Temporal Data Management

Christian S. Jensen, *Senior Member, IEEE*, and Richard T. Snodgrass, *Senior Member, IEEE*

Abstract—A wide range of database applications manage time-varying information. Existing database technology currently provides little support for managing such data. The research area of temporal databases has made important contributions in characterizing the semantics of such information and in providing expressive and efficient means to model, store, and query temporal data. This paper introduces the reader to temporal data management, surveys state-of-the-art solutions to challenging aspects of temporal data management, and points to research directions.

Index Terms—Query language, SQL, temporal database, temporal data model, time-constrained database, transaction time, TSQL2, user-defined time, valid time.

———————————— ✦ ————————————

1 INTRODUCTION

A *TEMPORAL DATABASE* records time-varying information. Most database applications are temporal in nature, e.g., financial applications such as portfolio management, accounting, and banking, record-keeping applications such as personnel, medical-record, and inventory management, scheduling applications such as airline, train, and hotel reservations and project management, and scientific applications such as weather monitoring.

The study of temporal databases is a vibrant research topic, with an active community of several hundred researchers who have produced some 2,000 papers over the last two decades. These papers are listed in a series of seven cumulative bibliographies (the last [23] provides pointers to the previous ones). The field has produced a comprehensive glossary of terminology [8], a book-length survey providing a snapshot circa 1993 [19], and three workshop proceedings [2], [3], [15]. The nascent SQL3 draft standard now includes Part 7, SQL/Temporal [11]. The topic of temporal databases is now included in textbooks and an encyclopedia [20].

The present paper examines a variety of central areas of temporal database research. Each area is first motivated, and then sample contributions are surveyed, to give the reader a feel for the type of challenges and issues that are faced in each particular area. None of the existing surveys provides a short entree into the field as presented by this paper. The paper concludes with the authors' outlook into the possible future of temporal database research.

Given the space limitation, we cannot survey all areas, let alone all contributions, and the presentation must be brief. Thus, we have omitted a wide range of contributions that we consider important. A recent survey [12] and the slightly older book on temporal databases [19] go into more depth. We have also found it useful to focus on relational databases. The relational model is well-known, and its simplicity is conducive to maintaining an emphasis on the temporal essence of past research.

2 ONTOLOGICAL FOUNDATIONS

Before we proceed to consider temporal data models and query languages, we examine in data model-independent terms the association of times and facts, which is at the core of temporal data management.

Initially, a brief description of terminology is in order. A database models and records information about a part of reality, termed the *miniworld*. Aspects of the miniworld are represented in the database by a variety of structures that we will simply term database entities. We will employ the term "fact" for any statement that can meaningfully be assigned a truth value, i.e., that is either true or false. In general, times are associated with database entities.

Our focus will be on the facts that databases record. Several different temporal aspects have been associated with these. Most importantly, the *valid time* of a fact is the collected times—possibly spanning the past, present, and future—when the fact is true in the miniworld [8]. Valid time thus captures the time-varying states of the miniworld. By definition, all facts have a valid time. However, the valid time may not necessarily be recorded in the database, for any of a number of reasons. For example, the valid time may not be known, or recording it may not be relevant for the applications supported by the database. If a database models different possible worlds, the database facts may have several valid times, one for each such world.

Next, the *transaction time* of a database fact is the time when the fact is current in the database. Unlike valid time, transaction time may be associated with any database entity, not only with facts. For example, transaction may be associated with objects and values that are not facts because they cannot be true or false in isolation. Thus, all database entities have a transaction-time aspect. This aspect may or may not, at the database designer's discretion, be captured in the database. The transaction-time aspect of a database entity has a duration: from insertion to (logical) deletion. Transaction time captures the time-varying states of the

---
- *C.S. Jensen is with the Department of Computer Science, Institute for Electronic Ssystems, Aalborg University, Fredrik Bajers Vej 7E, DK–9220, Aalborg Ø, Denmark. E-mail: csj@cs.auc.dk.*
- *R.T. Snodgrass is with the Department of Computer Science, University of Arizona, Tucson, AZ 85721-0077. E-mail: rts@cs.arizona.edu.*

database, and applications that demand accountability or traceability rely on databases that record transaction time.

Observe that the transaction time of a database fact, say "*F*," is the valid time of the related fact, "*F is current in the database*." This would indicate that supporting transaction time as a separate aspect is redundant. However, both valid and transaction time are aspects of the contents of all databases, and recording both of these is essential in a wide range of applications. In addition, transaction time, due to its special semantics, is particularly well-behaved and may be supplied automatically by the DBMS. Specifically, the transaction time of facts stored in the database marches monotonically forward, and is bounded by the time the database was created at one end and by the current time at the other end. This provides the rationale for the focus of most temporal database research on providing improved support for valid time and transaction time as separate aspects.

In addition, some other times have been considered, e.g., decision time. But the desirability of building decision time support into temporal database technologies is unclear, because the number and meaning of "the decision times" of a fact varies from application to application and because decision times, unlike transaction time, generally do not exhibit specialized properties.

Much research has been conducted on the semantics and representation of time, from quite theoretical topics such as temporal logic and infinite periodic time sequences to rather applied questions such as how to represent time values in minimal space and how to utilize calendars. Also, there is a large body of research on time data types, e.g., time points, time intervals (or "periods"), and temporal elements (sets of intervals).

3 TEMPORAL DATA MODELS

Temporal data management can be very difficult using conventional (nontemporal) data models and query languages. Accommodating the time-varying nature of the enterprise is largely left to the developers of database applications, leading to ineffective and inefficient ad hoc solutions that must be reinvented each time a new application is developed. The result is that data management is currently an excessively involved and error-prone activity.

The first step providing support for temporal data management is to extend the database structures of the data model supported by the DBMS. More specifically, means must be given for capturing the valid and transaction times of the facts recorded by the relations, leading to temporal relations.

Subsequent steps are to provide support for temporal data modeling and database design, and to design temporal query languages that operate on the databases of the temporal data models. These topics are covered in Sections 4 and 5, respectively.

Adding time to the relational model, then, has been a daunting task, and more than two dozen extended relational data models have been proposed [9]. Most of these models support valid time only; some also support transaction time. We will consider three of these latter models and related design issues.

As a simple example, consider a video store where customers, identified by the `CustomerID` attribute, rent video tapes, identified by the `TapeNum` attribute. We consider a few rentals during May 1997. On the Second of May, customer C101 rents tape T1234 for three days. The tape is subsequently returned on the Fifth. Also on the Fifth, customer C102 rents tape T1245 with an open-ended return date. The tape is eventually returned on the Eighth. On the Ninth, customer C102 rents tape T1234 to be returned on the 12th. On the 10th, the rental period is extended to include the 13th, but this tape is not returned until the 16th. The video store keeps a record of these rentals in a relation termed `CheckedOut`.

Fig. 1 illustrates a relation instance in the Bitemporal Conceptual Data Model (BCDM) [9] that describes the sample rental scenario. This data model timestamps tuples, corresponding to facts, with values that are sets of (*transaction time*, *valid time*)-pairs, captured using attribute `T` in the figure. Fig. 2 provides a graphical illustration of the three

CustomerID	TapeNum	T
C101	T1234	$\{(2,2),(2,3),(2,4),(3,2),(3,3),(3,4),$ $\ldots,(UC,2),(UC,3)(UC,4)\}$
C102	T1245	$\{(5,5),(6,5),(6,6),(7,5),(7,6),(7,7),$ $(8,5),(8,6),(8,7),\ldots,$ $(UC,5),(UC,6),(UC,7)\}$
C102	T1234	$\{(9,9),(9,10),(9,11),$ $(10,9),(10,10),(10,11),(10,12),$ $(10,13),\ldots,$ $(13,9),(13,10),(13,11),(13,12),$ $(13,13),(14,9),\ldots,(14,14),$ $(15,9),\ldots,(15,15),(16,9),\ldots,$ $(16,15),\ldots,(UC,9),\ldots,(UC,15)\}$

Fig. 1. Bitemporal conceptual `CheckedOut` instance.

Fig. 2. Bitemporal diagram of the CheckedOut instance.

CustomerID	TapeNum	T_s	T_e	V_s	V_e
C101	T1234	2	UC	2	4
C102	T1245	5	7	5	now
C102	T1245	8	UC	5	7
C102	T1234	9	9	9	11
C102	T1234	10	13	9	13
C102	T1234	14	15	9	now
C102	T1234	16	UC	9	15

CustomerID		TapeNum	
$[2,Now] \times [2,4]$	C101	$[2,Now] \times [2,4]$	T1234
$[5,7] \times [5,\infty]$	C102	$[5,7] \times [5,\infty]$	T1245
$[8,Now] \times [5,7]$		$[8,Now] \times [5,7]$	
$[9,9] \times [9,11]$		$[9,9] \times [9,11]$	T1234
$[10,13] \times [9,13]$		$[10,13] \times [9,13]$	
$[14,15] \times [9,\infty]$		$[14,15] \times [9,\infty]$	
$[16,Now] \times [9,15]$		$[16,Now] \times [9,15]$	

Fig. 3. Alternative representations of the CheckedOut instance.

timestamp values, which are termed bitemporal elements. In the general case of infinite and continuous time domains, these are finite unions of rectangles in the two-dimensional space spanned by valid and transaction time.

The presence of a pair (*tt, vt*) in a timestamp of a tuple means that the current state of the database at time *tt* records that the fact represented by the tuple is valid at time *vt*. The special value *UC* ("until changed") serves as a marker indicating that its associated facts remain part of the current database state, and the presence of this value results in new time pairs being included into the sets of pairs at each clock tick.

The timestamp of the second tuple is explained as follows. On the Fifth, it is believed that customer C102 has checked out tape T1245 on the Fifth. Then, on the Sixth, the rental period is believed to include the Fifth and the Sixth. On the Seventh, the rental period extends to also include the Seventh. From then on, the rental period remains fixed. The current time is the 17th, and as time passes, the region grows to the right; the arrows indicate this and correspond to the *UC* values in the textual representation.

The idea behind the BCDM is to retain the simplicity of the relational model while also capturing the temporal aspects of the facts stored in a database. Because no two tuples with mutually identical explicit attribute values (termed *value-equivalent*) are allowed in a BCDM relation instance, the full history of a fact is contained in exactly one tuple. In addition, BCDM relation instances that are syntactically different have different information content, and vice versa. This conceptual cleanliness is generally not obtained by other bitemporal models where syntactically different instances may record the same information.

However, when it comes to the internal representation and the display to users of temporal information, the BCDM falls short. Although it is arguably a first-normal-form relation, the varying length and voluminous timestamps of

tuples are impractical to manage directly, and the timestamp values are also hard to comprehend in the BCDM format. Better suited representations of temporal information exist for these purposes.

Fig. 3 illustrates the same temporal information as in Fig. 1, in two different data models. The model exemplified at the top uses a practical and popular (particularly when implementation is considered) fixed-length format for tuples. Attributes T_s and T_e record starting and ending transaction times, and V_s and V_e record starting and ending valid times. In this format, each tuple's timestamp then encodes a rectangular or stair-shaped bitemporal region, and it may take several such tuples to represent a single fact.

The relation format at the bottom in Fig. 3 is a typical non-first-normal-form representation. In this format, a relation is thought of as recording information about some type of objects. The present relation records information about customers and thus holds one tuple for each customer in the example, with a tuple containing all information about a customer. In this way, a single tuple records multiple facts. For example, the second tuple records two facts: rental information for customer C102 for the two tapes T1245 and T1234.

Unlike in the BCDM where relations must be updated at every clock tick, relations in the two other formats stay up-to-date; this is achieved by introducing variables (e.g., *now*) as database values that assume the (changing) current time value. It should be noted that all of the three types of bitemporal relations are equally expressive in that they may record the same facts. Put more formally (and briefly), the relation instances that these models may record are snapshot equivalent.

Finally, it should be noted that the sample relations illustrate the two predominant choices for where to enter time values into relations, namely at the level of tuples (tuple timestamping) and at the level of attribute values ("attribute" timestamping).

4 DESIGNING TEMPORAL DATABASES

Database design is typically considered in two contexts. In conceptual design, a database is modeled using a high-level design model that is independent of the particular (implementation) data model of the DBMS that is eventually to be used for managing the database. The second context of database design is the implementation data model, which is assumed to conform to the ANSI/X3/SPARC three-level architecture. In this context, database design must thus be considered at the view, logical, and physical (or, "internal") levels. We proceed to consider conceptual and logical design of temporal databases.

4.1 Conceptual Design

By far, most research on conceptual design of temporal databases has been in the context of the Entity-Relationship (ER) model. This model, in its varying forms, is enjoying a remarkable, and increasing, popularity in industry. Building on the example introduced in Section 3, Fig. 4 illustrates a conventional ER diagram for video rentals.

The research on temporal ER modeling is well motivated. It is widely known that the temporal aspects of the miniworld are very important in a broad range of applications, but are also difficult to capture using the ER model. Put simply, when attempting to capture the temporal aspects, these tend to obscure and clutter otherwise intuitive and easy-to-comprehend diagrams.

The diagram in the figure is nontemporal, capturing the miniworld at a single point in time. Attempting to capture the temporal aspects that are essential for this application clutters up the simple diagram. For example, since the same customer may check out the same tape at different times, the `CustomerID` and `TapeNum` attributes do not identify a single instance of `CheckedOut`. Instead, it is necessary to make `CheckedOut` a ternary relationship type, with the third entity type capturing start dates of rentals. There is also the issue of where to place the end-time attribute of rentals. Next, rental prices may vary over time, e.g., due to

promotions and films getting old. Finally, including transaction time further complicates matters.

The research community's response has been to develop temporally enhanced ER models. Indeed, about a dozen such models have been reported in the research literature [5]. These models represent attempts at modeling the temporal aspects of information more naturally and elegantly. The proposed extensions are based on quite different approaches. One approach is to devise new notational shorthands that replace some of the patterns that occur frequently in ER diagrams when temporal aspects are being modeled. One example is the pattern that occurs when modeling a time-varying attribute in the ER model. Another approach is to change the semantics of the existing ER model constructs, making them temporal. In its extreme form, this approach does not result in any new syntactical constructs—all the original constructs have simply become temporal. With this approach, it is also possible to add new constructs.

The ideal temporal ER model is easy to understand in terms of the ER model; does not invalidate legacy diagrams and database applications; and does not restrict the database to be temporal, but rather permits the designer to mix temporal and nontemporal aspects.

The existing models typically assume that their schemas are mapped to schemas in the relational model that serves as the implementation data model. The mapping algorithms are constructed to add appropriate time-valued attributes to the relation schemas. None of the models have one of the many time-extended relational models proposed [12] as their implementation model. These models have data-definition and query-language capabilities that better support the management of temporal data and would thus constitute natural candidate implementation platforms. Also, mappings to emerging models (e.g., SQL3) are missing. It is a challenge to design mappings that maximally exploit these and other candidate implementation platforms.

4.2 Logical Design

A central goal of conventional relational database design is to produce a database schema, consisting of a set of *relation schemas*. Normal forms constitute an attempt at characterizing "good" relation schemas. A wide variety of
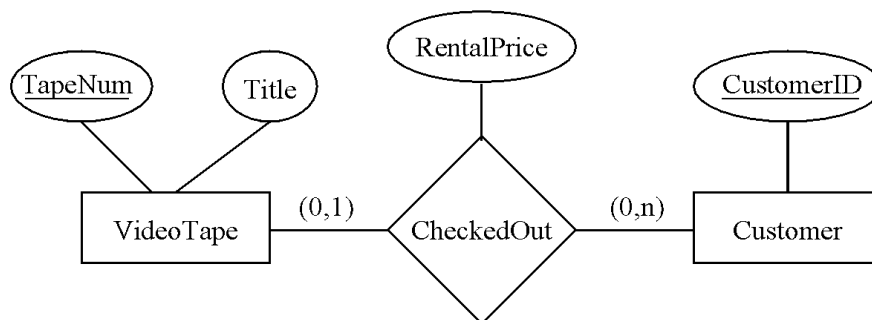

Fig. 4. Nontemporal conventional ER diagram for video rentals.

normal forms has been proposed, the most prominent being third normal form and Boyce-Codd normal form. An extensive theory has been developed to provide a solid formal footing.

The existing normalization concepts are not applicable to temporal relational data models because these models employ relational structures that are different from conventional relations. There is thus a need for new temporal normal forms and underlying concepts that may serve as important guidelines during temporal database design.

In response to this need, an array of temporal normalization concepts have been proposed [9], including temporal dependencies, keys, and normal forms. Consider the `CheckedOut` relation schema from Section 3, as exemplified in Figs. 1 and 3. Does `CustomerID` (temporally) determine `TapeNum` or vice versa? Looking at the first representation in Fig. 3 and applying conventional dependencies directly, the answer to both questions is no. (Note that the possible answers are 'no' and 'perhaps.') The second representation is so different from a regular relation that it makes little sense to directly apply conventional dependencies. The relation in Fig. 1 also rules out any of the dependencies when we apply regular dependencies directly.

Stepping back, it should be that the same dependencies hold for the `CheckedOut` relation independently of how it is represented. At *any point in time*, a customer may have checked out several tapes. In contrast, a tape can only be checked out by a single customer at a single point in time. With this view, `TapeNum` temporally determines `CustomerID`, but the reverse does not hold. This notion of dependency naturally generalizes conventional dependencies and may be applied to other dependencies than functional. With this notion of dependency, a temporal normalization theory may be built that parallels conventional normalization theory and that is independent of any particular representation of a temporal relation.

Dependencies and their associated normal forms can also be defined *between* time points, and taking into account temporal granularity [21], [22].

5   ADDING TIME TO QUERY LANGUAGES

Given the prevalence of applications that currently manage time-varying data, one might ask why a temporal query language is even needed. Is the existence of all this SQL code not proof that SQL is sufficient for writing such applications? The reality is that in conventional query languages like SQL, temporal queries *can* be expressed, but with great difficulty.

In addition to the `CheckedOut` relation from Section 3, we assume in this section a `VideoTape` relation with attributes `TapeNum`, `Title`, and `RentalPrice`. Consider first this database with only current information. To determine who has checked out which titles, SQL provides a natural solution:

```
SELECT CustomerID, Title
FROM CheckedOut, VideoTape
WHERE CheckedOut.TapeNum =
    VideoTape.TapeNum
```

We then extend the `VideoTape` and `CheckedOut` relations to record also past and future states by adding to each relation two additional attributes, `StartDate` and `EndDate`, specifying the interval of validity of the tuples. To request the *history* of who checked out which titles requires 25(!) lines of SQL: four `SELECT` statements, `UNION`ed together, performing a case analysis of how the interval of validity of `CheckedOut` overlaps the interval of validity of `VideoTape` (see pp. 106–107 of [18]).

As another example, specifying referential integrity on the nontemporal relations is trivial in SQL: "`CONSTRAINT TapeNum REFERENCES VideoTape`." When the two relations are time-varying, referential integrity requires a 28-line SQL assertion, with triply nested `EXISTS/NOT EXISTS` subqueries. (Readers are encouraged to try their hand at these two examples.) Ordinary queries on the nontemporal relations become extremely challenging when timestamp attributes are added. Even SQL experts would be hard pressed to express the following in SQL: What is the history of the average rental price for checked out video tapes?

Some 40 temporal query languages have been defined [18], most with their own data model. One of the most recent is TSQL2 [16], developed as a second-generation language by many of the designers of first-generation temporal query languages. The goal of TSQL2 was to consolidate approaches to temporal calculus-based query languages, to achieve a consensus extension to SQL-92 [10] upon which future research could be based.

With a temporal query language, simple queries should remain simple when time is added. The temporal join can be expressed in the variant of TSQL2 being proposed for inclusion into SQL3 [17] as follows:

```
VALIDTIME SELECT CustomerID, Title
FROM CheckedOut, VideoTape
WHERE CheckedOut.TapeNum =
    VideoTape.TapeNum
```

Similarly, referential integrity can be expressed as "`CONSTRAINT TapeNum VALIDTIME REFERENCES VideoTape`." Even with this minimal explanation, the reader should have no difficulty in expressing the average rental price query in this extension to SQL.

Early query languages were based on the relational algebra. Calculus-based, datalog-based, and object-oriented temporal query languages appeared later. Much of the recent work involves extensions to SQL.

As query languages are strongly influenced by the underlying data model, many of the issues raised in Section 3 have analogues in temporal query languages. As one example, whether the data model timestamps tuples or attribute values influences the language. The history of who checked out which titles can be expressed in the TempSQL [4] query language, which utilizes attribute-value timestamps, as shown in Fig. 3.

```
SELECT CustomerID, Title
WHILE ⟦CheckedOut.TapeNum =
    VideoTape.TapeNum⟧ ∩ [Now, Now] × [0, ∞]
FROM CheckedOut, VideoTape
```

Here, the `WHILE` construct restricts the time domain of the resulting attributes to those times when the equality was satisfied. The intersection ensures we only examine the data that is current in the database, that is, data with a transaction time of *now*.

Language design must consider the impact of the time-varying nature of data on all aspects of the language, including predicates on temporal values, temporal constructors, supporting states or events (or both) in the language, supporting multiple calendars, modification of temporal relations, cursors, views, integrity constraints, temporal indeterminacy, handling *now*, aggregates, schema versioning, vacuuming, and periodic data. Most of these topics have been the sole focus of one (or several) papers. However, these aspects interact in subtle ways, requiring consideration of all (or a substantial subset) to ensure that the design makes sense. Adequately documenting the design, rationale, and semantics of a comprehensive attack on the problem is daunting: the description of TSQL2 required some 700 pages [16].

Recently a set of criteria for temporal query languages has emerged. These include *temporal upward compatibility* (that is, conventional queries and modifications on temporal relations should act on the current state); support for *sequenced queries* (that request the history of something, such as the temporal join above); support for *point-based* and *interval-based* semantics; adequate expressive power; and the ability to be efficiently implemented.

6 TEMPORAL DBMS IMPLEMENTATION

There has been a vast amount of work in storage structures and access methods for temporal data, as well as a dozen-odd temporal DBMS prototypes [1]. There have been two basic approaches. Most authors assume an *integrated* approach, in which the internal modules of a DBMS are modified or extended to support time-varying data. More recently, there has been work using a *stratum* approach, in which a layer converts temporal query language statements into conventional statements executed by an underlying DBMS, which is itself not altered. While the former approach ensures maximum efficiency, the latter approach is more realistic in the medium term. In the following we will, consistent with the vast majority of papers on temporal DBMS implementation, assume an integrated approach, utilizing timestamping of tuples with time intervals.

6.1 Query Processing

Optimization of temporal queries is more involved than that of conventional queries, for several reasons. First, the relations that temporal queries are defined over are often larger; this justifies trying harder to optimize the queries, and spending more execution time to perform the optimization. Second, the predicates used in temporal queries are harder to optimize. In traditional database applications, predicates are usually equality predicates (hence the prevalence of equijoins and natural joins). In temporal queries, conjunctions of inequality predicates appear more frequently. As an example, the TSQL2 temporal join query given in Section 5 determines the overlap between validity intervals from the `CheckedOut` relation and the `VideoTape` relation. This overlap is translated into two "≤" predicates on the underlying timestamps, as follows:

```
BEGIN(CheckedOut) <= END(VideoTape) AND
    BEGIN(VideoTape) <= END(CheckedOut)
```

Conventional DBMSs focus on equality predicates and often implement inequality joins as Cartesian products, with their concomitant inefficiency.

There are new and unexploited opportunities for query optimization when time is present. The current time advances continuously and, for transaction time, the time value used most recently in updates is the largest value used so far. This implies that a natural clustering or sort order will manifest itself, which can be exploited during query optimization and evaluation. The integrity constraint `BEGIN(i) <= END(i)` holds for every interval *i*. Also, for many relations, the intervals associated with a key are contiguous in time, with one interval starting exactly when the previous interval ended (an example is the `VideoTape` relation). Semantic query optimization can exploit these integrity constraints, as well as additional ones that can be inferred.

6.2 Implementing Algebraic Operators

Attention has been directed at the common (and often expensive) temporal algebraic operators: selection, joins, aggregates, and duplicate elimination. We examine selection in the next section, on temporal indexes.

A wide variety of binary joins have been considered, including *time-join* and *time-equijoin* (TE-join), *event-join*, and *TE-outerjoin*, *contain-join*, *contain-semijoin* and *intersect-join*, and *temporal natural join* (e.g., [6]). The various algorithms proposed for these joins have generally been extensions to nested loop or merge joins that exploit sort orders or local workspace, as well as hash joins. Next, time-varying aggregates are especially challenging. While there has been much work on the topic in the data warehousing context, only a few papers have considered the more general problem. Finally, *coalescing* is an important operation in temporal databases. Coalescing merges value-equivalent tuples with intervals that overlap or meet. This operation may be implemented by first sorting the argument relation on the explicit attribute values as well as the valid time. In a subsequent scan, the merging is then accomplished.

6.3 Indexing Temporal Data

Conventional indexes have long been used to reduce the need to scan an entire relation to access a subset of its tuples, to support the selection algebraic operator and temporal joins. Indexes are even more important in temporal relations due to their size. Many temporal indexing strategies are available [14]. Many of the indexes are based on B^+-trees, which index on values of a single key; most of the remainder are based on R-trees, which index on ranges (intervals) of multiple keys. The worst-case performance for most proposals has been evaluated in terms of total space required, update per change, and several important queries. Most of this work is in the context of the selection operator. As also mentioned, indexes may be used to efficiently

implement temporal joins, coalescing, and aggregates—this is an area of active investigation.

7 SUMMARY

This paper has briefly introduced the reader to temporal data management, emphasizing what we believe are important concepts and surveying important results produced by the research community. In what remains, we first summarize the current state-of-the-art, then point to issues that remain challenges and which require further attention.

A great amount of research has been expended on temporal data models and query languages, which has shown itself as an extraordinarily complex challenge with subtle issues. We feel that the semantics of standard temporal relational schemas and their logical design are well understood, and the Bitemporal Conceptual Data Model is gaining acceptance as the appropriate model in which to consider data semantics.

Many languages have been proposed for querying temporal databases, half of which have a formal basis. The numerous types of temporal queries are fairly well understood. The TSQL2 query language has consolidated many years of research results into a single, comprehensive language. Constructs from that language are being incorporated into a new part of SQL3, called SQL/Temporal [11].

The semantics of the time domain, including its structure, dimensionality, and indeterminacy, is quite well understood, and representational issues of timestamps have recently been resolved. Operations on timestamps are now well understood, and efficient implementations exist.

Temporal joins, aggregates, and coalescing are well understood, and efficient implementations exist. More than a dozen temporal index structures have been proposed, supporting valid time, transaction time, or both. A handful of prototype temporal DBMS implementations have been developed [1].

8 OUTLOOK

While numerous important insights and results have been reported, there are still many research challenges, some of which we now consider. Quite frankly, insufficient attention is being paid, in our opinion, to many of these pressing problems, reducing the potential impact of earlier results. In many cases the core concepts are now available, but have yet to be applied, or shown how they can be applied, to simplify and automate the data management activities of those in the trenches.

We feel that there is a need for increased *legacy-awareness* in a number of areas within temporal databases. Research is needed that takes into account the reality that most databases are in fact legacy temporal databases and that the applications running on them are in fact legacy temporal database applications. In contrast, most research so far has assumed that applications will be designed using a new temporal data model, implemented using novel temporal query languages, and run on as yet nonexistent temporal DBMSs. In the short to medium term, this is an unrealistic assumption. Indeed, in part because of this and despite the

obvious need in the marketplace, no prominent commercial temporal relational DBMS exists.

The recent growth in database architectures, including the various types of middleware, prompts a need for increased *architecture-awareness*. Studies are needed to provide the concepts and approaches for third-party developers to efficiently and effectively implement temporal database technology while maximally exploiting available architectural infrastructure and the functionality already offered by existing DBMSs. The resulting temporal DBMS architectures will provide a highly relevant alternative to the standard integrated architecture that is generally assumed. As a next step, research is needed to exploit existing and novel performance-improving advances, such as temporal algebraic operator implementations and indexes, in these architectures. Finally, approaches for transitioning legacy applications will become increasingly sought after as temporal technology moves from research to practice.

Also, there has been little work on adding time to so-called fourth-generation languages that are revolutionizing user interfaces for commercially available DBMSs. Figs. 1 and 3 further emphasize the need for ways to *visualize temporal data*. Scrolling down a table with additional timestamp attributes is not likely to be effective.

We feel that results on the *conceptual design* of temporal databases have potential for finding application in practice, but additional research is needed. When database designers actually understand temporal database concepts, they are able to design better databases using existing models and tools. A central challenge is to provide complete conceptual models, with associated design tools, that cover all aspects of designing a temporal database; empirical evaluation of these by real users is needed to provide essential insights.

Concerning *performance*, more empirical studies are needed to compare temporal algebraic operator implementations, and to possibly suggest even more efficient implementations. While preliminary performance studies have been carried out for each of the proposed temporal indexes in isolation, there has been little effort to empirically compare them. More work is also needed on exploiting temporal indexes in algebraic operations other than selection. Finally, there has been little work in refining and validating cost models of temporal operators, or of developing and maintaining database statistics. For example, the cardinality (number of specific values) of an attribute is less useful than the average cardinality at a point in time. Another useful statistic is the number of *long-lived tuples*, the presence of which is the bane of some index structures and temporal algebraic operators.

Active databases, which include rules for responding to changes to the database and to external events, are being extended to take into account prior history (e.g., a temperature reading in a nuclear power plant may be acceptable if it is decreasing; that same temperature reading may signal a problem if it is increasing). As yet there has been little integration of rule constructs and temporal constructs.

We expect the area of *spatiotemporal databases* to become increasingly important. Providing built-in support for both space and time makes it convenient to manage objects with extents in physical space and time, opening up for new

database applications. For example, many "moving objects" such as rental cars will be equipped with GPS devices, and their trajectories will be stored in databases. The integration of temporal databases with spatial databases offers new challenges. As a single example, no really good index seems to exist for the trajectories of moving objects.

Temporal data mining is a largely unexplored area. While extracting static associations from a mass of data is an important goal, more effort needs to be focused on associations that capture time-varying behavior, such as "When stock A goes up, stock B goes up within two weeks."

The fairly recent focus among vendors, users, and researchers alike on *data warehousing* has brought new prominence to temporal databases. Inmon, known as the founder of data warehousing, cites time variance as one of four salient characteristics of a data warehouse [7], and there is general consensus that a data warehouse is likely to contain several years of time-referenced data. Being temporal, data warehouses are thus prime candidates to benefit from the advances in temporal databases. But there appears to be a lack of integration. In fact, some of the original impetus for a separate data model and query language for data warehouses arose from a perceived lack of temporal support in the relational model and SQL. We feel that increased support for time-varying data in SQL will enable greater use of relational databases directly in data warehousing applications. If the differentiation between relational and star schemas could be reduced, users and developers could exploit an overarching data model for the enterprise. Conversely, the special architecture of a data warehouse and the emphasis on supporting advanced query functionality, e.g., application-specific time-series analysis, bring novel challenges to temporal database researchers.

Another active area of commercial products is that of *time-series* abstract data types (i.e., Informix's datablades, Oracle's cartridges). These extensions are highly useful for specialized applications, particularly in the financial sphere, but do not address the general problem of easy expression of temporal constraints, queries, and modifications discussed in Section 5. Rather, a more comprehensive approach along the lines of the extensions being considered for SQL/Temporal appears to be more appropriate. Interestingly, the object-oriented features of SQL3 are largely orthogonal to and unaffected by these temporal extensions.

Adopting a longer term perspective, we find it likely that new database management technologies and application areas will continue to emerge that provide 'temporal' challenges. Due to the ubiquity of time and its importance to most database management applications, and because built-in temporal support generally offers many benefits and is challenging to provide, we expect research in temporal aspects of new database management technologies for existing and new application areas to continue to flourish. For example, in the relatively new area of web and databases, we expect research on temporal aspects to emerge. In fact, this universal concern for the specific needs of time-varying data could eventually mean that 'temporal databases' will cease to be considered a separate subdiscipline of database research: Ultimately, *all* aspects of data management will naturally accommodate the inherently temporal nature of data.

REFERENCES

[1] M. Böhlen, "Temporal Database System Implementations," *ACM SIGMOD Record*, vol. 24, no. 4, pp. 53–60, Dec. 1995.

[2] *Recent Advances in Temporal Databases: Proc. Int'l Workshop Temporal Databases*," J. Clifford and A. Tuzhilin, eds., Workshops in Computing Series, Springer-Verlag, 1995.

[3] *Temporal Databases: Research and Practice*, O. Etzion, S. Jajodia, and S. Sripada, eds., Lecture Notes in Computer Science 1399, Springer-Verlag, 1998.

[4] S.K. Gadia and G. Bargava, "SQL-Like Seamless Query of Temporal Data," in [15].

[5] H. Gregersen and C.S. Jensen, "Temporal Entity-Relationship Models—A Survey," *IEEE Trans. Knowledge and Data Eng.*, to appear.

[6] H. Gunadhi and A. Segev, "A Framework for Query Optimization in Temporal Databases," *Proc. Fifth Conf. Statistical and Scientific and Databases Management*, pp. 131–147, Charlotte, N.C., Apr. 1990.

[7] W.H. Inmon, *Building the Data Warehouse*, John Wiley and Sons, 1992.

[8] *A Consensus Glossary of Temporal Database Concepts—Feb. 1998 Version*, C.S. Jensen and C.E. Dyreson, eds., pp. 367–405, in [3].

[9] C.S. Jensen and R.T. Snodgrass, "Semantics of Time-Varying Information," *Information Systems*, vol. 21, no. 4, pp. 311–352, 1996.

[10] J. Melton and A.R. Simon, *Understanding the New SQL: A Complete Guide*, Morgan Kaufmann, 1993.

[11] *SQL/Temporal*, J. Melton, ed., ISO/IEC JTC 1/SC 21/WG 3 DBL-MCI-0012, July 1996.

[12] G. Özsoyoglu and R.T. Snodgrass, "Temporal and Real-Time Databases: A Survey," *IEEE Trans. Knowledge and Data Eng.*, vol. 7, no. 4, pp. 513–532, Aug. 1995.

[13] J.F. Roddick and J.D. Patrick, "Temporal Semantics in Information Systems—A Survey," *Information Systems*, vol. 17, no. 3, pp. 249–267, Oct. 1992.

[14] B. Salzberg and V.J. Tsotras, "A Comparison of Access Methods for Time Evolving Data," *ACM Computing Surveys*, to appear.

[15] *Proc. Int'l Workshop Infrastructure for Temporal Databases*, R.T. Snodgrass, ed., Arlington, Texas, June 1993.

[16] I. Ahn, G. Ariav, D. Batory, J. Clifford, C.E. Dyreson, R. Elmasri, F. Grandi, C.S. Jensen, W. Käfer, N. Kline, K. Kulkarni, T.Y. Leung, N. Lorentzos, J.F. Roddick, A. Segev, M.D. Soo, and S.M. Sripada, *The TSQL2 Temporal Query Language*, R.T. Snodgrass, ed., Kluwer, 1995; URL: http://www.cs.arizona.edu/people/rts/sql3.html.

[17] R.T. Snodgrass, M.H. Böhlen, C.S. Jensen, and A. Steiner, *Adding Valid Time to SQL/Temporal*, ANSI X3H2-96-501r2, ISO/IEC JTC 1/SC 21/WG 3 DBL-MAD-146r2, Nov. 1996.

[18] R.T. Snodgrass, "Temporal Databases," *Advanced Database Systems*, Part II, Morgan Kaufmann, 1997.

[19] *Temporal Databases: Theory, Design, and Implementation*, A. Tansel, J. Clifford, S. Gadia, S. Jajodia, A. Segev, and R.T. Snodgrass, eds., Benjamin/Cummings, 1994.

[20] V.J. Tsotras and X.S. Wang, "Temporal Databases," *Encyclopedia of Electrical and Electronics Eng.*, John Wiley and Sons, 1999.

[21] X.S. Wang, C. Bettini, A. Brodsky, and S. Jajodia, "Logical Design for Temporal Databases with Multiple Granularities," *ACM Trans. Database Systems*, vol. 22, no. 2, pp. 115–171, June 1997.

[22] J. Wijsen, "Temporal FDs on Complex Objects," *ACM Trans. Database Systems*, vol. 23, no. 4, Dec. 1998, to appear.

[23] Y. Wu, S. Jajodia, and X.S. Wang, "Temporal Database Bibliography Update," pp. 338–366, in [3].

Christian S. Jensen is a research professor of computer science at Aalborg University, where he directs the Nykredit Center for Database Research. He was with the University of Arizona during several sabbaticals and, prior to joining the faculty of Aalborg University, he conducted his graduate studies at the University of Maryland. His research interests are in the area of database systems and include issues of semantics, modeling, and performance. With his colleagues, he receives substantial national and international funding for his research. He is a member of the Editorial Board of *IEEE Transactions on Knowledge and Data Engineering*. He was general chair of the 1995 International Workshop on Temporal Databases, was a Program Committee vice chair and a Best Papers Awards Committee member for the 1998 IEEE International Conference on Data Engineering, and co-chairs the Program Committee for the 1999 Workshop on Spatio-Temporal Database Management. He continues to serve on the Program Committee and other committees for a number of conferences, including ACM SIGMOD, CAiSE, EDBT, IEEE Data Engineering, SSDBM, and VLDB, and he serves regularly as a reviewer for all the major database journals. He co-directs the TimeCenter, an international center for the support of temporal database applications on traditional and emerging DBMS technologies. He is a senior member of the IEEE, and a member of the ACM and the IEEE Computer Society.

Richard T. Snodgrass received his PhD degree from Carnegie Mellon University in 1982 and joined the University of Arizona in 1989, where he is a professor of computer science. He is chair of ACM SIGMOD. He is an associate editor of *ACM Transactions on Database Systems*, and is on the Editorial Boards of *IEEE Transactions on Knowledge and Data Engineering* and the *International Journal on Very Large Databases*. He chaired the Program Committees for the 1994 ACM SIGMOD Conference and the 1993 International Workshop on an Infrastructure for Temporal Databases; was a vice-chair of the Program Committee for the 1993 and 1994 International Conferences on Data Engineering; and will chair the American Program Committee for the 2001 International Conference on Very Large Databases. He chaired the TSQL2 Language Design Committee; edited the book, *The TSQL2 Temporal Query Language* (Kluwer Academic Press); and is now working closely with the ISO SQL3 Committee to add temporal support to that language. He initiated the SQL/Temporal part of the SQL3 Draft Standard. He is a co-author of *Advanced Database Systems* (Morgan Kaufmann), and a co-editor of *Temporal Databases: Theory, Design, and Implementation* (Benjamin/Cummings). He co-directs the TimeCenter. His research interests include temporal databases, query language design, query optimization and evaluation, storage structures, and database design. He is a senior member of the IEEE, a member of the IEEE Computer Society, and a fellow of the ACM.