maximally builds on the existing data model and makes the temporal query language easy to use for programmers familiar with the existing query language. The requirement thus helps protect investments in programmer training. It also turns out that this property makes the semantics of tables with valid-time support straight-forward to specify, as shown in Section 7, and enables a wide range of implementation alternatives, some of which are listed in Section 8.2.

These requirements induce four levels of temporal functionality, to be defined in SQL/Temporal.

Level 1 This lowest level captures the minimum functionality necessary for the temporal query language to satisfy upward compatibility with SQL3. Thus, there is support for legacy SQL3 statements, but there are no tables with valid-time support and no temporal queries. Put differently, the functionality at this level is identical to that of SQL3.

Level 2 This level adds to the previous level solely by allowing for the presence of tables with valid-time support. The temporal upward compatibility requirement is applicable to this subset of SQL/Temporal. This level adds no new syntax for queries or modifications—only queries and modifications with SQL3 syntax are possible.

Level 3 The functionality of Level 2 is enhanced with the possibility of giving sequenced temporal functionality to queries, views, constraints, assertions, and modifications on tables with valid-time support. This level of functionality is expected to provide adequate support for many applications. Starting at this level, temporal queries exist, so SQL/Temporal must be a sequenced-consistent extension of SQL3.

Level 4 Finally, the full temporal functionality normally associated with a temporal language is added, specifically, non-sequenced temporal queries, assertions, constraints, views, and modifications. These additions include temporal queries and modifications that have no syntactic counterpart in SQL3.

6 Tables with Valid-Time Support in SQL3

This section informally introduces the new constructs of SQL/Temporal. These constructs are an improved and extended version of those in the consensus temporal query language TSQL2 [13]. The improvements concern guaranteeing the properties listed in Section 5, to support easy migration of legacy SQL3 application code [3]. The extensions concern views, assertions, and constraints (specifically temporal upward compatible and sequenced and non-sequenced extensions) that were not considered in the original TSQL2 design.

The presentation is divided into four levels, where each successive level adds temporal functionality. The levels correspond to those discussed informally in the previous section. Throughout, the functionality is exemplified with input to and corresponding output from a prototype system [14]. The reader may find it instructive to execute the sample statements on the prototype. In the examples, executable statements are displayed in `typewriter style` on a line of their own starting with the prompt "`> `".

6.1 Level 1: Upward Compatibility

Level 1 ensures upward compatibility (see Figure 1), i.e., it guarantees that legacy SQL3 statements evaluated over databases without temporal support return the result dictated by SQL3.

6.1.1 SQL3 Extensions

Obviously there are no syntactic extensions to SQL3 at this level.

6.1.2 A Quick Tour

The following statements are executed on January 1, 1995. A company creates two tables, an employee table and a monthly salary table. Every employee must have a salary. These schema changes can be easily expressed in SQL3.

```
> CREATE TABLE employee(ename VARCHAR(12), eno INTEGER PRIMARY KEY,
                        street VARCHAR(22), city VARCHAR(10), birthday DATE);
> CREATE TABLE salary(eno INTEGER REFERENCES employee(eno), amount INTEGER);

> CREATE ASSERTION emp_has_sal CHECK
  (NOT EXISTS ( SELECT *
                FROM employee AS e
                WHERE NOT EXISTS ( SELECT *
                                   FROM salary AS s
                                   WHERE e.eno = s.eno)));
```

These tables are populated.

```
> INSERT INTO employee
        VALUES ('Therese', 5873, 'Bahnhofstrasse 121', 'Zurich', DATE '1961-03-21');
> INSERT INTO employee
        VALUES ('Franziska', 6542, 'Rennweg 683', 'Zurich', DATE '1963-07-04');

> INSERT INTO salary VALUES (6542, 3200);
> INSERT INTO salary VALUES (5873, 3300);
```

A view identifies those employees with a monthly salary greater than $3500.

```
> CREATE VIEW high_salary AS SELECT * FROM salary WHERE amount > 3500;
```

Employee Therese is given a 10% raise. Since the salary table has no temporal support, Therese's previous salary is lost.

```
> UPDATE salary s
        SET amount = 1.1 * amount
  WHERE s.eno = (SELECT e.eno FROM employee e WHERE e.ename = 'Therese');

> COMMIT;
```

6.2 Level 2: Temporal Upward Compatibility

Level 2 ensures temporal upward compatibility as depicted in Figure 2. Temporal upward compatibility is straightforward for queries. They are evaluated over the current state of a database with valid-time support.

6.2.1 SQL3 Extensions

The create table statement is extended to define tables with valid-time support. Specifically, this statement can be followed by the clause "`AS VALID` <datetime field>", e.g., "`AS VALID DAY`". This specifies that the table has valid-time support, with states indexed by particular days. The alter table statement is extended to permit valid-time support to be added to a table without such support or dropped from a table with valid-time support.

A table with valid-time support is conceptually a sequence of states indexed with valid-time granules at the specified granularity. This is the view of a table with valid-time support adopted in temporal upward compatibility and sequenced semantics. At a more specific logical level, a table with valid-time support is *also* a collection of rows associated with valid-time periods.

Indeed, our definition of the semantics of the addition to SQL/Temporal being proposed satisfies temporal upward compatibility and sequenced semantics.

6.2.2 A Quick Tour

The following statements are executed on February 1, 1995.

```
> ALTER TABLE salary ADD VALID DAY;
> ALTER TABLE employee ADD VALID DAY;
```

The following statements are typed in the next day (February 2, 1995).

```
> INSERT INTO employee
        VALUES('Lilian', 3463, '46 Speedway', 'Tuscon', DATE '1970-03-09');
> INSERT INTO salary VALUES(3463, 3400);
> COMMIT;
```

The `employee` table contains the following rows. (In these examples, we used open-closed ("`[...)`") for periods.)

ename	eno	street	city	birthday	Valid
Therese	5873	Bahnhofstrasse 121	Zurich	1961-03-21	[1995-02-01 - 9999-12-31)
Franziska	6542	Rennweg 683	Zurich	1963-07-04	[1995-02-01 - 9999-12-31)
Lilian	3463	46 Speedway	Tuscon	1970-03-09	[1995-02-02 - 9999-12-31)

Note that the valid time extends to *forever*, which in SQL3 is the largest date.

The `salary` table contains the following rows.

eno	amount	Valid
6542	3200	[1995-02-01 - 9999-12-31)
5873	3630	[1995-02-01 - 9999-12-31)
3463	3400	[1995-02-02 - 9999-12-31)

We continue, still on February 2. Tables, views, and queries act like before, because temporal upward compatibility is satisfied. To find out where the high-salaried employees live, use the following.

```
> SELECT ename, city
  FROM   high_salary AS s, employee AS e
  WHERE  s.eno = e.eno;
```

Evaluated over the current state, this returns the employee Therese, in Zürich.

Assertions and referential integrity act like before, applying to the current state. The following transaction will abort due to (1) a violation of the `PRIMARY KEY` constraint, (2) a violation of the `emp_has_sal` assertion and (3) a referential integrity violation, respectively.

```
> INSERT INTO employee
        VALUES ('Eric', 3463, '701 Broadway', 'Tucson', DATE '1988-01-06');
> INSERT INTO employee
        VALUES ('Melanie', 1234, '701 Broadway', 'Tucson', DATE '1991-03-08');
> INSERT INTO salary VALUES(9999, 4900);
> COMMIT;
```

6.3 Level 3: Sequenced Language Constructs

Level 3 adds syntactically similar, sequenced counterparts of existing queries, modifications, views, constraints, and assertions (see Figure 3). Sequenced SQL/Temporal queries produce tables with valid-time support. The state of a result table at each time is computed from the state of the underlying table(s) at the same time, via the semantics of the contained SQL3 query. In this way, users are able to express temporal queries in a natural fashion, exploiting their knowledge of SQL3. Temporal views, assertions and constrains can likewise be naturally expressed.

6.3.1 SQL3 Extensions

Temporal queries, modifications, views, assertions, and constraints are signalled by the reserved word `VALID`. This reserved word can appear in a number of locations; Section 10 supplies the details.

Derived table in a from clause In the from clause, one can prepend `VALID` to a <query expression>.

View definition Temporal views can be specified, with sequenced semantics.

Assertion definition A sequenced assertion applies to each of the states of the underlying table(s). This is in contrast to a snapshot assertion, which is only evaluated on the current state. In both cases, the assertion is checked before a transaction is committed.

Table and column constraints When specified with `VALID`, such constraints must apply to all states of the table with valid-time support.

Cursor expression Cursors can range over tables with valid-time support.

Single-row select Such a select can return a row with an associated valid time.

Fetch statement The period associated with a row with valid-time support can be placed in a local variable in embedded SQL.

Modification statements When specified with `VALID`, the modification applies to each state comprising the table with valid-time support.

In all cases, the `VALID` reserved word indicates that sequenced semantics is to be employed.

6.3.2 A Quick Tour

We evaluate the following statements on March 1, 1995.

Prepending `VALID` to any `SELECT` statement evaluates that query on all states, in a sequenced fashion. The first query provides the history of the monthly salaries paid to employees. This query is constructed by first writing the snapshot query, then prepending `VALID`.

```
> VALID
     SELECT ename, amount
     FROM   salary AS s, employee AS e
     WHERE  s.eno = e.eno;
```

This evaluates to the following.

ename	amount	Valid
Franziska	3200	[1995-02-01 - 9999-12-31)
Therese	3630	[1995-02-01 - 9999-12-31)
Lilian	3400	[1995-02-02 - 9999-12-31)

List those for which no one makes a higher salary in a different city, over all time.

```
> VALID
     SELECT ename
     FROM   employee AS e1, salary AS s1
     WHERE  e1.eno = s1.eno
     AND NOT EXISTS (SELECT ename
                     FROM   employee AS e2, salary AS s2
                     WHERE  e2.eno = s2.eno
                         AND    s2.amount > s1.amount
                         AND    e1.city <> e2.city);
```

This gives the following result.

ename	Valid
Therese	[1995-02-01 - 9999-12-31)
Franziska	[1995-02-01 - 1995-02-02)

Therese is listed because the only person in a different city, Lilian, makes a lower salary. Franziska is listed because for that one day, there was no one in a different city (Lilian didn't join the company until February 2).

We then create a temporal view, similar to the non-temporal view defined earlier. In fact, the only difference is the use of the reserved word `VALID`.

```
> CREATE VIEW high_salary_history AS
      VALID SELECT * FROM salary WHERE s.salary > 3500;
```

Finally, we define a temporal column constraint.

```
> ALTER TABLE salary ADD VALID CHECK (amount > 1000 AND amount < 12000);
> COMMIT;
```

Rather than being checked on the current state only, this constraint is checked on each state of the `salary` table. This is useful to restrict *retroactive* changes [6], i.e., changes to past states and *proactive* changes, i.e., changes to future states. This constraint is satisfied for all states in the table.

Sequenced modifications are similarly handled. To remove employee #5873 for all states of the database, we use the following statement.

```
> VALID DELETE FROM employee
        WHERE eno = 5873;
> VALID DELETE FROM salary
        WHERE eno = 5873;
> COMMIT;
```

To correct the common misspelling of Tucson, we use the following statement.

```
> VALID UPDATE employee
           SET city = 'Tucson'
        WHERE city = 'Tuscon';
> COMMIT;
```

This updates all incorrect values, at all times, including the past and future. Lillian's city is thus corrected.

6.4 Level 4: Non-Sequenced Language Constructs

Level 4 accounts for non-sequenced queries (see Figure 5) and non-sequenced modifications (see Figure 6). Many useful queries and modifications are in this category. However, their semantics is necessarily more complicated than that of sequenced queries, because non-sequenced queries cannot exploit that useful property. Instead, they must support the formulation of special-purpose user-defined temporal relationships between implicit timestamps, datetime values expressed in the query, and stored datetime columns in the database.

Nonsequenced SQL/Temporal queries can produce tables with or without valid-time support, depending on whether the valid-time period of the resulting rows is provided in the query. The state of a result table, if a table is without valid-time support, or the state of a result table at each time, if a table has valid-time support, is computed from potentially all of the states of the underlying table(s), at any time. The semantics are quite simple. A nonsequenced evaluation treats a table with valid-time support as a table without temporal support, but with an additional column containing the timestamp.

6.4.1 SQL3 Extensions

Nonsequenced valid queries are signalled by the new reserved word `NONSEQUENCED` preceding the reserved word `VALID`. This applies analogously to nonsequenced modifications, views, assertions, and constraints. This reserved word can appear in a number of locations; Section 10 supplies the details.

Derived table in a from clause In the from clause, one can prepend `NONSEQUENCED VALID` to a <query expression>. This results in a table without temporal support, and is the means of removing the valid-time support of a table.

View definition Nonsequenced views can be specified.

Assertion definition A nonsequenced assertion applies simultaneously to all of the states of the underlying table(s). This is in contrast to a snapshot assertion, which is only evaluated on the current state. In both cases, the assertion is checked before a transaction is committed.

Table and column constraints When specified with `NONSEQUENCED VALID`, such constraints must apply to the table with valid-time support as a whole.

Cursor expression Cursors can range over the result of a nonsequenced select.

Single-row select A nonsequenced single-row select will return a row without temporal support, even when evaluated over tables with valid-time support.

Modification statements When specified with `NONSEQUENCED VALID`, the modification applies simultaneously to all states comprising the table with valid-time support.

In all cases, the `NONSEQUENCED` reserved word indicates that nonsequenced semantics is to be employed.

This portion includes other useful, related constructs.

- An optional period expression after `VALID` specifies that the valid-time period of each row of the result is intersected with the value of the expression. This allows one to restrict the result of a select statement, cursor expression, or view definition to a specified period, and to restrict the time for which assertion definitions, table constraints and column constraints are checked.
- An optional period expression after `NONSEQUENCED VALID` specifies the valid-time period of each row of the result, and thus renders the resulting table to have valid-time support. This enables a table without temporal support to be converted into a table with valid-time support within a query or other statement.
- For modification statements, the period expression after `VALID` and `VALID NONSEQUENCED` specifies the temporal scope of the modification: the times at which the modification is to be applied.
- The value expression "`VALID(`<correlation name>`)`" evaluates to the valid-time period of the row associated with the correlation or table name. This is required because valid-time periods of tables with valid-time support are not explicit columns (the alternative violates temporal upward compatibility).

The following quick tour provides examples of these constructs.

6.4.2 A Quick Tour

This quick tour starts with the database as it was when we last left it, in the previous quick tour. The `employee` table has the following contents.

ename	eno	street	city	birthday	Valid
Franziska	6542	Rennweg 683	Zurich	1963-07-04	[1995-02-01 - 9999-12-31)
Lilian	3463	46 Speedway	Tucson	1970-03-09	[1995-02-02 - 9999-12-31)

The `salary` table has the following contents.

eno	amount	Valid
6542	3200	[1995-02-01 - 9999-12-31)
3463	3400	[1995-02-02 - 9999-12-31)

A period expression after `VALID` specifies the temporal scope of the result. List those who were employed sometime during the first six months.

```
> VALID PERIOD '[1995-01-01 - 1995-07-01)' SELECT ename FROM employee;
```

This returns the following table.

ename	Valid
Franziska	[1995-02-01 - 1995-07-01)
Lilian	[1995-02-02 - 1995-07-01)

On April 1, 1995, we give Lilian a 5% raise, starting immediately. This is a temporally upward compatible modification, and so is already expressible in SQL.

```
> UPDATE salary
  SET amount = 1.05 * amount
  WHERE eno = (SELECT S.eno
               FROM salary AS S, employee as E
               WHERE ename = 'Lilian' AND E.eno = S.eno);

> COMMIT;
```

This results in the following `salary` table.

eno	amount	Valid
6542	3200	[1995-02-01 - 9999-12-31)
3463	3400	[1995-02-02 - 1995-04-01)
3463	3570	[1995-04-01 - 9999-12-31)

To determine who was given salary *raises*, we must simultaneously consider two consecutive states of the `salary` table, before and after the raise. This requires a nonsequenced query.

```
> NONSEQUENCED VALID SELECT ename
  FROM employee AS E, salary AS S1, salary AS S2
  WHERE E.eno = S1.eno AND E.eno = S2.eno
     AND S1.amount < S2.amount AND VALID(S1) MEETS VALID(S2);
```

`MEETS` ensures that the valid-time period associated with `S1` is immediately followed by the valid-time period associated with `S2`. Since the valid-time period of a row is not in an explicit column (as this would violate temporal upward compatibility), `VALID()` is used to extract the associated valid-time period. The result is a table without temporal support, because `NONSEQUENCED` is not followed by a period expression.

ename
Lilian

If we instead wish to get back a table with valid-time support, i.e., "Who was given salary raises, and when did they receive the higher salary?", we place a <value expression> after `VALID` to specify when each resulting row is valid. Our first try is the following.

```
> NONSEQUENCED VALID VALID(S2) SELECT ename
  FROM employee AS E, salary AS S1, salary AS S2
  WHERE E.eno = S1.eno AND E.eno = S2.eno
     AND S1.amount < S2.amount AND VALID(S1) MEETS VALID(S2);
```

This isn't quite correct, because the period expression following `VALID` can only mention the columns of the following select statement. So we put the value in the select list, and use an enclosing (sequenced) select statement to get rid of this extra column.

```
> VALID SELECT ename
  FROM (NONSEQUENCED VALID S2valid SELECT ename, VALID(S2) AS S2valid
        FROM employee AS E, salary AS S1, salary AS S2
        WHERE E.eno = S1.eno AND E.eno = S2.eno
           AND S1.amount < S2.amount AND VALID(S1) MEETS VALID(S2) ) AS S;
```

This query has the following result.

ename	Valid
Lilian	[1995-04-01 - 9999-12-31)

If we had desired the time when the person had received the *lower* salary, we would simply specify `VALID(S1)` instead.

Following `VALID` with a period expression in a modification (whether sequenced or not) specifies the temporal scope of the modification. Two applications of this are retroactive and future changes. Assume it is now May 1, 1995. Franziska, employee 6542, will be taking a leave of absence the last half of the year.

```
> VALID PERIOD '[1995-07-01 - 1996-01-01)'
  DELETE FROM salary
  WHERE eno = 6542;

> VALID PERIOD '[1995-07-01 - 1996-01-01)'
  DELETE FROM employee
  WHERE eno = 6542;

> COMMIT;
```

The `salary` table now has the following contents.

eno	amount	Valid
6542	3200	[1995-02-01 - 1995-07-01)
6542	3200	[1996-01-01 - 9999-12-31)
3463	3400	[1995-02-02 - 1995-04-01)
3463	3570	[1995-04-01 - 9999-12-31)

The `employee` table has the following contents.

ename	eno	street	city	birthday	Valid
Franziska	6542	Rennweg 683	Zurich	1963-07-04	[1995-02-01 - 1995-07-01)
Franziska	6542	Rennweg 683	Zurich	1963-07-04	[1996-01-01 - 9999-12-31)
Lilian	3463	46 Speedway	Tucson	1970-03-09	[1995-02-02 - 9999-12-31)

Note that these deletions split single periods into two, with a lapse between them. Many modifications are greatly simplified in this way. Also note that previously specified sequenced valid referential integrity and other constraints and assertions must apply to each state. Hence, if the first `DELETE` was performed, but not the second, the `COMMIT` will abort because the `emp_has_sal` constraint is violated for certain states, such as the one on August 1, 1995.

The period expression following `VALID` is also allowed for assertions and constraints. Assume that no employee may make less than 3000 during 1996.

```
> CREATE ASSERTION salary_check
  VALID PERIOD '[1996-01-01 - 1997-01-01)' CHECK
     (NOT EXISTS ( SELECT * FROM salary WHERE amount < 3000 ) );
```

This is a sequenced assertion, and thus applies separately to each state (at least, those in 1996). Nonsequenced assertions and constraints apply to all states at once. To assert that there is only one employee with a particular name, we use the following constraint within the `employee` table definition.

```
> CONSTRAINT unique_name UNIQUE (ename)
```

This is interpreted with temporal upward compatible semantics, and so applies only to the current state. If all we do is temporal upward compatible modifications, this will be sufficient. However, if we perform future updates, violations may be missed. To always check all states, a sequenced constraint is used.

```
> CONSTRAINT unique_name_per_time VALID UNIQUE (ename)
```

This will ensure that at any time, each ename value is unique.

To ensure that each ename is unique, *across all states simultaneously*, a nonsequenced constraint is required.

```
> CONSTRAINT unique_name_over_all_time NONSEQUENCED VALID UNIQUE (ename)
```

The above `employee` table satisfies the first two constraints, but not the third (the nonsequenced one), because there are two rows with an ename of Franziska.

As with `VALID`, `NONSEQUENCED VALID` can appear in a from clause. To give employees a 5% raise if they never had a raise before, we first write a temporal upward compatible modification (i.e., without `VALID`) to give the raise.

```
> UPDATE salary AS S
  SET amount = 1.05 * amount;
```

We can augment this statement to use a non-sequenced query in the from clause to look for raises in the past.

```
> UPDATE salary AS S
  SET amount = 1.05 * amount
  WHERE NOT EXISTS (SELECT *
                    FROM (NONSEQUENCED VALID SELECT *
                          FROM salary AS S1, salary AS S2
                          WHERE S1.amount < S2.amount
                              AND VALID(S1) MEETS VALID(S2)
                              AND S1.eno = S.eno and S2.eno = S.eno) AS S3
                    );
> COMMIT;
```

The `NOT EXISTS` was added. Assume that the update was entered on June 1, 1995. The following `salary` table results.

eno	amount	Valid
6542	3200	[1995-02-01 - 1995-06-01)
6542	3360	[1995-06-01 - 1995-07-01)
6542	3360	[1996-01-01 - 9999-12-31)
3463	3400	[1995-02-02 - 1995-04-01)
3463	3570	[1995-04-01 - 9999-12-31)

Since the update is evaluated with temporal upward compatible semantics, if changes the salary for valid times after June 1.

Finally, we wish to define a snapshot view of the `salary` table in which the row's timestamp appears as an explicit column.

```
> CREATE VIEW snapshot_salary (eno, amount, validtime) AS
  NONSEQUENCED VALID SELECT S.*, VALID(S) FROM salary AS S;
```

Coming around full circle, we can define a valid-time view on `snapshot_salary` that uses the explicit column `validtime` as an implicit timestamp.

```
> CREATE VIEW temporal_salary (eno, amount) AS
  VALID SELECT eno, amount
  FROM (NONSEQUENCED VALID validtime SELECT * FROM snapshot_salary AS S) AS S2;
```

This conversion can also be applied within queries and cursors.

7 Formal Semantics of SQL/Temporal

In this section, we provide a formal semantics for the constructs introduced into SQL/Temporal, expressed in terms of the relational algebraic semantics for SQL3.

We use $\langle t \| VT \rangle$ to denote a row in a table with valid-time support. The vertical double-bar "$\|$" is used to separate valid-time from explicit attributes. If VT is a period, then VT^- is its beginning delimiting timestamp and VT^+ is the granule following its ending delimiting timestamp.

7.1 Translating SQL/Temporal Queries to Relational Algebra Expressions

We first provide the semantics of an SQL3 query over tables without temporal support. In the definition given next, let $r_1, \ldots, r_n$ denote tables without temporal support. We base the definition of the semantics on the semantics of SQL3, expressed in terms of the relational algebra.

$$[\![\text{<query expression>}]\!]_{\text{SQL/T}}(r_1, \ldots, r_n) \triangleq [\![\text{<query expression>}]\!]_{\text{standard}}(r_1, \ldots, r_n)$$

Here, $[\![$<query expression>$]\!]_{\text{standard}}$, which evaluates to the relational algebra expression that corresponds to <query expression>, is assumed to be given. This definition satisfies upward compatibility.

EXAMPLE 8: We start with a non-temporal query, i.e., a query evaluated with standard semantics. Assume `p` and `q` are both tables without temporal support. The query Q_1

```
SELECT p.X
FROM p, q
WHERE p.X = q.X
```

is equivalent to the relational algebra expression

$$[\![Q_1]\!]_{\text{SQL/T}}(p,q) = [\![Q_1]\!]_{\text{standard}}(p,q) = \pi_{p.X}(p \bowtie_{p.X=q.X} q) \quad .$$

□

The semantics of an SQL3 query over a combination of snapshot and tables with valid-time support is very similar. For every table with valid-time support r_i appearing as an argument, replace it with $\tau^{vt}_{now}(r_i)$ on the right hand side. The valid-timeslice operator τ^{vt}_c extracts the current snapshot state from a table with valid-time support.

$$\tau^{vt}_c(r) \quad \triangleq \quad \{t \mid \exists VT(\langle t \| VT\rangle \in r \wedge VT^- \leq c \wedge c < VT^+)\}$$

EXAMPLE 9: We now examine a non-temporal query over a combination of tables with and without temporal support, with standard semantics. Assume `p` is a table without temporal support and `t` is a table with valid-time support. The query Q_1

```
SELECT p.X
FROM p, t
WHERE p.X = t.X
```

is equivalent to the relational algebra expression

$$[\![Q_1]\!]_{\text{SQL/T}}(p,t) = [\![Q_1]\!]_{\text{standard}}(p, \tau^{vt}_{now}(t)) = \pi_{p.X}(p \bowtie_{p.X=t.X} (\tau^{vt}_{now}(t))) \quad .$$

□

This definition satisfies temporal upward compatibility.

Next, we define the semantics of sequenced SQL/Temporal additions in terms of the snapshot semantics. This allows these extensions to be consistent with all snapshot constructs defined in SQL3.

$$[\![\texttt{VALID}\ \text{<query expression>}]\!]_{\text{SQL/T}}(r_1,\ldots,r_n) \triangleq [\![\text{<query expression>}]\!]_{\text{temporal}}(r_1,\ldots,r_n)$$

In this definition, $[\![$<query expression>$]\!]_{\text{temporal}}$ is equivalent to $[\![$<query expression>$]\!]_{\text{standard}}$, except that every non-temporal relational algebra operator (e.g., $\bowtie, \sigma, \pi$) is replaced by a corresponding temporal relational algebra operator (e.g., $\bowtie^{vt}, \sigma^{vt}, \pi^{vt}$). We provide definitions of the temporal algebra in Section 7.3.

EXAMPLE 10: An SQL/Temporal query $Q_2 =$ `VALID` Q_1 is evaluated with temporal semantics, due to its leading valid clause. Both `p` and `q` must be tables with valid-time support. Thus,

```
VALID
  SELECT p.X
  FROM p, q
  WHERE p.X = q.X
```

is equivalent to the temporal relational algebra expression

$$[\![Q_2]\!]_{\mathrm{SQL/T}}(p,q) = [\![\texttt{VALID } Q_1]\!]_{\mathrm{SQL/T}}(p,q) = [\![Q_1]\!]_{\mathrm{temporal}}(p,q) = \pi^{vt}_{p.X}(p \bowtie^{vt}_{p.X=q.X} q) \quad .$$

Note that apart from the ^{vt}-superscripts, which are added to relational algebra operators, the translation between SQL queries and relational algebra expressions has not changed at all. □

The definitions above satisfy sequenced semantics if the temporal relational operators are sequenced with respect to their conventional relational counterparts. The next step is to define a temporal relational algebra with this property.

7.2 The Conventional Relational Algebra

As a precursor to defining the temporal relational algebra, we review Codd's relational algebra.

$$\begin{array}{lcl}
\sigma_c(r) & \triangleq & \{t \mid t \in r \wedge c(t)\} \\
\pi_f(r) & \triangleq & \{t_1 \mid t_2 \in r \wedge t_1 = f(t_2)\} \\
r_1 \cup r_2 & \triangleq & \{t \mid t \in r_1 \vee t \in r_2\} \\
r_1 \bowtie_c r_2 & \triangleq & \{t_1 \circ t_2 \mid t_1 \in r_1 \wedge t_2 \in r_2 \wedge c(\langle t_1 \circ t_2 \rangle)\} \\
r_1 \setminus r_2 & \triangleq & \{t \mid t \in r_1 \wedge t \notin r_2\} \\
AG_{agg,f}(r) & \triangleq & \{t \circ a \mid t \in r \wedge a = agg(\{t_1 | t_1 \in r \wedge f(t_1) = f(t)\})\}
\end{array}$$

In this formalism, c is a predicate, f is a list of attributes (for the aggregate operator, a list of the `GROUP BY` attributes), and agg is a function (e.g., sum_3) that when applied to a set of rows returns the single value of the aggregate (e.g., `SUM`) evaluated over the indicated attribute (e.g., the third attribute).

Observe that the algebra defined above is based on sets and thus does not permit duplicates. We have chosen to assume a set-based framework in the semantics given here because this yields a short definition where the general approach stands out more clearly. The complications that follow from giving up the set-based basis have been explored in the past and are omitted. We emphasize that the proposed additions to SQL/Temporal follow the data model of SQL3 and are not strictly set based.

7.3 The Temporal Relational Algebra

The next step is to define the temporal relational algebra operators. Informally, each definition respects sequenced semantics. In addition to that, the algebra features two properties which we would like to point out. First, the algebra preserves the periods entered into the database, i.e., it *matters* for the query results whether we store, e.g., one row with a valid-time period of $[10-20]$ or two (value-equivalent) rows with valid-time periods $[10-15]$ and $[16-20]$, respectively. Second, care was taken to only consider end points of valid-time periods of rows when implementing the operators—intermediate time points are never used. This allows for an efficient (essentially, granularity independent) implementation.

In Figure 7, the constructor *intersect* (over two periods) returns a period containing those chronons in both underlying periods, and the predicate *overlaps* (over two periods) returns true if the two periods overlap and false, otherwise. Both operations are easily expressed as operations on the beginning and ending delimiting timestamps of periods. The symbol "$\circ$" denotes concatenation. The definition of AG^{vt} is especially complex. It determines *constant periods*, during which no row starts or ends [10]. A constant period can go from the start of one row to the start of another, from the start of one row to the end of another, or from the end of one row to the end of another.

7.4 Nonsequenced Semantics

When `NONSEQUENCED VALID` is used, each table with valid-time support is converted to a table without temporal support via the SN function.

$$SN(r) \triangleq \{\langle t, VT\rangle \mid \langle t \,||\, VT\rangle \in r\}$$

$$
\begin{array}{lcl}
\sigma_c^{vt}(r) & \triangleq & \{\langle t \| VT\rangle \mid \langle t \| VT\rangle \in r \wedge c(\langle t \| VT\rangle)\} \\
\pi_f^{vt}(r) & \triangleq & \{\langle t_1 \| VT\rangle \mid \langle t_2 \| VT\rangle \in r \wedge t_1 = \langle f(t_2) \| VT\rangle\} \\
r_1 \cup^{vt} r_2 & \triangleq & \{\langle t \| VT\rangle \mid \langle t \| VT\rangle \in r_1 \vee \langle t \| VT\rangle \in r_2\} \\
r_1 \bowtie_c^{vt} r_2 & \triangleq & \{\langle \langle t_1, VT_1\rangle \circ \langle t_2, VT_2\rangle \| VT\rangle \mid \langle t_1 \| VT_1\rangle \in r_1 \wedge \langle t_2 \| VT_2\rangle \in r_2 \wedge \\
 & & \quad c(\langle t_1, VT_1\rangle \circ \langle t_2, VT_2\rangle) \wedge \\
 & & \quad VT = \mathit{intersect}(VT_1, VT_2) \wedge VT_1 \mathit{\ overlaps\ } VT_2\} \\
r_1 \setminus^{vt} r_2 & \triangleq & \{\langle t \| VT\rangle \mid \langle t \| VT_1\rangle \in r_1 \wedge \\
 & & \quad (\exists VT_2(\langle t \| VT_2\rangle \in r_2 \wedge VT_1^- \leq VT_2^+ \wedge VT^- = VT_2^+) \vee VT^- = VT_1^-) \wedge \\
 & & \quad (\exists VT_3(\langle t \| VT_3\rangle \in r_2 \wedge VT_1^+ \geq VT_3^- \wedge VT^+ = VT_3^-) \vee VT^+ = VT_1^+) \wedge \\
 & & \quad VT^- < VT^+ \wedge \\
 & & \quad \neg\exists VT_4(\langle t \| VT_4\rangle \in r_2 \ \wedge VT_4^+ > VT^- \ \wedge VT_4^- < VT^+)\} \\
AG_{agg,f}^{vt}(r) & \triangleq & \{\langle t \circ a \| VT\rangle \mid \langle t \| VT_1\rangle \in r \ \wedge \langle t_2 \| VT_2\rangle \in r \ \wedge f(t) = f(t_2) \wedge \\
 & & \quad ((VT^- = VT_1^- \ \wedge VT^+ = VT_2^-) \vee \\
 & & \quad\quad (VT^- = VT_1^- \ \wedge VT^+ = VT_2^+) \vee \\
 & & \quad\quad (VT^- = VT_1^+ \ \wedge VT^+ = VT_2^+)) \ \wedge VT^- < VT^+ \wedge \\
 & & \quad \neg\exists\langle t_4 \| VT_4\rangle \in r(f(t) = f(t_4) \wedge \\
 & & \quad\quad ((VT^- < VT_4^- < VT^+) \vee (VT^- < VT_4^+ < VT^+))) \wedge \\
 & & \quad a = agg(\{t_3 | \langle t_3 \| VT_3\rangle \in r \ \wedge VT_3 \mathit{\ overlaps\ } VT \wedge f(t) = f(t_3)\})\}
\end{array}
$$

Figure 7: Semantics of the temporal algebra

Then, the query is evaluated with the conventional semantics. Assume in the following that all the tables are tables with valid-time support.

$$[\![\texttt{NONSEQUENCED}\ \text{<query expression>}]\!]_{\text{SQL/T}}(r_1, \ldots, r_n) \triangleq [\![\text{<query expression>}]\!]_{\text{standard}}(SN(r_1), \ldots, SN(r_n))$$

Example 11: Let Q be the following non-temporal query.

```
SELECT p.X
FROM p, q
WHERE p.X = q.X
```

Assume that `p` is a table without temporal support and `q` is a valid-time table. To evaluate this query according to temporal upward compatibility, we timeslice `q` as of now.

$$[\![Q]\!]_{\text{SQL/T}}(p, q) = [\![Q]\!]_{\text{standard}}(p, \tau_{now}^{vt}(q)) = \pi_{p.X}(p \bowtie_{p.X=q.X} \tau_{now}^{vt}(q)) \ .$$

Now consider the SQL/Temporal query $Q_2 =$ `NONSEQUENCED VALID` Q. `q` is converted to a table without temporal support with an additional column, containing the valid-time period of the row, then the query is evaluated according to the standard SQL semantics.

$$
\begin{array}{lcl}
[\![Q_2]\!]_{\text{SQL/T}}(p, q) & = & [\![\texttt{NONSEQUENCED VALID}\ Q]\!]_{\text{SQL/T}}(p, q) \\
 & = & [\![Q]\!]_{\text{standard}}(p, SN(q)) \\
 & = & \pi_{p.X}(p \bowtie_{p.X=q.X} SN(q)) \ .
\end{array}
$$

Completing this example, let's examine sequenced valid semantics. Assume that `p` is also a table with valid-time support (sequenced semantics requires all underlying tables to be tables with valid-time support).

Consider $Q_3 =$ `VALID` Q.

$$\begin{aligned}[\![Q_3]\!]_{\mathrm{SQL/T}}(p,q) &= [\![\texttt{VALID } Q]\!]_{\mathrm{SQL/T}}(p,q) \\ &= [\![Q]\!]_{\mathrm{temporal}}(p,q) \\ &= \pi^{vt}_{p.X}(p \bowtie^{vt}_{p.X=q.X} q) \quad .\end{aligned}$$

Note that apart from the ^{vt}-superscripts, which are added to relational algebra operators, the translation between SQL queries and relational algebra expressions for all three types of queries has not changed at all. □

The semantics of `VALID(`c`)` are quite simple. If c is associated with a row $\langle t \mid\mid VT\rangle$ of a table with valid-time support, then

$$[\![\texttt{VALID}(c)]\!] = VT \quad .$$

If a period expression follows `VALID`, the result of the select statement is restricted to the period specified. This can be accomplished by an extension of the timeslice operator τ to take a period expression as its subscript.

$$\begin{aligned}&[\![\texttt{VALID } p \text{ <query expression>}]\!]_{\mathrm{SQL/T}}(r_1,\ldots,r_n) \\ &\quad \triangleq \{\langle t \mid\mid VT\rangle \mid \langle t \mid\mid VT'\rangle \in [\![\texttt{VALID } \text{<query expression>}]\!]_{\mathrm{SQL/T}}(r_1,\ldots,r_n) \wedge VT = VT' \cap [\![p]\!] \wedge VT \neq \emptyset\}\end{aligned}$$

If a period expression follows `NONSEQUENCED VALID`, a valid-time table results, with the valid-time period as specified. Assume in the following that all the tables are tables with valid-time support.

$$\begin{aligned}&[\![\texttt{NONSEQUENCED VALID } p \text{ <query expression>}]\!]_{\mathrm{SQL/T}}(r_1,\ldots,r_n) \\ &\qquad = \{\langle t \mid\mid [\![p]\!]\rangle \mid t \in ([\![\texttt{NONSEQUENCEDVALID } \text{<query expression>}]\!]_{\mathrm{SQL/T}}(r_1,\ldots,r_n)\}\end{aligned}$$

8 A Foundation for Implementing the Extensions

We first provide a mapping of temporal relational operations to conventional relational algebra expressions. We then list a range of alternatives for implementing the temporal relational operators.

8.1 Implementing the Temporal Algebra

Here, we give the conventional algebraic equivalents for the temporal algebraic operators. We emphasize that conventional operators range over a different domain (tables without temporal support) than do temporal operators (tables with valid-time support). In Figure 8, the set A_{r_i} contains the explicit attributes of table r_i, and a is the attribute appended by AG.

A "^{vt}" superscript on a table indicates that it is a table with valid-time support; those tables without such a superscript are tables without temporal support, each with an explicit VT column. The auxiliary function $SN(r) = \{\langle t, VT\rangle \mid \langle t \| VT\rangle \in r\}$ maps a table with valid-time support into a table without temporal support with valid-time being an explicit attribute. It is assumed that tables with valid-time support are mapped into tables without temporal support using SN before conventional algebraic operators are applied. Note that function SN is not needed at the implementation level. However, it is required here because Codd's relational algebra operators are only well-defined over tables without temporal support. Finally, there is a rename operator, $\rho_i(r)$ that gives table r the name i. Again, the aggregate operator is the most complex. The relational difference and the outer Cartesian product are the analogs of "$\neg\exists$" in the calculus; the inner Cartesian product and unions (to compute t_2) are the analogs of the two row variables in the calculus.

Example 12: We continue by mapping the temporal algebraic equivalent of the SQL/Temporal query $Q_2 =$ `VALID` Q_1 into the snapshot algebra. Here, we assume that the table `p` has a single column, `X`, and

$$
\begin{array}{lcl}
\sigma^{vt}_{c}(r^{vt}) & \rightsquigarrow & \sigma_{c}(r) \\
\pi^{vt}_{f}(r^{vt}) & \rightsquigarrow & \pi_{f,VT}(r) \\
r_1^{vt} \cup^{vt} r_2^{vt} & \rightsquigarrow & r_1 \cup r_2 \\
r_1^{vt} \bowtie^{vt}_{c} r_2^{vt} & \rightsquigarrow & \pi_{A_{r_1},VT_{r_1},A_{r_2},VT_{r_2},VT=intersect(VT_{r_1},VT_{r_2})}(r_1 \bowtie_{c \wedge VT_{r_1}\ overlaps\ VT_{r_2}} r_2) \\
r_1^{vt} \setminus^{vt} r_2^{vt} & \rightsquigarrow & t_2 \setminus \pi_{A_{t_2},VT_{t_2}}(t_2 \bowtie_{A_{t_2}=A_{r_2} \wedge VT_{t_2}\ overlaps\ VT_{r_2}} r_2) \\
 & & t_2 = t_1 \cup \pi_{A_{r_1},period(VT^{+}_{r_2},VT^{-}_{t_1})}(t_1 \bowtie_{A_{t_1}=A_{r_2} \wedge VT^{-}_{t_1} \leq VT^{+}_{r_2} \wedge VT^{+}_{r_2} < VT^{+}_{t_1}} r_2) \\
 & & t_1 = r_1 \cup \pi_{A_{r_1},period(VT^{-}_{r_1},VT^{-}_{r_2})}(r_1 \bowtie_{A_{r_1}=A_{r_2} \wedge VT^{-}_{r_1} < VT^{-}_{r_2} \wedge VT^{-}_{r_2} \leq VT^{+}_{r_1}} r_2) \\
AG^{vt}_{agg,f}(r^{vt}) & \rightsquigarrow & AG_{agg,f}(t_2 - \pi_{A_1,A_2,a,VT^{-}_{1},VT^{+}_{1}}(\\
 & & \qquad \sigma_{f(A_1)=f(A_2) \wedge ((VT^{-}_{1} < VT^{-}_{2} < VT^{+}_{1}) \vee (VT^{-}_{1} < VT^{+}_{2} < VT^{+}_{1}))}(\rho_1(t_2) \times \rho_2(r)))) \\
 & & t_2 = \pi_{A_1,a,period(VT^{-}_{1},VT^{-}_{2})}(t_1) \cup \pi_{A_1,a,period(VT^{-}_{1},VT^{+}_{2})}(t_1) \cup \pi_{A_1,a,period(VT^{+}_{1},VT^{+}_{2})}(t_1) \\
 & & t_1 = \pi_{A_1,a,VT_1,VT_2}(\sigma_{f(A_1)=f(A_2)}(\rho_1(r)) \times \rho_2(r))
\end{array}
$$

Figure 8: Snapshot equivalents of the the temporal algebra operators

that the table `q` has two columns, `X` and `Y`.

$$
\begin{array}{lcl}
[\![Q_2]\!]_{\mathrm{SQL/T}}(p,q) & = & [\![\texttt{VALID } Q_1]\!]_{\mathrm{SQL/T}}(p,q) = [\![Q_1]\!]_{\mathrm{temporal}}(p,q) = \pi^{vt}_{p.X}(p \bowtie^{vt}_{p.X=q.X} q) \\
 & = & \pi_{X_p,VT}(\pi_{X_p,VT_p,X_q,Y_q,VT_q,VT=intersect(VT_p,VT_q)}(SN(p) \bowtie_{p.X=q.X \wedge VT_p\ overlaps\ VT_q} SN(q))) \\
 & = & \pi_{X_p,VT=intersect(VT_p,VT_q)}(SN(p) \bowtie_{p.X=q.X \wedge VT_p\ overlaps\ VT_q} SN(q))
\end{array}
$$

□

8.2 Alternatives for Implementing SQL/Temporal

The transformations from the temporal algebra to the conventional algebra gives us several options for implementing SQL/Temporal.

1. Map temporal queries into temporal algebra, then into regular algebra, according to Figure 8, then back into SQL.
2. Map temporal queries into temporal algebra, then, according to Figure 7, directly into SQL.
3. Map temporal queries directly into SQL, utilizing the temporal algebra implicitly in the query rewrite phase (this is what the prototype does).

EXAMPLE 13: Continuing with the previous example, the SQL/Temporal query Q_2 on tables with valid-time support can be mapped to an SQL3 query on tables without temporal support where the implicit timestamps are now placed in explicit attributes.

```
SELECT p.X, INTERSECT(p.VT, q.VT) AS VT
FROM p, q
WHERE p.X = q.X AND p.VT OVERLAPS q.VT
```

Here, we use the `OVERLAPS` predicate and the `INTERSECT` operator already present in SQL/Temporal. □

Finally, we point out some possibilities for query optimization.

1. Map temporal queries into temporal algebra, optimize as with SQL algebra (with existing transformations and/or new cost formulas), then map back in SQL.
2. Map temporal queries into temporal algebra, then into SQL algebra, then optimize, then evaluate.
3. Map temporal queries into temporal algebra, then optimize (again, using new cost formulas), then evaluate the temporal algebra directly, with the concomitant increase in performance.

8.3 Implementing Temporal Assertions and Constraints

The general approach to checking an assertion is to negate it and then execute it as a query [2]. If the query result is empty, i.e., if no rows are returned, the assertion is respected, otherwise it is violated.

EXAMPLE 14: To check the assertion `emp_has_sal` from Section 6.1.2 we execute the query

```
SELECT *
FROM employee AS e
WHERE NOT EXISTS ( SELECT *
                   FROM salary AS s
                   WHERE e.eno = s.eno)
```

A non-empty result indicates a violation of the assertion. □

Temporal assertions and constraints, specified with `VALID`, can be checked in a similar way, with a `VALID SELECT` statement.

First, note that database systems have to improve the sketched mechanism to achieve acceptable performance. Well-known techniques include incremental consistency checking, simplification of assertions, and special-purpose checking algorithms for, e.g., column constraints. Second, it becomes obvious how important it is to address *all* aspects of a query language when transitioning from a nontemporal to a temporal database system. Negation, which might be used rarely in queries asked by users, is crucial for answering assertions because these usually involve some form of implication, i.e., involve negation. In our approach, it is no harder to state a temporal negation than it is to state a temporal join. This makes specification (and implementation) of assertions particularly elegant.

8.4 Implementing Nonsequenced Semantics

Interestingly, supporting nonsequenced semantics is much easier than supporting sequenced semantics.

Assume as before that a table with valid-time support is *represented* at the physical level as a conventional table with an additional, unnamed column containing the valid time. Then a nonsequenced query can be evaluated as before, with no special treatment. `VALID()` simply returns the value of the unnamed column. The optional period expression is treated as an additional unnamed column in the evaluation of a nonsequenced query.

A major advantage of nonsequenced semantics is its ease of implementation. The disadvantage of nonsequenced semantics is that the user is responsible for handling the time dimension explicitly, which is why *sequenced* semantics is so important.

9 Summary

In this change proposal, we first outlined several desirable features of SQL/Temporal relative to SQL3: upward compatibility, temporal upward compatibility, and sequenced semantics. A series of four levels of increasing functionality was elaborated. The specific syntactic additions were outlined and examples given to illustrate these constructs. The extensions involve (a) the use of the `VALID` reserved words, to indicate valid-time support (in the case of schema specification statements) and sequenced semantics (in the case of queries, modifications, views, assertions and constraints), (b) the use of the `NONSEQUENCED` reserved word for nonsequenced semantics, and (c) the use of a period expression to temporally scope sequenced and nonsequenced queries, modifications, views, cursors, constraints, and assertions. We provided a formal semantics, in terms of the formal semantics of SQL3, that satisfied the sequenced semantics correspondence between temporal queries and snapshot queries, and also provided the semantics for nonsequenced queries. Finally, we listed alternative implementation approaches which vary in the degree of implementation difficulty and the achievable performance efficiency, and showed that implementing nonsequenced semantics is straightforward.

Appendix A provides formal definitions of the properties discussed in Section 5.

We end by listing some of the advantages of the approach espoused here.

- Upward compatibility is assured, permitting existing constructs to operate exactly as before.
- Only two new reserved words, `NONSEQUENCED` and `VALID`, are required.
- Satisfaction of temporal upward compatibility ensures that existing applications do not break when tables without temporal support have such support added.
- Satisfaction of sequenced semantics ensures that temporal queries, modifications, views, assertions, and constraints are easy to specify, formalize, and implement.
- Nonsequenced semantics permits tables with valid-time support to be converted to tables without such support, with an explicit timestamp column, and such for valid-time support to be added to tables, even within a query.
- Since the semantics is defined in terms of the non-temporal semantics, the extensions are compatible with *all* the facilities of SQL3.
- A simple period expression permits the temporal scope to be specified.
- A prototype implementation exists [14]; this prototype was invaluable in refining the language additions.
- Transaction time support will require few syntactic or semantic extensions, and will be fully compatible and consistent with these valid-time features.

10 Proposed Language Extensions

The syntax is given as extensions to "Database Language SQL — Part 7: Temporal" [8].

11 Clause 3 Definitions, notations, and conventions

11.1 Subclause 3.1 Definitions

1) Add the following terms.

g) **valid time**: The valid time of a value is the time when the value is valid.

h) **row with valid-time support**: A row of a table with an associated valid time, which is a value of data type PERIOD.

i) **table with valid-time support**: A table with valid-time support is one in which each row is a row with valid-time support.

j) **state of a table with valid-time support at a valid time**: The state of a table with valid-time support, TV, at a specified valid time T, is the table without valid-time support comprising the rows of TV associated with valid-time periods that overlap T.

k) **current state of a table with valid-time support**: The current state of a table with valid-time support is the state of that table at valid time CURRENT_TIMESTAMP.

l) **temporal upward compatibility**: A derived table or cursor is said to be evaluated with temporal upward compatibility if its result is the same as it would be were each of its leaf generally underlying tables with valid-time support with no intervening <from clause> replaced with its current state.

An integrity constraint is said to be evaluated with temporal upward compatibility if its result is the same as it would be were each of its leaf generally underlying tables with valid-time support with no intervening <from clause> of every derived table in the integrity constraint replaced with its current state.

A delete statement or update statement is said to be evaluated with temporal upward compatibility on a table with valid-time support if it is effectively executed for the state of the table at every valid time equal to or following CURRENT_TIMESTAMP and with each of its leaf generally underlying tables with valid-time support with no intervening <from clause> of every derived table in the statement replaced with its current state.

m) **sequenced valid semantics**: A derived table or cursor is said to be evaluated with sequenced valid semantics if it is effectively evaluated for the state at every valid time of every leaf generally underlying table with valid-time support with no intervening <from clause>.

An integrity constraint is said to be evaluated with sequenced valid semantics if it is effectively evaluated for the state at every valid time of every leaf generally underlying table with valid-time support with no intervening <from clause>.

A delete statement or update statement is said to be evaluated with sequenced valid semantics if it is effectively executed at the state at every valid time of the table being modified and of every leaf generally underlying table with valid-time support of every derived table with no intervening <from clause> in the statement.

n) **nonsequenced valid semantics**: A derived table or cursor is said to be evaluated with nonsequenced valid semantics if its result is the same as it would be were each of its leaf generally underlying tables with valid-time support with no intervening <from clause> replaced with a table with no valid-time support with rows with identical values for the columns, except that there is an additional, unnamed column whose value for each row is the valid-time period of that row.

An integrity constraint is said to be evaluated with nonsequenced semantics if its result is the same as it would be were each of its leaf generally underlying tables with valid-time support with no intervening <from clause> of every derived table in the integrity constraint replaced with a table without valid-time support with rows with identical values for the columns, except that there is an additional, unnamed column whose value for each row is the valid-time period of that row.

o) **precision of a table with valid-time support**: The precision of a table with valid-time support is the precision of the valid-time period of its rows.

p) **forever**: Forever is the maximum datetime value, 9999-12-31 23:59:59.999999.

12 Clause 4 Concepts

1) Insert this new Subclause at the end of Clause 4, "Concepts".

12.1 Subclause 4.3 Semantics of Statements on Tables with Valid-Time Support

The semantics of the statements on tables with valid-time support is defined in terms of the semantics of statements on tables without valid-time support, which is already defined in this International Standard. In this way, the definitions exploit the other definitions and rules in this International Standard, without replication.

In this Clause the meaning is described only for queries, but the concepts apply to views, cursors, constraints, assertions, and modification statements.

Temporal upward compatible queries (i.e., SELECT without VALID) treat each underlying table that has valid-time support as a table without valid-time support, by using instead the current state of the table. Hence, a query evaluated with temporal upward compatibility on a table with valid-time support will use only the current state in the evaluation.

Sequenced valid semantics (e.g., VALID SELECT) is used only on tables with valid-time support, and queries with sequenced valid semantics result in tables with valid-time support. Sequenced valid semantics is defined in terms of the semantics of the statement on tables without valid-time support. Let Q be a sequenced valid query, with Q = VALID Q1, where Q1 is a query without VALID. The meaning of Q1 on tables without temporal support is already defined by this International Standard. Let R be the valid-time table that is the result of Q on one or more tables with valid-time support. For all times T, the state of R at time T is the result of evaluating Q1 on the states of the underlying tables at time T. Any R that satisfies this property is a valid result of Q.

Nonsequenced valid semantics (e.g., NONSEQUENCED VALID SELECT) treats each underlying table that has valid-time support as a table without valid-time support, but with an additional unnamed column whose value for a row in the table is the valid-time period associated with the corresponding row in the original table.

The <value expression> following VALID is used in two ways. Within a nonsequenced valid query, it supplies the valid-time period of the computed rows. Everywhere else, it specifies the "temporal scope" of the sequenced valid query, integrity constraint, or modification. For queries, the result is computed with sequenced valid semantics, then the valid-time periods of the result are intersected with the value of the <value expression> to determine the final valid-time period. For integrity constraints and modifications, the effect of the constraint or modification is restricted to the value of <value expression>.

Note to proposal reader: Very informally, what is going on is that the SQL3 semantics of a query Q is treated as a black box. Put a query and a database (conventional, without valid-time support) in one side, and out comes a (carefully specified) table (without valid-time support) on the other side.

The semantics of the nonsequenced, sequenced, and temporally upward compatible queries, modifications, etc., are specified by using this black box.

The advantages are numerous. (1) We don't have to modify each page of the specification; instead, the additions to the specification are small and isolated. (2) As SQL3 grows, with new constructs, the temporal query variants still work fine. (3) The intuition of the user is aided by the fact that a temporal version of a query is defined in terms of the original, nontemporal version of the query.
End of note.

1) Insert the following new Subclause immediately after Subclause 4.3, "Semantics of Statements on Tables with Valid-Time Support".

12.2 Section 4.4 Concepts used in Modification Statements

To remove a period RP from the valid-time period P of a row R of a table T, there are five cases.

a) If RP and P do not overlap, then do nothing.

b) If the beginning delimiting timestamp of UP is less than the beginning delimiting timestamp of P and the ending delimiting timestamp of UP is less than the ending delimiting timestamp of P, then replace the beginning delimiting timestamp of P for row R with the granule following the ending delimiting timestamp of RP.

c) If the beginning delimiting timestamp of UP is less than or equal to the beginning delimiting timestamp of P and the ending delimiting timestamp of UP is greater than or equal to the ending delimiting timestamp of P, then mark row R for deletion.

d) If the beginning delimiting timestamp of UP is greater than the beginning delimiting timestamp of P and the ending delimiting timestamp of UP is less than the ending delimiting timestamp of P, then replace the ending delimiting timestamp of P for row R with the granule preceding the beginning delimiting timestamp of UP. Let NP be a period of precision of that of P, with a beginning delimiting timestamp of the granule following the ending delimiting timestamp of RP, and an ending delimiting timestamp of that of the original value of P. The new rows NR, with column values identical to R, and with an associated valid-time period of NP, is inserted into T.

e) If the beginning delimiting timestamp of UP is greater than the beginning delimiting timestamp of P and less than the ending delimiting timestamp of P and the ending delimiting timestamp of UP is greater than the ending delimiting timestamp of P, then replace the beginning delimiting timestamp of P for row R with the granule preceding the beginning delimiting granule of RP.

To update the valid-time period P of a row R of a table T for a period UP, there are five cases.

a) If UP and P do not overlap, then do nothing.

b) If the beginning delimiting timestamp of UP is less than the beginning delimiting timestamp of P and the ending delimiting timestamp of UP is less than the ending delimiting timestamp of P, then replace the ending delimiting timestamp of P for row R with the granule following the ending delimiting timestamp of UP, and perform the update on R. Let NP be the period of precision of that of P with a beginning delimiting timestamp of that of the original value of P, and an ending delimiting timestamp of that of UP. A new row NR, with column values identical to R, and with an associated valid-time period of NP, is inserted into T. The update is performed only on NR.

c) If the beginning delimiting timestamp of UP is less than or equal to the beginning delimiting timestamp of P and the ending delimiting timestamp of UP is greater than or equal to the ending delimiting timestamp of P, then perform the update on R.

d) If the beginning delimiting timestamp of UP is greater than the beginning delimiting timestamp of P and the ending delimiting timestamp of UP is less than the ending delimiting timestamp of P, then replace the ending delimiting timestamp of P for row R with the preceding granule of the beginning delimiting timestamp of UP. Let NP1 and NP2 be periods of precision of that of P. Let NP1 have a beginning delimiting timestamp of that of UP, and an ending delimiting timestamp of the ending delimiting timestamp of UP. Let NP2 have a beginning delimiting timestamp of the granule following the ending delimiting timestamp of UP, and an ending delimiting timestamp of that of the original value of P. Two new rows, NR1 and NR2, with column values identical to R, and with an associated valid-time periods of NP1 and NP2, are inserted into T. The update is performed only on NR1.

e) If the beginning delimiting timestamp of UP is greater than the beginning delimiting timestamp of P and less than the ending delimiting timestamp of P and the ending delimiting timestamp of UP is greater than the ending delimiting timestamp of P, then replace the ending delimiting timestamp of P for row R with the granule preceding the beginning delimiting timestamp of UP. Let NP be the period of precision of that of P with a beginning delimiting timestamp of that of UP, and an ending delimiting timestamp of that of the original value of P. A new row NR, with column values identical to R, and with an associated valid-time period of NP, is inserted into T. The update is performed only on NR.

Within a sequenced valid semantics, for a given row R a period P satisfies a <search condition> if the states of the row at all granules of P satisfy the <search condition>.

13 Clause 5 Lexical elements

13.1 Subclause 5.1 <token> and <separator>

Function

Specify lexical units (tokens and separators) that participate in SQL language.

Format

1) In the Format, add the following BNF production:

```
<reserved word> ::=
        !! All alternatives from part 2 of ISO/EIC 9075
      | NONSEQUENCED
      | VALID
```

Syntax Rules

No additional Syntax Rules.

Access Rules

No additional Access Rules.

General Rules

No additional General Rules.

14 Clause 6 Scalar expressions

14.1 Subclause 6.1 <item reference>

1) Insert this Subclass, "<item reference>" immediately preceding Subclause 6.2, "<set function specification>".

Function

Reference a column, parameter, or variable.

Format

No additional Format items.

Syntax Rules

1. (Insert this SR) If a <column name> CN is contained in a <valid expression> simply contained in a <query expression> QE, then the scope clause SC of CN is the <query expression body> simply contained in QE.

Access Rules

No additional Access Rules.

General Rules

No additional General Rules.

14.2 Subclause 6.4 <period value function>

Function

Specify a function yielding a value of type period.

Format

1) In the Format, add the following two BNF productions:

```
<period value function> ::=
          !! All alternatives from part 2 of ISO/EIC 9075
        | <valid function>

<valid function> ::=
          VALID <left paren> [ <table name> | <correlation name> ]
                <right paren>
```

Syntax Rules

1. (Insert this SR) The <table name> or <correlation name> of a <valid function> shall be associated with a table with valid-time support.

2. (Insert this SR) The data type of <valid function> shall be PERIOD, with a precision of that of the table associated with the <table name> or <correlation name>.

Note to proposal reader: The precision of a table with valid-time support was defined in Clause 3.1 "Definitions".

Access Rules

No additional Access Rules.

General Rules

1. (Insert this GR) The <valid function> evaluates to the valid-time period of a given row of the table associated with the <table name> or <correlation name>.

15 Clause 7 Query Processing

15.1 Subclause 7.7 <query expression>

1) Insert this new Subclause immediately after Subclause 7.6, "<meets predicate>".

Function

Specify a table.

Format

1) In the Format, replace the <query expression> BNF production with:

<query expression> ::=
[<with clause>] <temporal query expression body>

2) Add the following BNF production:

<temporal query expression body> ::=
[<time option>] <query expression body>

Note to proposal reader: This adds an optional <time option> to <query expression>.

3) In the Format, add the following two BNF productions:

<time option> ::=
<valid option>

<valid option> ::=
[NONSEQUENCED] VALID [<value expression>]

Note to proposal reader: the <time option> is expanded in the transaction-time change proposal, in a similar fashion to <valid option> proposed here. There are five cases:

1. `SELECT`
   - works on anything
   - evaluates to a table with no temporal support
2. `VALID SELECT`
   - works only on tables with valid-time support
   - evaluates to a table with valid-time support
3. `VALID` <period exp> `SELECT`
   - like `VALID SELECT`, but only returns timestamps within <period exp> (a simple example is `VALID PERIOD 1995 SELECT`)
4. `NONSEQUENCED VALID SELECT`
   - works on anything
   - acts like tables with valid-time support have an explicit (unnamed) timestamp column
   - evaluates to a table with no temporal support

5. `NONSEQUENCED VALID` <period exp> `SELECT`

   - like `NONSEQUENCED VALID SELECT`, but uses the <period exp> as a timestamp, and thus returns a table with valid-time support

To convert a table with valid-time support to a table with no temporal support, use `SELECT` (if only the current state is of interest) or `NONSEQUENCED VALID SELECT`.

To convert a table without valid-time support to a table with valid-time support, use `NONSEQUENCED VALID` <period exp> `SELECT`.
End of note.

Syntax Rules

1. Replace <query expression body> with <temporal query expression body> in SR1b.

2. (Insert this SR) If VALID without NONSEQUENCED is specified in <query expression>, then each exposed table, query, or correlation name contained in the <query expression body> without an intervening <from clause> shall identify a table with valid-time support and with identical precision P.

   Note to proposal reader: This ensures that sequenced valid queries are only evaluated "over" tables with valid-time support.

3. (Insert this SR) If VALID is specified in the <time option> of a <query expression> Q, then either Q shall be simply contained in a <from clause> or Q shall not be contained in a second query expression.

   Note to proposal reader: VALID is allowed in only two places: prepended to the outermost <query expression>, or immediately within a <from clause>. The reason is that the <time option> applies to the entire <query expression>, evaluated on the exposed tables, queries, and correlation names, which are specified in the <from clause>s of the query.

4. (Insert this SR) The data type of the <value expression> contained in <time option> shall be PERIOD. Subclause 6.3 "<item reference>" restricts the scope of column names in <value expression>.

5. (Insert this SR) Let T be the result of the <query expression>.

   Case:

   a) If VALID without NONSEQUENCED is specified in <time option>, then T shall be a table with valid-time support and with a precision of P. The precision of the <value expression> contained in the <valid option> of <time option> shall be P.

   b) If NONSEQUENCED VALID is specified in <time option>, then

      Case:

      i) If <value expression> is specified in the <valid option> of <time option>, then T shall be a table with valid-time support and with a precision of that of <value expression>.

      ii) Otherwise, T shall be a table without valid-time support.

   c) Otherwise, T shall be a table without valid-time support.

      Note to proposal reader: If NONSEQUENCED VALID is specified, the precision of the <value expression> of <valid option> is arbitrary.

Access Rules

No additional Access Rules.

General Rules

1. (Insert this GR) Case:

   a) If VALID without NONSEQUENCED is specified in <time option>, then <query expression body> is evaluated with sequenced valid semantics. If <value expression> is specified in the <valid option> of <time option>, then for each resulting row, the valid-time period of the row is the intersection of the value of <value expression> and the original valid-time period of the row.

   b) If NONSEQUENCED VALID is specified in <time option>, then <query expression body> is evaluated with nonsequenced valid semantics. If <value expression> is specified in the <valid option> of <time option>, then the valid-time period of each row has the value of <value expression>.

   c) Otherwise, the <query expression body> is evaluated with temporal upward compatibility.

16 Clause 9 Schema definition and manipulation

16.1 Subclause 9.2 <table definition>

1) Insert this new Subclause immediately after Subclause 9.1, "<default clause>".

Function

Define a persistent base table, a created local temporary table, or a global temporary table.

Format

1) In the Format, replace the <table definition> BNF production with:

<table definition> ::=
CREATE [<table scope>] [EXTENT] TABLE <table name>
[<table element list> | <subtable clause>]
[<temporal definition>]
[ON COMMIT <table commit action> ROWS]

Note to proposal reader: This augments the production for the non-terminal <table definition> with an additional, optional clause to specify that the new table is to be a table with valid-time support.

2) In the Format, add the following two BNF productions.

<temporal definition> ::=
AS VALID [<table precision>]

<table precision> ::=
<left paren> <period precision> <right paren>

Syntax Rules

No additional Syntax Rules.

Access Rules

No additional Access Rules.

General Rules

1. (Insert this GR) Case:

   a) If VALID is specified in <temporal definition>, then the table is a table with valid-time support with a precision of <table precision>.

   b) Otherwise, the table does not have valid-time support.

16.2 Subclause 9.3 <column definition>

1) Insert this new Subclause immediately after Subclause 9.2, "<table definition>".

Function

Define a column of a table.

Format

1) In the Format, replace the <column constraint definition> BNF production with:

<column constraint definition> ::=
[<constraint name definition>]
[<time option>]
<column constraint> [<constraint attributes>]

Note to proposal reader: This adds an optional <time option> to column constraints.

Syntax Rules

1. (Insert this SR) If VALID is specified in <time option>, then T shall be a table with valid-time support.
2. (Insert this SR) If VALID without NONSEQUENCED is specified,
   Case:
   a) If <column constraint> is <references specification>, then the table identified by <table name> simply contained in the <referenced table and columns> of <references specification> shall be a table with valid-time support.
   b) If <column constraint> is <check constraint definition>, then each table associated with an exposed <table name>, <query expression>, or <correlation name> contained in the <column constraint> without an intervening <from clause> shall be a table with valid-time support and with identical precision.
3. (Insert this SR) The precision of <value expression> contained in the <valid option> of <time option> shall be the precision of T. The scope for <valid expression> shall be the columns of T.
4. (Insert this SR) If NONSEQUENCED is specified in the <time option> TO contained in <column constraint definition> , then TO shall not contain <value expression>.

Access Rules

No additional Access Rules.

General Rules

1. (Insert this GR) Case:
   a) If VALID without NONSEQUENCED is specified in <time option>, then <column constraint> is evaluated with sequenced valid semantics. If <value expression> is contained in the <valid option> of <time option>, then the <column constraint> of a <column constraint definition> applies only to those states at valid times overlapping the value of <value expression>.
   b) If NONSEQUENCED is specified in <time option>, then <column constraint> is evaluated with nonsequenced valid semantics.
   c) Otherwise, <column constraint> is evaluated with temporal upward compatibility.

16.3 Subclause 9.4 <table constraint definition>

1) Insert this new Subclause immediately after Subclause 9.3, "<column definition>".

Function

Specify an integrity constraint.

Format

1) In the Format, replace the <table constraint definition> BNF production with:

<table constraint definition> ::=
[<constraint name definition>]
[<time option>]
<table constraint> [<constraint attributes>]

Note to proposal reader: This adds an optional <time option>. For constraints and assertions, there are four cases:

1. `CHECK`
   - works on anything
   - only considers current state
2. `VALID CHECK`
   - works only on tables with valid-time support
   - the assertion must be true for the state at every valid time
3. `VALID` <period exp> `CHECK`
   - like `VALID CHECK`, but only considers the times in <period exp> (a simple example is `VALID PERIOD '[1995-01-01 - 1995-12-31]' CHECK`)
4. `NONSEQUENCED VALID CHECK`
   - works on anything
   - acts like tables with valid-time support have an explicit (unnamed) timestamp column; all rows are considered at once

`NONSEQUENCED VALID` <period exp> `CHECK` is not allowed.
End of note.

Syntax Rules

1. (Insert this SR) Let T be the table defined by the <table definition> containing this <table constraint definition>.
2. (Insert this SR) If VALID is specified in <time option>, then T shall be a table with valid-time support. The precision of <value expression> contained in the <value expression> of <time option> shall be the precision of T.
3. (Insert this SR) If VALID without NONSEQUENCED is specified in <table constraint definition>, then each exposed table, query, or correlation name contained in the <table constraint> without an intervening <from clause> shall identify a table with valid-time support and with identical valid-time period precision.

4. (Insert this SR) If <time option> contained in <column constraint definition> contains NONSEQUENCED, then it shall not contain <value expression>.

Access Rules

No additional Access Rules.

General Rules

1. (Insert this GR) Case:

   a) If VALID without NONSEQUENCED is specified in <time option>, then <table constraint> is evaluated with sequenced valid semantics. If <value expression> is contained in the <valid option> of <time option>, then the <table constraint> of a <column constraint definition> only applies to those states at valid times overlapping the value of <value expression>.

   b) If NONSEQUENCED VALID is specified in <time option>, then <table constraint> is evaluated with nonsequenced valid semantics.

   c) Otherwise, <table constraint> is evaluated with temporal upward compatibility.

16.4 Subclause 9.5 <alter table statement>

1) Insert this new Subclause immediately after Subclause 9.4, "<table constraint definition>".

Function

Change the definition of a table.

Format

1) In the Format, add the following BNF production:

```
<alter table action> ::=
        !! All alternatives from ISO/EIC 9075
      | <add valid definition>
      | <drop valid definition>
      | <convert valid definition>
```

2) In the Format, add the following three BNF productions:

```
<add valid definition> ::=
        ADD VALID [ <table precision> ]

<drop valid definition> ::=
        DROP VALID

<convert valid definition> ::=
        CAST VALID AS <table precision>
```

Syntax Rules

1. (Insert this SR) Let T be the table identified by the <table name> immediately contained in <alter table statement>.
2. (Insert this SR) For the <add valid definition>, T shall be a table without valid-time support.
3. (Insert this SR) For the <drop valid definition>, T shall be a table with valid-time support.
4. (Insert this SR) For the <convert valid definition>, T shall be a table with valid-time support.

Access Rules

No additional Access Rules.

General Rules

1. (Insert this GR) If <add valid definition> is specified, then valid-time support is added to each row of T, by associating with that row a valid-time period from the current timestamp to forever with a precision of <table precision>.
2. (Insert this GR) If <drop valid definition> is specified, then valid-time support is removed from T, by replacing T with the result of

   SELECT * FROM T

   Note to proposal reader: That is, only the current state is retained.

3. (Insert this GR) If <convert valid definition> is specified, T is converted to the new precision specified, via the following statements. Let TC be the columns of T, and let P be the <period precision> of <table precision>.

```
CREATE TABLE Temp ( TC ) AS VALID P

INSERT INTO Temp
VALID SELECT ( TC )
FROM (NONSEQUENCED VALID CAST( validtime AS PERIOD(TIMESTAMP P ))
        SELECT *, VALID(T) AS validtime
        FROM T) AS S

DROP TABLE T

CREATE TABLE T ( TC ) AS VALID P

INSERT INTO T
        VALID SELECT * FROM Temp

DROP TABLE Temp
```

16.5 Subclause 9.6 <assertion definition>

1) Insert this new Subclause immediately after Subclause 9.5, "<alter table statement>".

Function

Specify an integrity constraint by means of an assertion and specify when the assertion is to be checked.

Format

1) In the Format, replace the <triggered assertion> BNF production with:

<triggered assertion> ::=
[<time option>]
CHECK <left paren> <search condition> <right paren>

Note to proposal reader: This adds an optional <time option>.

Syntax Rules

1. (Insert this SR) If VALID without NONSEQUENCED is specified in <time option>, then each exposed table, query, or correlation name contained in the <search condition> without an intervening <from clause> shall identify a table with valid-time support and with identical valid-time period precision P. The precision of the <value expression> contained in <valid option> of <time option> shall be P.
2. (Insert this SR) If NONSEQUENCED is specified in the <time option> TO of <triggered assertion>, then the <valid option> of TO shall not contain <value expression>.

Access Rules

No additional Access Rules.

General Rules

1. (Insert this GR) Case:
   a) If VALID without NONSEQUENCED is specified in <time option>, then <search condition> is evaluated with sequenced valid semantics. If <value expression> is contained in the <valid option> of <time option>, then the <triggered assertion> of a <column constraint definition> only applies to those states at times overlapping the value of <value expression>
   b) If NONSEQUENCED VALID is specified in <time option>, then <search condition> is evaluated with nonsequenced valid semantics.
   c) Otherwise, <search condition> is evaluated with temporal upward compatibility.

17 Clause 11 Data manipulation

1) Insert this new Clause immediately before Clause 10, "Information Schema and Definition Schema".

17.1 Subclause 11.1 <fetch statement>

1) Insert this new Subclause at the beginning of Clause, "Data manipulation".

Function

Position a cursor on a specified row of a table and retrieve values from that row.

Format

1) In the Format, replace the <fetch statement> BNF production with:

<fetch statement> ::=
FETCH [[<fetch orientation>] FROM] <cursor name>
[INTO <fetch target list>]
[INTO VALID <target specification>]

Note to proposal reader: This extends the <fetch statement> to also allow the valid-time period of the row to be accessed.

Syntax Rules

1. (Insert this SR) <fetch statement> shall contain "INTO <fetch target list>", "INTO VALID <target specification>", or both.
2. (Insert this SR) If INTO VALID is specified in the <fetch statement>, then T shall have valid-time support. The data type of the <target specification> of <fetch statement> shall be PERIOD with a precision of that of T.

Access Rules

No additional Access Rules.

General Rules

1. (Insert this GR) If INTO VALID is specified, then the valid-time period associated with the current row is assigned to the target of the <target specification>.

17.2 Subclause 11.2 <select statement: single row>

1) Insert this new Subclause immediately after Subclause 11.1, "<fetch statement>".

Function

Retrieve values from a specified row of a table.

<select statement: single row> ::=
[<time option>]
SELECT [<set quantifier>] <select list>
INTO <select target list>
<table expression>

Note to proposal reader: This adds an optional <time option>.

Syntax Rules

1. (Insert this SR) If VALID without NONSEQUENCED is specified in <time option>, then each exposed <table name>, <query expression>, or <correlation name> contained in the <table expression> without an intervening <from clause> shall identify a table with valid-time support and with identical precisions P.

2. (Insert this SR) If VALID is specified in a <time option> of a <query expression> Q that is contained in the <table expression> of <select statement: single row>, then Q shall be simply contained in a <from clause>.

3. (Insert this SR) Let T be the result of <select statement: single row>.

   Case:

   a) If VALID without NONSEQUENCED is specified in <time option>, then T shall be a table with valid-time support and with precision P. The precision of <value expression> of the <valid option> of <time option> shall be P.

   b) If NONSEQUENCED VALID is specified in <time option>, then Case:

      i) If <value expression> is specified in the <valid option> of <time option>, then T shall be a table with valid-time support and with a precision of <value expression>.

      ii) Otherwise, T shall be a table without valid-time support.

   c) Otherwise, T shall be a table without valid-time support.

4. (Insert this SR) The scope for the <valid expression> contained in the <time option> of <select statement: single row> shall be the columns of the <select list>.

   Note to proposal reader: If NONSEQUENCED is specified, then the precision of the <value expression> contained in the <valid option> of <time option> is arbitrary.

Access Rules

No additional Access Rules.

General Rules

1. (Insert this GR) Case:

   a) If VALID without NONSEQUENCED is specified in <time option>, then the <query expression body> is evaluated with sequenced valid semantics. If <value expression> is specified in the <valid option> of <time option>, then for each resulting row, the valid-time period of the row is the intersection of the valid of <value expression> and the original valid-time period of the row.

   b) If NONSEQUENCED VALID is specified in <time option>, then the <query expression body> is evaluated with nonsequenced valid semantics. If <value expression> is specified in the <valid option> of <time option>, then the valid-time period of each row has the value of <value expression>.

   c) Otherwise, the <query expression body> is evaluated with temporal upward compatibility.

17.3 Subclause 11.3 <delete statement: positioned>

1) Insert this new Subclause immediately after Subclause 11.2, "<select statement: single row>".

Function

Delete a row of a table.

Format

1) In the Format, replace the <delete statement: positioned> BNF production with:

<delete statement: positioned> ::=
[<time option>]
DELETE [FROM <table reference>]
WHERE CURRENT OF <cursor name>

Note to proposal reader: This augments the production for <delete statement: positioned> with an additional, optional <time option> clause. For deletions and updates, there are four cases:

1. `DELETE`
   - works on anything
   - when applied to a table with valid-time support, deletes for times from now to forever
2. `VALID DELETE`
   - works only on tables with valid-time support
   - applies to the state at each time
3. `VALID` <period exp> `DELETE`
   - like `VALID DELETE`, but only considers the times in <period exp> (a simple example is `VALID PERIOD '[1995-01-01 - 1995-12-31]' DELETE`)
4. `NONSEQUENCED VALID DELETE`
   - works only on tables with valid-time support
   - acts like the table has an explicit (unnamed) timestamp column (in the <search condition>, if present)
5. `NONSEQUENCED VALID` <period exp> `DELETE`
   - like `NONSEQUENCED VALID DELETE`, but only considers the times in <period exp>

End of note.

Syntax Rules

1. (Insert this SR) Let T be the subject table of the <delete statement: positioned>.
2. (Insert this SR) If VALID is specified in <time option>, then T shall be a table with valid-time support.
3. (Insert this SR) The precision of <value expression> contained in the <valid option> of <time option> shall be the precision of T. The scope for the <valid expression> contained in the <time option> of <delete statement:positioned> shall be the columns of T.

Access Rules

No additional Access Rules.

General Rules

1. (Insert this GR) Case:

   a) If VALID is specified in <time option>, then
      Case:

      i) If a <value expression> is specified in the <valid option> of <time option>, then the value of <value expression> is removed from the valid-time period of the row.

      ii) Otherwise, the row is marked for deletion.

   b) Otherwise,
      Case:

      i) If T is a table with valid-time support, then the period from the current timestamp to forever is removed from the valid-time period of the row.

      ii) Otherwise, the row is marked for deletion.

17.4 Subclause 11.4 <delete statement: searched>

1) Insert this new Subclause immediately after Subclause 11.3, "<delete statement: positioned>".

Function

Delete rows of a table.

Format

1) In the Format, replace the <delete statement: searched> BNF production with:

```
<delete statement: searched> ::=
        [ <time option> ]
        DELETE FROM <table reference>
            [ WHERE <search condition> ]
```

Note to proposal reader: This augments the production for <delete statement: searched> with an additional, optional clause.

Syntax Rules

1. (Insert this SR) Let T be the subject table of the <delete statement: searched>.
2. (Insert this SR) If VALID is specified in <time option>, then T shall be a table with valid-time support.
3. (Insert this SR) If VALID is specified in a <time option> of a <query expression> Q that is contained in the <search condition> of <delete statement:searched>, then Q shall be simply contained in a <from clause>.
4. (Insert this SR) The precision of <value expression> contained in the <valid option> of <time option> shall be the precision of T. The scope for the <valid expression> contained in the <time option> of <delete statement:searched> shall be the columns of T.
5. (Insert this SR) If VALID without NONSEQUENCED is specified in <time option>, then each exposed <table name>, <query expression>, or <correlation name> contained in the <search condition> without an intervening <from clause> shall identify a table with valid-time support and with identical precisions P.

Access Rules

No additional Access Rules.

General Rules

1. (Insert this GR) Case:
    a) If VALID without NONSEQUENCED is specified in <time option>, then the <search condition> is evaluated with sequenced valid semantics.

       Case:

       i) If <value expression> is specified in the <valid option> of <time option>, then let PS be the set of those periods for which the <search condition> is satisfied. For each period DP in PS, the intersection of DP and the value of <value expression> is removed from the valid-time period of the row, as well as from the valid-time period of new rows that were inserted in the processing of previous periods from PS for this row.

ii) Otherwise, the valid time associated with each period satisfying the <search condition> is removed from the valid-time period of the row, as well as from the valid-time period of new rows that were inserted in the processing of previous periods for this row.

b) If NONSEQUENCED VALID is specified in <time option>, then the <search condition> is evaluated with nonsequenced valid semantics.

Case:

i) If a <value expression> is specified in <time option>, and if <search condition> is satisfied, then the value of <value expression> is removed from the valid-time period of the row.

ii) Otherwise, the row is marked for deletion.

c) Otherwise, the <search condition> is evaluated with temporal upward compatibility. If the <search condition> is satisfied for the relevant row and T is a table with valid-time support, the period from the current timestamp to forever is removed from the valid-time period of the row.

17.5 Subclause 11.5 <insert statement>

1) Insert this new Subclause immediately after Subclause 11.4, "<delete statement: searched>".

Function

Create new rows in a table.

Format

No additional Format items.

Syntax Rules

1. (Insert this SR) T is the subject table of the <insert statement>.
2. (Insert this SR) If T is a table with valid-time support, then <insert columns and source> shall evaluate to a table with valid-time support and with a precision identical to that of T.

Access Rules

No additional Access Rules.

General Rules

No additional General Rules.

17.6 Subclause 11.6 <update statement: positioned>

1) Insert this new Subclause immediately after Subclause 11.5, "<insert statement>".

Function

Update a row of a table.

Format

1) In the Format, replace the <update statement: positioned> BNF production with:

```
<update statement: positioned> ::=
        [ <time option> ]
        UPDATE [ <table reference> ]
              SET <set clause list>
              WHERE CURRENT OF <cursor name>
```

Note to proposal reader: This adds an optional <time option>.

Syntax Rules

1. (Insert this SR) Let T be the subject table of the <update statement: positioned>.
2. (Insert this SR) If VALID is specified in <time option>, then T shall be a table with valid-time support.
3. (Insert this SR) The precision of <value expression> contained in the <valid option> of <time option> shall be the precision of T. The scope for the <valid expression> contained in the <time option> of <update statement:positioned> shall be the columns of T.

Access Rules

No additional Access Rules.

General Rules

1. (Insert this GR) Case:

   a) If VALID is specified in <time option>,
      Case:
      i) If <value expression> is specified, then the valid-time period of the row is updated for the value of the <value expression>.
      ii) Otherwise, the update is performed on the row.

   b) Otherwise,
      Case:
      i) If T is a table with valid-time support, then the valid-time period of the row is updated for the period from the current timestamp to forever.
      ii) Otherwise, the update is performed on the row.

17.7 Subclause 11.7 <update statement: searched>

1) Insert this new Subclause immediately after Subclause 11.6, "<update statement:positioned>".

Function

Update rows of a table.

Format

1) In the Format, replace the <update statement: searched> BNF production with:

```
<update statement: searched> ::=
        [ <time option> ]
        UPDATE <table reference>
            SET <set clause list>
            [ WHERE <search condition> ]
```

Note to proposal reader: This adds an optional <time option>.

Syntax Rules

1. (Insert this SR) Let T be the subject table of the <update statement: searched>.
2. (Insert this SR) If VALID is specified in <time option>, then T shall be a table with valid-time support.
3. (Insert this SR) If VALID is specified in a <time option> of a <query expression> Q that is contained in the <search condition> of <update statement:searched>, then Q shall be simply contained in a <from clause>.
4. (Insert this SR) The precision of <value expression> contained in the <valid option> of <time option> shall be the precision of T. The scope for the <valid expression> contained in the <time option> of <update statement: searched> shall be the columns of T.
5. (Insert this SR) If VALID without NONSEQUENCED is specified in <time option>, then each exposed <table name>, <query expression>, or <correlation name> contained in the <search condition> without an intervening <from clause> shall identify a table with valid-time support and with identical precisions P.

Access Rules

No additional Access Rules.

General Rules

1. (Insert this GR) Case:
    a) If VALID is specified without NONSEQUENCED in <time option>, then the <search condition> is evaluated with sequenced valid semantics.

       Case:

        i) If <value expression> is specified in the <valid option> of <time option>, then let PS be the set of those periods for which the <search condition> is satisfied. For each period DP in PS, the valid-time period of the row, as well as of new rows that were inserted in the processing of previous periods from PS for this row, are updated for the intersection of DP and the value of <value expression>.

ii) Otherwise, the valid-time period of the row, as well as of new rows that were inserted in the processing of previous periods for this row, are updated for each period satisfying the <search condition>.

b) If NONSEQUENCED VALID is specified in <time option>, then the <search condition> is evaluated with nonsequenced valid semantics.

Case:

i) If <value expression> is specified and if the <search condition> is satisfied, then the valid-time period of the row is updated for the value of the <value expression>.

ii) Otherwise, if the <search condition> is satisfied, then the update is performed on the row.

c) Otherwise,

Case:

i) If T is a table with valid-time support, then the <search condition> is evaluated with temporal upward compatibility. If the <search condition> is satisfied, then the valid-time period of the row is updated for the period from the current timestamp to forever.

ii) Otherwise, if the <search condition> is satisfied, then the update is performed on the row.

18 Clause 12 Information Schema and Definition Schema

18.1 Subclause 12 Information Schema

18.1.1 Subclause 12.1.1 TABLES view

1) Insert this new Table immediately preceding Subclause 10.2, "Definition Schema".

Function

Identify the tables defined in this catalog that are accessible to a given user.

Definition

1. Replace the TABLES view with the following.

```
CREATE VIEW TABLES
         AS SELECT
               TABLE_CATALOG, TABLE_SCHEMA, TABLE_NAME, TABLE_TYPE,
               VALID_TIME_SUPPORT, VALID_TIME_PRECISION,
         FROM DEFINITION_SCHEMA.TABLES
         WHERE ( TABLE_CATALOG, TABLE_SCHEMA, TABLE_NAME )
         IN (
               SELECT TABLE_CATALOG, TABLE_SCHEMA, TABLE_NAME
               FROM DEFINITION_SCHEMA.TABLE_PRIVILEGES
               WHERE GRANTEE IN ( 'PUBLIC', CURRENT_USER )
              UNION
               SELECT TABLE_CATALOG, TABLE_SCHEMA, TABLE_NAME
               FROM DEFINITION_SCHEMA.COLUMN_PRIVILEGES
               WHERE GRANTEE IN ( 'PUBLIC', CURRENT_USER ) )
               AND TABLE_CATALOG
                = ( SELECT CATALOG_NAME FROM INFORMATION_SCHEMA_CATALOG_NAME )
```

Note to proposal reader: This adds two columns: VALID_TIME_SUPPORT and VALID_TIME_PRECISION.

Leveling Rules

No additional Leveling Rules.

18.2 Subclause 12.2 Definition Schema

18.2.1 Subclause 12.2.2 TABLES base table

1) Insert this new Table immediately following Subclause 10.2.1, "DATA_TYPE_DESCRIPTOR base table".

Function

The TABLES table contains one row for each table including views. It effectively contains a representation of the table descriptors.

Definition

1. Replace the TABLES table with the following.

```
CREATE TABLE TABLES
         (
         TABLE_CATALOG INFORMATION_SCHEMA.SQL_IDENTIFIER,
         TABLE_SCHEMA  INFORMATION_SCHEMA.SQL_IDENTIFIER,
         TABLE_NAME   INFORMATION_SCHEMA.SQL_IDENTIFIER,
         TABLE_TYPE   INFORMATION_SCHEMA.CHARACTER_DATA,
               CONSTRAINT TABLE_TYPE_NOT_NULL NOT NULL,
         VALID_TIME_SUPPORT CHARACTER_DATA
               CONSTRAINT VALID_TIME_SUPPORT_CHECK
         CHECK (VALID_TIME_SUPPORT IN ('STATE','NONE')),
         VALID_TIME_PRECISION  INFORMATION_SCHEMA.CARDINAL_NUMBER,
               CONSTRAINT TABLE_TYPE_CHECK CHECK ( TABLE_TYPE IN
              ( 'BASE TABLE', 'VIEW', 'GLOBAL TEMPORARY',
                'LOCAL TEMPORARY', 'EXTENT' ) ),
        ------------------------------------------------------------
         ISO Only-caused by ANSI changes not yet considered by ISO
         ------------------------------------------------------------
               CONSTRAINT TABLE_TYPE_CHECK CHECK ( TABLE_TYPE IN
              ( 'BASE TABLE', 'VIEW', 'GLOBAL TEMPORARY', 'LOCAL TEMPORARY' ) ),

         ------------------------------------------------------------

               CONSTRAINT CHECK_TABLE_IN_COLUMNS
         CHECK ( ( TABLE_CATALOG, TABLE_SCHEMA, TABLE_NAME ) IN
               ( SELECT TABLE_CATALOG, TABLE_SCHEMA, TABLE_NAME
         FROM COLUMNS ) ),

         CONSTRAINT TABLES_PRIMARY_KEY
               PRIMARY KEY ( TABLE_CATALOG, TABLE_SCHEMA, TABLE_NAME ),

         CONSTRAINT TABLES_FOREIGN_KEY_SCHEMATA
               FOREIGN KEY ( TABLE_CATALOG, TABLE_SCHEMA ) REFERENCES SCHEMATA,

         CONSTRAINT TABLES_CHECK_NOT_VIEW CHECK ( NOT EXISTS
               ( SELECT TABLE_CATALOG, TABLE_SCHEMA, TABLE_NAME
         FROM TABLES
         WHERE TABLE_TYPE = 'VIEW'
         EXCEPT
         SELECT TABLE_CATALOG, TABLE_SCHEMA, TABLE_NAME
         FROM VIEWS ) )
```

)

Note to proposal reader: This adds two columns: VALID_TIME_SUPPORT and VALID_TIME_PRECISION.

Leveling Rules

No additional Leveling Rules.

18.2.2 Subclause 12.2.3 VIEWS base table

1) Insert this new Table immediately following Subclause 12.2.2, "TABLES base table".

Function

The VIEWS table contains one row for each row in the TABLES table with a TABLE_TYPE of 'VIEW'. Each row describes the query expression that defines a view. The table effectively contains a representation of the view descriptors.

Definition

1. Replace the VIEWS table with the following.

```
CREATE TABLE VIEWS
         (
         TABLE_CATALOG   INFORMATION_SCHEMA.SQL_IDENTIFIER,
         TABLE_SCHEMA INFORMATION_SCHEMA.SQL_IDENTIFIER,
         TABLE_NAME     INFORMATION_SCHEMA.SQL_IDENTIFIER,
         VIEW_DEFINITION INFORMATION_SCHEMA.CHARACTER_DATA,
         CHECK_OPTION INFORMATION_SCHEMA.CHARACTER_DATA
               CONSTRAINT CHECK_OPTION_NOT_NULL NOT NULL
               CONSTRAINT CHECK_OPTION_CHECK CHECK ( CHECK_OPTION IN
         ( 'CASCADED', 'LOCAL', 'NONE' ) ),
         IS_UPDATABLE    INFORMATION_SCHEMA.CHARACTER_DATA
               CONSTRAINT IS_UPDATABLE_NOT_NULL NOT NULL
               CONSTRAINT IS_UPDATABLE_CHECK CHECK ( IS_UPDATABLE IN ( 'YES', 'NO' ) ),
         VALID_TIME_SUPPORT CHARACTER_DATA
       CONSTRAINT VALID_TIME_SUPPORT_CHECK CHECK (VALID_TIME_SUPPORT IN ('STATE','NONE'))
         VALID_TIME_PRECISION  INFORMATION_SCHEMA.CARDINAL_NUMBER,
         CONSTRAINT VIEWS_PRIMARY_KEY
               PRIMARY KEY ( TABLE_CATALOG, TABLE_SCHEMA, TABLE_NAME ),

         CONSTRAINT VIEWS_IN_TABLES_CHECK
               CHECK ( ( TABLE_CATALOG, TABLE_SCHEMA, TABLE_NAME ) IN
               ( SELECT TABLE_CATALOG, TABLE_SCHEMA, TABLE_NAME
         FROM TABLES
         WHERE TABLE_TYPE = 'VIEW' ) ),

         CONSTRAINT VIEWS_IS_UPDATABLE_CHECK_OPTION_CHECK
               CHECK ( ( IS_UPDATABLE, CHECK_OPTION ) NOT IN
              ( VALUES ( 'NO', 'CASCADED' ), ( 'NO', 'LOCAL' ) ) )
         )
```

Note to proposal reader: This adds two columns: VALID_TIME_SUPPORT and VALID_TIME_PRECISION.

Leveling Rules

No additional Leveling Rules.

18.2.3 Subclause 12.2.4 TABLE_CONSTRAINTS base table

1) Insert this new Table immediately following Subclause 12.2.3, "VIEWS base table".

Function

The TABLE_CONSTRAINTS table has one row for each table constraint associated with a table. It effectively contains a representation of the table constraint descriptors.

Definition

1. Replace the TABLE_CONSTRAINTS table with the following.

```
CREATE TABLE TABLE_CONSTRAINTS
         (
         CONSTRAINT_CATALOG   INFORMATION_SCHEMA.SQL_IDENTIFIER,
         CONSTRAINT_SCHEMA     INFORMATION_SCHEMA.SQL_IDENTIFIER,
         CONSTRAINT_NAME         INFORMATION_SCHEMA.SQL_IDENTIFIER,
         CONSTRAINT_TYPE         INFORMATION_SCHEMA.CHARACTER_DATA
         CONSTRAINT CONSTRAINT_TYPE_NOT_NULL NOT NULL
               CONSTRAINT CONSTRAINT_TYPE_CHECK

         CHECK ( CONSTRAINT_TYPE IN
                ( 'UNIQUE',
                  'PRIMARY KEY',
                  'FOREIGN KEY',
                  'CHECK' ) ),

         TABLE_CATALOGINFORMATION_SCHEMA.SQL_IDENTIFIER
               CONSTRAINT TABLE_CONSTRAINTS_TABLE_CATALOG_NOT_NULL NOT NULL,
         TABLE_SCHEMA           INFORMATION_SCHEMA.SQL_IDENTIFIER
               CONSTRAINT TABLE_CONSTRAINTS_TABLE_SCHEMA_NOT_NULL NOT NULL,
         TABLE_NAME  INFORMATION_SCHEMA.SQL_IDENTIFIER
               CONSTRAINT TABLE_CONSTRAINTS_TABLE_NAME_NOT_NULL NOT NULL,
         IS_DEFERRABLE INFORMATION_SCHEMA.CHARACTER_DATA
               CONSTRAINT TABLE_CONSTRAINTS_IS_DEFERRABLE_NOT_NULL NOT NULL,
         INITIALLY_DEFERRED   INFORMATION_SCHEMA.CHARACTER_DATA
               CONSTRAINT TABLE_CONSTRAINTS_INITIALLY_DEFERRED_NOT_NULL
         VALID_TIME_SUPPORT CHARACTER_DATA
               CONSTRAINT VALID_TIME_SUPPORT_CHECK
         CHECK (VALID_TIME_SUPPORT IN ('SEQUENCED','NONSEQUENCED','NONE')),
         VALID_TIME_PERIOD  INFORMATION_SCHEMA.CARDINAL_NUMBER,

         CONSTRAINT TABLE_CONSTRAINTS_PRIMARY_KEY
               PRIMARY KEY ( CONSTRAINT_CATALOG, CONSTRAINT_SCHEMA, CONSTRAINT_NAME ),

         CONSTRAINT TABLE_CONSTRAINTS_DEFERRED_CHECK
               CHECK ( ( IS_DEFERRABLE, INITIALLY_DEFERRED ) IN
              ( VALUES ( 'NO',  'NO' ),
              ( 'YES', 'NO' ),
              ( 'YES', 'YES' ) ) ),

         CONSTRAINT TABLE_CONSTRAINTS_CHECK_VIEWS
               CHECK ( TABLE_CATALOG
              <> ANY ( SELECT CATALOG_NAME FROM SCHEMATA )
```

```
    OR
      ( ( TABLE_CATALOG, TABLE_SCHEMA, TABLE_NAME ) IN
( SELECT TABLE_CATALOG, TABLE_SCHEMA, TABLE_NAME
 FROM TABLES
 WHERE TABLE_TYPE <> 'VIEW' ) ) ),

   CONSTRAINT TABLE_CONSTRAINTS_UNIQUE_CHECK
         CHECK ( 1 =( SELECT COUNT (*)
FROM ( SELECT CONSTRAINT_CATALOG, CONSTRAINT_SCHEMA,
            CONSTRAINT_NAME FROM TABLE_CONSTRAINTS
WHERE CONSTRAINT_TYPE IN
          ( 'UNIQUE', 'PRIMARY KEY' )
UNION ALL
SELECT CONSTRAINT_CATALOG, CONSTRAINT_SCHEMA, CONSTRAINT_NAME
FROM REFERENTIAL_CONSTRAINTS
UNION ALL
SELECT CONSTRAINT_CATALOG, CONSTRAINT_SCHEMA, CONSTRAINT_NAME
FROM CHECK_CONSTRAINTS ) AS X
WHERE ( CONSTRAINT_CATALOG, CONSTRAINT_SCHEMA, CONSTRAINT_NAME )
        = ( X.CONSTRAINT_CATALOG, X.CONSTRAINT_SCHEMA, X.CONSTRAINT_NAME ) ) ),

   CONSTRAINT UNIQUE_TABLE_PRIMARY_KEY_CHECK
         CHECK ( UNIQUE ( SELECT TABLE_CATALOG, TABLE_SCHEMA, TABLE_NAME
        FROM TABLE_CONSTRAINTS
        WHERE CONSTRAINT_TYPE = 'PRIMARY KEY' ) )
   )
```

Note to proposal reader: This adds two columns: VALID_TIME_SUPPORT and VALID_TIME_PERIOD.

Leveling Rules

No additional Leveling Rules.

18.2.4 Subclause 12.2.5 CHECK_CONSTRAINTS base table

1) Insert this new Table immediately following Subclause 12.2.4, "TABLE_CONSTRAINTS base table".

Function

The CHECK_CONSTRAINTS table has one row for each domain constraint, table check constraint, and assertion.

Definition

1. Replace the CHECK_CONSTRAINTS table with the following.

```
CREATE TABLE CHECK_CONSTRAINTS
         (
         CONSTRAINT_CATALOG INFORMATION_SCHEMA.SQL_IDENTIFIER,
         CONSTRAINT_SCHEMA  INFORMATION_SCHEMA.SQL_IDENTIFIER,
         CONSTRAINT_NAME      INFORMATION_SCHEMA.SQL_IDENTIFIER,
         CHECK_CLAUSE    INFORMATION_SCHEMA.CHARACTER_DATA,
         VALID_TIME_SUPPORT CHARACTER_DATA
         CONSTRAINT VALID_TIME_SUPPORT_CHECK
                CHECK (VALID_TIME_SUPPORT IN ('SEQUENCED','NONSEQUENCED,'NONE'))
         VALID_TIME_PERIOD  INFORMATION_SCHEMA.CARDINAL_NUMBER,

         CONSTRAINT CHECK_CONSTRAINTS_PRIMARY_KEY
               PRIMARY KEY ( CONSTRAINT_CATALOG, CONSTRAINT_SCHEMA, CONSTRAINT_NAME ),

         CONSTRAINT CHECK_CONSTRAINTS_SOURCE_CHECK
               CHECK ( ( CONSTRAINT_CATALOG, CONSTRAINT_SCHEMA, CONSTRAINT_NAME )
     IN
       ( SELECT * FROM (
               SELECT CONSTRAINT_CATALOG, CONSTRAINT_SCHEMA, CONSTRAINT_NAME
               FROM ASSERTIONS
               UNION
               SELECT CONSTRAINT_CATALOG, CONSTRAINT_SCHEMA, CONSTRAINT_NAME
               FROM TABLE_CONSTRAINTS
               UNION
               SELECT CONSTRAINT_CATALOG, CONSTRAINT_SCHEMA, CONSTRAINT_NAME
               FROM DOMAIN_CONSTRAINTS ) ) )
         )
```

Note to proposal reader: This adds two columns: VALID_TIME_SUPPORT and VALID_TIME_PERIOD.

Leveling Rules

No additional Leveling Rules.

18.2.5 Subclause 12.2.6 ASSERTIONS base table

1) Insert this new Table immediately following Subclause 12.2.5, "CHECK_CONSTRAINTS base table".

Function

The ASSERTIONS table has one row for each assertion. It effectively contains a representation of the assertion descriptors.

Definition

1. Replace the TABLE_ASSERTIONS table with the following.

```
CREATE TABLE ASSERTIONS
         (
         CONSTRAINT_CATALOG   INFORMATION_SCHEMA.SQL_IDENTIFIER,
         CONSTRAINT_SCHEMA     INFORMATION_SCHEMA.SQL_IDENTIFIER,
         CONSTRAINT_NAME         INFORMATION_SCHEMA.SQL_IDENTIFIER,
         IS_DEFERRABLEINFORMATION_SCHEMA.CHARACTER_DATA
               CONSTRAINT ASSERTIONS_IS_DEFERRABLE_NOT_NULL NOT NULL
         INITIALLY_DEFERRED   INFORMATION_SCHEMA.CHARACTER_DATA
               CONSTRAINT ASSERTIONS_INITIALLY_DEFERRED_NOT_NULL NOT NULL
         CHECK_TIME  INFORMATION_SCHEMA.CHARACTER_DATA
               CONSTRAINT ASSERTIONS_CHECK_TIME_CHECK
         CHECK ( CHECK_TIME IN ('IMMEDIATE', 'DEFERRED' ) ),
         VALID_TIME_SUPPORT CHARACTER_DATA
               CONSTRAINT VALID_TIME_SUPPORT_CHECK
          CHECK (VALID_TIME_SUPPORT IN ('SEQUENCED','NONSEQUENCED','NONE')),
         VALID_TIME_PERIOD  INFORMATION_SCHEMA.CARDINAL_NUMBER,

         CONSTRAINT ASSERTIONS_PRIMARY_KEY
               PRIMARY KEY ( CONSTRAINT_CATALOG, CONSTRAINT_
         SCHEMA, CONSTRAINT_NAME ),

         CONSTRAINT ASSERTIONS_FOREIGN_KEY_CHECK_CONSTRAINTS
               FOREIGN KEY (CONSTRAINT_CATALOG, CONSTRAINT_
         SCHEMA, CONSTRAINT_NAME ) REFERENCES CHECK_CONSTRAINTS,

         CONSTRAINT ASSERTIONS_FOREIGN_KEY_SCHEMATA
               FOREIGN KEY ( CONSTRAINT_CATALOG, CONSTRAINT_SCHEMA )
          REFERENCES SCHEMATA,

         CONSTRAINT ASSERTIONS_DEFERRED_CHECK
               CHECK ( ( IS_DEFERRABLE, INITIALLY_DEFERRED ) IN
            VALUES ( ( 'NO',  'NO' ),
                     ( 'YES', 'NO' ),
                     ( 'YES', 'YES' ) ) )
          )
```

Note to proposal reader: This adds two columns: VALID_TIME_SUPPORT and VALID_TIME_PERIOD.

19 Acknowledgements

This change proposal was written by the four authors listed on the title page. The first author was supported in part by NSF grant ISI-9202244 and by grants from IBM, the AT&T Foundation, and DuPont. The second and third authors were supported in part by the Danish Natural Science Research Council, grant 9400911. In addition, the third author was supported by grants 11–1089–1 and 11–0061–1, also provided by the Danish Natural Science Research Council. The document was produced in part during visits by the first author to Aalborg University and by the second author to the University of Arizona.

This change proposal presents an improved and extended version of some of the constructs in TSQL2, which was designed by a committee consisting of Richard T. Snodgrass (chair), Ilsoo Ahn, Gad Ariav, Don S. Batory, James Clifford, Curtis E. Dyreson, Ramez Elmasri, Fabio Grandi, Christian S. Jensen, Wolfgang Käfer, Nick Kline, Krishna Kulkanri, T.Y. Cliff Leung, Nikos Lorentzos, John F. Roddick, Arie Segev, Michael D. Soo and Suryanarayana M. Sripada. Their participation in the TSQL2 design was critical.

We thank Curtis Dyreson for helpful comments, and Hugh Darwen and Mike Sykes for their suggested changes to the proposed definitions in Section 11.1. Jim Melton provided extensive help on all of the proposed language extensions; the authors greatly appreciate the time Jim took to explain the intracacies of writing change proposals.

A Formal Definition of Compatibility Properties

We have adopted the convention that a data model consists of two components, namely a set of data structures and a language for querying the data structures [15]. 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 queries in *M* that may be formulated on some database in *DS*. We use *db* to denote a database and *q* to denote a query.

A.1 Upward Compatibility

One data model is syntactically upward compatible with another data model if all data structures and legal query expressions of the latter model are contained in the former model.

DEFINITION 1: (syntactical upward compatibility) 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
- $\forall q_2 \in QL_2\ (q_2 \in QL_1)$. □

When transitioning from one system to a new system, it is important that the new data model contains the existing data model. If that is the case, all existing application code will remain syntactically correct.

For a query expression q and an associated database db, both legal elements of QL and DS of data model $M = (DS, QL)$, define $\langle\langle q(db) \rangle\rangle_M$ as the result of evaluating q on 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. In addition to the previous syntactical requirement, we add the requirement that all queries expressible in the existing model must evaluate to the same results in the existing and new models.

DEFINITION 2: (upward compatibility) 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
- $\forall db_2 \in DS_2\ (\forall q_2 \in QL_2\ (\langle\langle q_2(db_2) \rangle\rangle_{M_2} = \langle\langle q_2(db_2) \rangle\rangle_{M_1}))$. □

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_2, to a new system, with data model M_1.

A.2 Temporal Upward Compatibility

Intuitively, the requirement is that a query q will return the same result on an associated snapshot database db as on the counterpart of the database with valid-time support, $\mathcal{T}(db)$. Further, modifications should not affect this. The precise definitions given next is explained in the following.

DEFINITION 3: (temporal upward compatibility) 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 support of a table without temporal support to the table with valid-time support with the same explicit attributes. Next, let $u_1, u_2, \ldots, u_n$ 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 and
- $\forall db_S \in DS_S\ (\forall q_S \in QL_S\ (\langle\langle q_S(u_n(u_{n-1}(\ldots(u_1(db_S)\ldots)))) \rangle\rangle_{M_S} =$
  $(\langle\langle q_S(u_n(u_{n-1}(\ldots(u_1(\mathcal{T}(db_S))))))) \rangle\rangle_{M_T})))$. □

Assume that, when moving to the new system, some of the existing (snapshot) tables are transformed into tables with valid-time support, using `ALTER`, without changing the existing set of (explicit) attributes. This transformation is denoted by $\mathcal{T}$ in the definition. Then the same sequence of modification statements, denoted by the u_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.

A.3 Sequenced Valid Semantics

We first define the notion of sequenced valid semantics among query languages. We use r and r^{vt} for denoting an instance without and with valid-time support, respectively. Similarly, db and db^{vt} are sets of table instances without and with valid-time support, respectively.

The definition uses a valid-timeslice operator $\tau_c^{M^v,M}$ (e.g., [9, 1]) which takes as arguments a table with valid-time support r^{vt} (in the data model M^{vt}) and a valid-time granule c and returns a table without temporal support r (in the data model M) containing all rows valid at time c. In other words, r consists of all rows of r^{vt} whose valid time includes the time granule c, but without the valid time. This operator was already introduced in Section 7.1; here, we simply have emphasized the models involved by using them as superscripts.

DEFINITION 4: (sequenced valid semantics) Let $M = (DS, QL)$ be a snapshot relational data model, and let $M^{vt} = (DS^{vt}, QL^{vt})$ be a valid-time data model. Data model M^{vt} is *sequenced valid with respect to* data model M if

$$\forall q \in QL\ (\exists q^{vt} \in QL^{vt}\ (\forall db^{vt} \in DS^{vt}\ (\forall c\ (\tau_c^{M^v,M}(q^{vt}(db^{vt})) = q(\tau_c^{M^v,M}(db^{vt})))))).$$ □

Graphically, sequenced valid semantics implies that for all query expressions q in the snapshot model, there must exist a query q^{vt} in the temporal model so that for all db^{vt} and for all c, the commutativity diagram shown in Figure 9 holds.

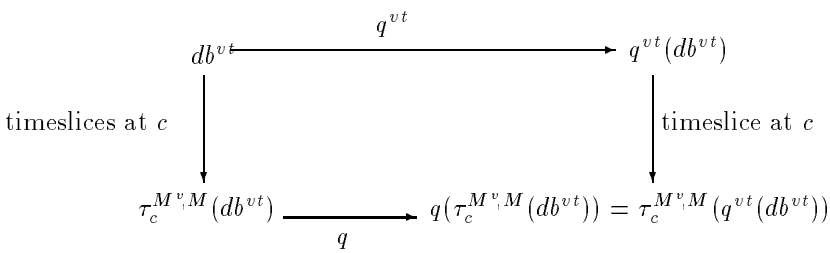

Figure 9: sequenced valid semantics of query q^{vt} with respect to query q at a chronon c

We require that each query q in the snapshot model has a counterpart q^{vt} in the temporal model that is sequenced valid with respect to it. Observe that q^{vt} being sequenced valid with respect to q poses no syntactical restrictions on q^{vt}. It is thus possible for q^{vt} to be quite different from q, and q^{vt} 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 require that q^{vt} and q be syntactically identical.

DEFINITION 5: (syntactically identical sequenced-valid extension) Let $M = (DS, QL)$ be a snapshot data model, and let $M^{vt} = (DS^{vt}, QL^{vt})$ be a valid-time data model. Data model M^{vt} is a *syntactically identical sequenced-valid extension* of model M if both of the following conditions hold.

1. Data model M^{vt} is sequenced valid with respect to data model M and

2. Each query in QL^{vt} that is sequenced valid with respect to a query in QL is syntactically identical to that query. □

If the valid-time data model treats tables with valid-time support as such, it is possible to use the same syntactical constructs (i.e., q^{vt} and q are identical) for querying tables with and without valid-time support. In this case, the type of support determines the meaning of the syntactical construct.

However, the identity property is incompatible with also requiring temporal upward compatibility. This latter property requires that a query from the snapshot model, when applied to a database with valid-time support, returns a table without temporal support. The property just defined requires the snapshot query to return a table with valid-time support when evaluated on the database with valid-time support.

Thus, not both of these properties can be satisfied by a temporal data model and the snapshot model it generalizes.

Our solution is to slightly relax the identity requirement, leading to the property defined below. With that property satisfied, the temporal queries may still exploit the programmers' intuition about the snapshot query language as much as possible.

DEFINITION 6: (syntactically similar sequenced-valid extension) Let $M = (DS, QL)$ be a snapshot data model, and let $M^{vt} = (DS^{vt}, QL^{vt})$ be a valid-time data model. Data model M^{vt} is a *syntactically similar sequenced-valid extension* of model M if both of the following conditions hold.

1. Data model M^{vt} is sequenced valid with respect to data model M and
2. For each query q^{vt} in QL^{vt} that is sequenced valid with respect to a query q in QL, $q^{vt} = S_1 q S_2$, where S_1 and S_2 are text strings that depend on QL^{vt} but not on q^{vt}. □

This property is consistent with temporal upward compatibility; the language designer simply has to select at least one of S_1 or S_2 as being non-empty. For the addition to SQL/Temporal proposed here, S_1 is simply "`VALID`", and S_2 is the empty string.

A.4 Properties of SQL/Temporal

We have developed a formal denotational semantics for SQL/Temporal, in terms of the semantics of SQL3. This semantics allowed us to prove the following important properties.

- SQL/Temporal is upward compatible with SQL3.
- SQL/Temporal is temporally upward compatible with SQL3.
- The `VALID` reserved word prepended to the `SELECT` statement ensures (syntactically similar) sequenced valid semantics [1, 3].
- SQL/Temporal is temporally ungrouped [4].