Schema Evolution and the Relational Algebra

Edwin McKenzie and Richard Snodgrass

Department of Computer Science
University of North Carolina
Chapel Hill, NC 27514

March 13, 2009

Abstract

In this paper we discuss extensions to the conventional relational algebra to support both aspects of transaction time, evolution of a database's *contents* and evolution of a database's *schema*. We define a relation's schema to be the relation's temporal *signature*, a function mapping the relation's attribute names onto their value domains, and *class*, indicating the extent of support for time. We also introduce commands to change a relation, now defined as a triple consisting of a sequence of classes, a sequence of signatures, and a sequence of states. A semantic type system is required to identify semantically incorrect expressions and to enforce consistency constraints among a relation's class, signature, and state following update. We show that these extensions are applicable, without change, to historical algebras that support valid time, yielding an algebraic language for the query and update of temporal databases. The additions preserve the useful properties of the conventional algebra.

A database's *schema* describes the *structure* of the database; the *contents* of the database must adhere to that structure [Date 1976, Ullman 1982]. *Schema evolution* refers to changes to the database's schema over time. Conventional databases allow only one schema to be in force at a time, requiring *restructuring* (also termed *logical reorganization* [Sockut & Goldberg 1979]) when the schema is modified. With the advent of databases storing past states [McKenzie 1986], it becomes desirable to accommodate multiple schemas, each in effect for an interval in the past. Schema *versioning* refers to retention of past schemas resulting from schema evolution.

In an earlier paper [McKenzie & Snodgrass 1987A] we proposed extensions to the conventional relational algebra [Codd 1970] that model the evolution of a database's contents. We did not, however, consider the evolution of a database's schema. In this paper, we provide further extensions to the conventional relational algebra that model the evolution of a database's schema. The extensions that support evolution of a database's contents are repeated here for completeness and because the extensions supporting schema evolution are best explained in concert with those earlier extensions.

1 Approach

Languages for database query and update exist at no less than three levels of database abstraction. At the user-interface level, calculus-based languages such as SQL are available for expressing query and update operations. At the algebraic level, the relational algebra is the formal, abstract language for expressing these same operations. Finally, at the physical level, query and update operations can be defined in terms of data structures and access strategies. In this paper we focus on language definition at the algebraic level. Our goal here is to define formally an algebraic language for database query and update that supports evolution (and versioning) of a database's schema, as well as its contents. To do so, we extend the relational algebra to handle one aspect of time in databases.

Time must be added to the underlying data model before it can be added to the relational algebra. In previous papers, we identified three orthogonal aspects of time that a database management system (DBMS) needs to support: valid time, transaction time, and user-defined time [Snodgrass & Ahn 1985, Snodgrass & Ahn 1986]. *Valid time* concerns modeling time-varying reality. The valid time of, say, an event is the clock time when the event occurred in the real world, independent of the recording of that event in some database. *Transaction time*, on the other hand, concerns the storage of information in the database. The transaction time of an event is the transaction number (an integer) of the transaction that stored the information about the event in the database. *User-defined time* is an uninterpreted domain for which the DBMS supports the operations of input, output, and perhaps comparison. As its name implies, the semantics of user-defined time is provided by the user or application program. These three types of time are orthogonal in the support required of the DBMS.

In these same papers, we defined four classes of relations depending on their support for valid time and transaction time: snapshot relations, rollback relations, historical relations, and temporal relations. User-defined time, unlike valid time and transaction time, is already supported by the relational algebra, in that it is simply another domain, such as integer or character string, provided by the DBMS [Bontempo 1983, Overmyer & Stonebraker 1982, Tandem 1983]. *Snapshot relations* support neither valid time nor transaction time. They model an enterprise at one particular point in time. As a snapshot relation is changed to reflect changes in the enterprise being modeled, past states of the relation, representing past states of the enterprise, are discarded. A snapshot relation consists of a set of *tuples* with the same set of *attributes*, and is usually represented as a two-dimensional table with attributes as columns and tuples as rows, as shown in Figure 1. Note that snapshot relations are exactly those relations supported by the relational algebra. Hence, for clarity, we will refer to the relational algebra hereafter as the snapshot algebra. *Rollback relations* support transaction time but do not support valid time. They may be represented as a sequence of snapshot states indexed by transaction time, as shown in Figure 2. (Here, the last transaction deleted one tuple and appended another.) Because they record the history of database activity, rollback relations can be rolled back to one of their past snapshot states for querying.

Historical relations support valid time but do not support transaction time. They model the history, as it is best known, of an enterprise. When an historical relation is changed, however, its past state, like that of a snapshot relation, is discarded. An historical relation may be represented as a three-dimensional solid, as shown in Figure 3. Because they record the history of the enterprise

Figure 1: Snapshot Relation

Figure 2: Rollback Relation

Figure 3: Historical Relation

Figure 4: Temporal Relation

being modeled, historical relations support historical queries. They do not, however, support rollback operations. *Temporal relations* support both valid time and transaction time. They may be represented as a sequence of historical states indexed by transaction time, as shown in Figure 4. Because they record both the history of the enterprise being modeled and the history of database activities, temporal relations support both historical queries and rollback operations.

Data models that support these four classes of relations have several important properties. First, a relation's schema can no longer be defined in terms of the relation's attributes alone; it must also include the relation's *class* (i.e., snapshot, rollback, historical, or temporal). Second, rollback and temporal relations, unlike snapshot and historical relations, are *append-only* relations. Information, once added to a rollback or temporal relation, cannot be deleted; otherwise, rollback operations could not be supported. Third, valid time and transaction time are orthogonal aspects of time. A relation may support either valid time or transaction time without supporting both. Also, the time when an enterprise changes (i.e., valid time) need not be, and usually will not be, the same as the time when the database is updated (i.e., transaction time) to reflect that change.

4

Finally, the same measures of time need not be used for valid and transaction time. For example, a temporal relation will have a variable granularity, which changes with each update, for transaction time but could have a fixed granularity (e.g., second) for valid time.

Fortunately, since valid time and transaction time are orthogonal, they may be studied in isolation. There have already been several proposals for adding valid time to the snapshot algebra [Ben-Zvi 1982, Clifford & Croker 1987, Gadia 1984, Gadia 1986, Jones et al. 1979, Lorentzos & Johnson 1987, McKenzie 1988, Navathe & Ahmed 1987, Tansel 1986, Yeung 1986], so we will not consider valid time in detail. We focus here on extensions to support transaction time.

In a previous paper [McKenzie & Snodgrass 1987A] we discussed extensions to the snapshot algebra to enable it to handle one aspect of transaction time: evolution of a database's contents. To handle evolution of the contents of a database containing both snapshot and rollback relations, we defined a relation to be a sequence of snapshot states, ordered by transaction number. Snapshot relations were modeled as single-element sequences while rollback relations were modeled as multi-element, append-only sequences. We also defined a database to be an ordered pair whose first component was a function from identifiers (i.e., relation names) to relations and whose second component was the transaction number of the most recently committed transaction on the database. We then augmented the algebra with a rollback operator to make past states of rollback relations available in the algebra and encapsulated this extended algebra in a language of commands for database update. Finally, we showed that the same approach could be used to extend an arbitrary historical algebra (i.e., an algebra supporting valid time) to handle evolution of the contents of a database containing both historical and temporal relations.

We now extend the relational algebra to handle the other aspect of transaction time: evolution of a database's schema. Schema evolution is associated solely with transaction time, since it defines how reality is modeled by the database. For example, a person's marital status is a (time-varying) aspect of reality, but the decision whether to record marital status, encoded in the schema, is a (time-varying) aspect of the database. Hence, we add the schema to the domain of database states. The schema and the contents of the database combine to define the database's state.

A relation's structure can no longer be defined in terms of the relation's attributes alone; it must also be defined in terms of the relation's class. Hence, we define a relation's schema to be a pair consisting of the relation's class and a function, which we refer to as the relation's *signature*, that maps the relation's attribute names onto their value domains. (If the identification of primary keys is desirable, this would also properly go into the signature.) The relation's contents, which we refer to as the relation's *state*, always must be consistent with both the relation's class and the relation's signature.

Our model of transaction time is predicated on two assumptions. First, we assume that a database may contain snapshot, rollback, historical, and temporal relations. Second, we assume that the class and signature, as well as the contents, of each relation in the database may change over time. For example, a relation defined initially as a snapshot relation could be changed to be a historical, rollback, or temporal relation. Later, it could be changed to be a snapshot relation once again.

A model of transaction time in a database containing relations of all four classes, must main-

tain, for each relation, its current class, signature, and state. The model also must retain, for each relation, its signature and state for those intervals during which its class was either rollback or temporal. Hence, we define a relation to be a triple consisting of a sequence of classes, a sequence of signatures, and a sequence of states, all ordered by transaction number. The class sequence records the relation's current class and intervals when the relation's class was either rollback or temporal. Similarly, the signature and state sequences record the relation's current signature and state and all changes in signature and state during intervals when the relation's class was either rollback or temporal. We also define a database state to be a function from identifiers (i.e., relation names) to relations. Finally, we define a database to be an ordered pair whose first component is a database state and whose second component is the transaction number of the most recently committed transaction on the database.

We define four commands for database update: `define_relation`, `modify_relation`, `destroy`, and `rename_relation`. The `define_relation` command assigns a new class and signature, along with the empty snapshot or historical state, to an undefined relation. The `modify_relation` command changes the current class, signature, and state of a defined relation. The `destroy` command is the counterpart of the `define_relation` command. It either physically or logically deletes from the database the current class, signature, and state of a relation, depending on the relation's class when the command is executed. The `rename_relation` command binds the current class, signature, and state of a relation to a new identifier. We assume that these commands execute in the context of a single, previously created database. Hence, no commands are necessary to create or delete the database. Since we are considering modeling transaction time from a functional, rather than from a performance, viewpoint, commands affecting access methods, storage mechanisms, or index maintenance are also not relevant.

Allowing a database's schema, as well as its contents, to change increases the complexity of the language. If we allow the database's schema to change, an algebraic expression that is semantically correct for the database's schema when one command executes may not be semantically correct for the database's schema when another command executes. We need a mechanism for identifying semantically incorrect algebraic expressions relative to the database's schema when each command executes and a way of ensuring that the schema and contents of the database state resulting from the command's execution are compatible. To identify semantically incorrect expressions, we introduce a *semantic type system* and augment all commands to do type-checking.

Finally, we encapsulate commands within a system of transactions to provide for both single-command and multiple-command transactions. A multiple-command transaction, like a single-command transaction, is treated as an atomic update operation, whether it changes one relation or several relations. Transactions are specified by the keywords `begin_transaction` and either `commit_transaction` or `abort_transaction`, the later depending on whether the transaction commits or aborts.

This language was designed to satisfy several other objectives as well. First, the language subsumes the expressive power of the snapshot algebra. For every expression in the snapshot algebra, there is an equivalent expression in the language. Second, the language subsumes the expressive power of any arbitrary historical algebra. For every expression in an historical algebra, there is an equivalent expression in the language. Third, the language ensures that all data stored in a relation when its class was either rollback or temporal are retained permanently and are accessible

6

via a rollback operator, even after the relation is logically deleted from the database. Fourth, commands change only a relation's class, signature, and state current at the start of a transaction. Past data that are retained to support rollback operations, once saved, are never changed. Hence, the language accommodates implementations that use write-once-read-many (WORM) optical disk to store non-current class, signature, and state information.

We employ denotational semantics to define the semantics of the language, due to its success in formalizing operations involving side-effects, such as assignment, in programming languages [Gordon 1979, Stoy 1977]. In defining the semantics of commands and algebraic operators, we have favored simplicity of semantics at the expense of efficient direct implementation. The language would be inefficient, in terms of storage space and execution time, if mapped directly into an implementation. However, the semantics do not preclude more efficient implementations using optimization strategies for both storage and retrieval of information. Elsewhere, we review briefly some of the techniques for efficient implementation, compatible with our semantics, that have been proposed by others [McKenzie 1988]. We also, without loss of generality, assume that transactions are executed sequentially in a single-user environment. Our approach applies equally to environments that permit the concurrent execution of transactions as long as their concurrency control mechanisms induce a serialization of transactions.

Our language for supporting the above extensions will be the topic of the next section. Additional aspects of the rollback operators are discussed briefly in Section 3. Section 4 will review related work and compare our approach with those of others.

2 The Language

In this section we provide the syntax and denotational semantics of our language for database query and update. In denotational semantics, a language is described by assigning to each language construct a *denotation* – an abstract entity that models its meaning [Gordon 1979, Scott 1976, Stoy 1977, Strachey 1966]. We chose denotational semantics to define the language because denotational semantics combines a powerful descriptive notation with rigorous mathematical theory [Gordon 1979], permitting the precise definition of database state. First, we define the syntax of the language. Then we define the language's semantic domains and a semantic type system for expressions. Finally, we define the semantic functions that map the language constructs onto their denotations.

2.1 Syntax

The language has three basic types of constructs: programs, commands, and expressions. A program is a sequence of one or more transactions. Both single-command and multi-command transactions are allowed. Commands occur within transactions; they change relations (e.g., define a relation, modify a relation, delete a relation). Expressions occur within commands and denote a single snapshot or historical state. We represent these three types of constructs by the syntactic categories:

$\mathcal{PROGRAM}$	Category of programs
$\mathcal{COMMAND}$	Category of commands
$\mathcal{EXPRESSION}$	Category of expressions

We use Backus-Naur Form to specify here the syntax of programs, commands, and expressions in terms of their *immediate constituents* (i.e., the highest-level constructs that make up programs, commands, and expressions). The complete syntax of the language, including definitions of the lower-level constituents such as identifiers and snapshot states is given elsewhere [McKenzie 1988].

$$P \quad ::= \quad \texttt{begin\_transaction } C \texttt{ commit\_transaction}$$
$$| \texttt{ begin\_transaction } C \texttt{ abort\_transaction}$$
$$| P_1 ; P_2$$

$$C \quad ::= \quad \texttt{define\_relation}(I, Y, Z) \, | \, \texttt{modify\_relation}(I, Y', Z', E)$$
$$| \texttt{ destroy}(I) \, | \, \texttt{rename\_relation}(I_1, I_2) \, | \, C_1, C_2$$

$$E \quad ::= \quad [\texttt{snapshot}, Z, S] \, | \, [\texttt{historical}, Z, H] \, | \, I$$
$$| E_1 \cup E_2 \, | \, E_1 - E_2 \, | \, E_1 \times E_2 \, | \, \boldsymbol{\pi} \, X(E) \, | \, \boldsymbol{\sigma} \, F(E)$$
$$| E_1 \,\hat{\cup}\, E_2 \, | \, E_1 \,\hat{-}\, E_2 \, | \, E_1 \,\hat{\times}\, E_2 \, | \, \hat{\boldsymbol{\pi}} \, X(E) \, | \, \hat{\boldsymbol{\sigma}} \, F(E)$$
$$| \boldsymbol{\rho}(I, N) \, | \, \hat{\boldsymbol{\rho}}(I, N) \, | \, (E)$$

$$Y' \quad ::= \quad Y \, | \, *$$

$$Y \quad ::= \quad \texttt{snapshot} \, | \, \texttt{rollback} \, | \, \texttt{historical} \, | \, \texttt{temporal}$$

$$Z' \quad ::= \quad Z \, | \, *$$

$$Z \quad ::= \quad (I_{1,1} : I_{1,2}, \ldots, I_{m,1} : I_{m,2})$$

where,

$C, C_1,$ and C_2 range over the category $\mathcal{COMMAND}$;

$E, E_1,$ and E_2 range over the category $\mathcal{EXPRESSION}$;

F ranges over the category $\mathcal{SIGMA \ EXPRESSION}$ of boolean expressions
of elements from the categories $\mathcal{IDENTIFIER}$ and $\mathcal{STRING}$ (i.e., the category
of strings in an alphabet), the relational operators, and the logical operators;

H ranges over the category $\mathcal{H\text{-}STATE}$ of alphanumeric representations of historical
states in an arbitrary historical algebra;

$I, I_{1,1}, \ldots, I_{m,2}$ range over the category $\mathcal{IDENTIFIER}$ of alphanumeric
identifiers;

N ranges over the category $\mathcal{NUMERAL}$ of decimal numerals;

$P, P_1,$ and P_2 range over the category $\mathcal{PROGRAM}$;

S ranges over the category $\mathcal{S}\text{-}\mathcal{STATE}$ of alphanumeric representations of snapshot states;

X ranges over the category $\mathcal{IDENTIFIER\ LIST}$;

Y ranges over the category $\mathcal{CLASS}$ of character strings denoting relation classes; and

Z ranges over the category $\mathcal{SIGNATURE}$ of alphanumeric representations of signatures.

An expression, which evaluates to either a snapshot or historical state, may be a constant (i.e., an ordered triple consisting of a relation class, signature, and state); an identifier I, representing the current state of the relation denoted by I; or an algebraic operator on either one or two other expressions. The allowable operators include the five operators that serve to define the snapshot algebra and the operators that serve to define an arbitrary historical algebra [McKenzie & Snodgrass 1987B]. Any operator in a given historical algebra may be included in the language as long as expressions involving that operator evaluate to a single historical state. Because historical algebras have different sets of operators, we show here only the historical counterparts to the conventional algebraic operators, simply for illustration. Each is represented as $\hat{op}$ to distinguish it from its snapshot algebra counterpart op. We also have included two additional operators, a rollback operator ρ and its historical counterpart $\hat{\rho}$. The rollback operator ρ takes two arguments, an identifier I and a transaction number N, and retrieves from the relation denoted by I the snapshot state current at the time of transaction N. Similarly, the rollback operator $\hat{\rho}$ retrieves from the relation denoted by I the historical state current at the time of transaction N.

EXAMPLES. The following are two examples of syntactically correct expressions in the language. The first is a constant and the second is an expression involving both a rollback operator and a constant. Their semantics will be specified in Sections 2.3 and 2.4. Because the historical algebras all define historical relations differently, we show in this paper only examples involving snapshot and rollback relations. Each example, however, has, for a given arbitrary historical algebra, an analogue involving historical and temporal relations.

```
[snapshot, (sname:string, class:string), (sname:"Phil", class:"junior"),
                                          (sname:"Linda", class:"senior"),
                                          (sname:"Ralph", class:"senior")]
```

π (sname)(ρ(R1,4))$\times$[snapshot, (course:string), (course:"English")]

Note that the alphanumeric representation of a signature includes both the names of attributes and the names of the attributes' value domains. □

There are four commands in the language. We present here a brief description of each command, with some examples. The semantics of commands will be defined formally in Section 2.5.

The `define_relation` command binds a class, a signature, and an empty relation state to an identifier I.

EXAMPLE.

```
define_relation(R1, snapshot, (sname:string, class:string))
```

Here, the identifier `R1` is defined to denote a snapshot relation with two attributes, sname and class. The contents of the relation is, by default, the empty set. □

The `modify_relation` command may change the current class, signature, or state of a rela-tion. Command parameters specify the new class, signature, and state. The special symbol "∗" represents, depending on context, either the current class or the current signature of a relation. It may appear as a parameter in a `modify_relation` command to indicate that a relation's new class (or signature) is simply the relation's current class (signature), unchanged.

EXAMPLES.

```
modify_relation(R1, *, *, [snapshot, (sname:string, class:string),
                          (sname:"Phil", class:"junior"),
                          (sname:"Linda", class:"senior"),
                          (sname:"Ralph", class:"senior")])

modify_relation(R1, *, (sname:string, course:string),
                     π (sname)(R1) × [snapshot, (course:string),
                                      (course:"English")])

modify_relation(R1, rollback, *, R1)
```

The first command changes the state of the relation denoted by `R1` but leaves the relation's class and signature unchanged. The second command changes the relation's signature and state, but not its class. The third command changes only the relation's class, as the expression `R1` evaluates to the current state of the relation. □

The `destroy` command deletes, either physically or logically, the current class, signature, and state of a relation, depending on the relation's class when the command is executed. The `rename_relation` command renames a relation by binding its current class, signature, and state to a new identifier.

EXAMPLES.

```
destroy(R1)

rename_relation(R2, R1)
```

Here we first delete the relation denoted by `R1` and then rename the relation denoted by `R2` as `R1`. □

Programs in the language contain two types of transactions, committed transactions and aborted transactions. Committed transactions are transactions, which the user initiates, that eventually commit. Aborted transactions are transactions, which the user initiates, that for some reason, dictated either by the user or by the system, abort rather than commit. The semantics of programs will be defined formally in Section 2.6.

2.2   Semantic Domains

A program denotes the database resulting from the execution of one or more transactions, in order, on an empty database. By defining the database that results from the execution of an arbitrary sequence of transactions, we specify the semantics of that transaction sequence, and hence the semantics of the language. In this section, we will define formally the flat domain (i.e, a domain with a trivial partial ordering [Schmidt 1986]) of databases; later sections will provide the connection between the syntactic category of programs and the semantic domain of databases. All domains introduced are flat domains and the notation $\{\cdots\}$ is used to represent flat domains.

Assume that we are given the domain $\mathcal{D} = \{\mathcal{D}_1, \ldots, \mathcal{D}_e\}$, where each domain $\mathcal{D}_u$, $1 \leq u \leq e$, is an arbitrary, non-empty, finite or denumerable set. Then, we can define the following semantic domains.

$$\mathcal{TRANSACTION\ NUMBER} = \{0, 1, ...\}$$

A transaction number is a non-negative integer that identifies a transaction that changes the database. The transaction number assigned to a transaction can be viewed as that transaction's time-stamp.

$$\mathcal{RELATION\ CLASS} = \{\text{UNDEFINED, SNAPSHOT, ROLLBACK, HISTORICAL, TEMPORAL}\}$$

A relation is either undefined or defined to be a snapshot, rollback, historical, or temporal relation.

$$\mathcal{RELATION\ SIGNATURE} = \mathcal{IDENTIFIER} \rightarrow [\mathcal{D} + \{\text{UNBOUND}\}]$$

where the notation "+" on domains means the disjoint union of domains. A relation's signature is a function that maps identifiers either onto a domain $\mathcal{D}_u$, $1 \leq u \leq e$ or onto UNBOUND. If a signature maps an identifier onto UNBOUND, then the identifier is unbound in that signature (i.e., it is associated with no domain). If, however, a signature maps an identifier onto a domain, then that mapping defines an attribute.

11

$\mathcal{SNAPSHOT}$ $\mathcal{STATE}$ = Domain of all semantically correct snapshot states (sets of m-tuples), as defined in the snapshot algebra [Maier 1983], for elements of the domain $\mathcal{RELATION}$ $\mathcal{SIGNATURE}$ and the domain $\{\mathcal{D}_1 + \cdots + \mathcal{D}_e\}$, where $\emptyset$ is the empty snapshot state. Hence, a snapshot state s on a relation signature z is a finite set of mappings from $\{I \mid z(I) \neq \text{UNBOUND}\}$ to $\mathcal{D}$, with the restriction that for each mapping $st \in s$, $st(I) \in z(I)$.

$\mathcal{HISTORICAL}$ $\mathcal{STATE}$ = Domain of all semantically correct historical states as defined in an arbitrary historical algebra (e.g., as defined in [Ben-Zvi 1982, Clifford & Croker 1987, Gadia 1984, Gadia 1986, Jones et al. 1979, Lorentzos & Johnson 1987, McKenzie 1988, Navathe & Ahmed 1987, Tansel 1986, Yeung 1986]).

$$\mathcal{RELATION} = [\,\mathcal{RELATION} \, \mathcal{CLASS} \times \mathcal{TRANSACTION} \, \mathcal{NUMBER} \times$$
$$[\,\mathcal{TRANSACTION} \, \mathcal{NUMBER} + \{-\}\,]\,]^{+} \times$$
$$[\,\mathcal{RELATION} \, \mathcal{SIGNATURE} \times \mathcal{TRANSACTION} \, \mathcal{NUMBER}\,]^{*} \times$$
$$[\,[\,\mathcal{SNAPSHOT} \, \mathcal{STATE} \times \mathcal{TRANSACTION} \, \mathcal{NUMBER}\,] +$$
$$[\,\mathcal{HISTORICAL} \, \mathcal{STATE} \times \mathcal{TRANSACTION} \, \mathcal{NUMBER}\,]\,]^{*}$$

where the special element "–" stands for the *present* time. A relation is thus an ordered triple consisting of

- a sequence of (relation class, transaction number, transaction number or "–") triples,

- a sequence of (relation signature, transaction number) pairs, and

- a sequence of (relation state, transaction number) pairs.

Relations are dynamic objects whose class, signature, and state are all allowed to change over time. For example, a relation defined initially as a snapshot relation could be modified to be a historical, rollback, or temporal relation. Later, the relation could be modified to be a snapshot relation once again. Every relation always has at least one element in its class sequence, the last element recording the relation's current class (i.e., undefined, snapshot, rollback, or temporal). Any other elements in the sequence record intervals when the relation's class was either rollback or temporal.

A relation's signature (state) sequence will be empty only if the relation is currently undefined and it was never a rollback or temporal relation. If a relation is currently other than undefined, there is at least one element in its signature (state) sequence, the last element recording the relation's current signature (state). Any other elements in the sequence record the signature (state) of the relation when its class was either rollback or temporal.

The transaction-number components of all elements, but the last element, in a relation's class sequence can be viewed as time-stamps defining a fixed, closed interval during which the

element's class component was the relation's class. In contrast, the third component of the last element in the sequence is always "–"; it is used to define an interval of dynamic length that always extends to the present. The transaction-number component of each element in a relation's signature (state) sequence can be viewed as a time-stamp indicating when the element's signature (state) was entered into the database and became the relation's current signature (state). Since we assume that database changes occur sequentially, the transaction-number components of a signature (state) sequence, while not necessarily consecutive, will be nevertheless strictly increasing. Thus, we can interpolate on the transaction-number component of elements in a relation's signature (state) sequence to determine the signature (state) of the relation at any time its class was either rollback or temporal.

EXAMPLE. The following is a sample relation. For notational convenience in this and later examples, we show only the attribute portion of a signature (i.e., the partial function from attribute names to value domains). Each signature maps all identifiers not shown onto UNBOUND. Also for notational convenience, we assume the natural mapping from attribute names onto attribute values for each tuple (e.g., (ename → "Phil", ssn → 250861414) is the first tuple shown at transaction 4).

class	*signature*	*state*
⟨(ROLLBACK, 2, 6),	⟨((sname → string, ssn → integer), 2),	⟨(∅, 2),
		({("Phil", 250861414), ("Linda", 147894290), ("Ralph", 459326889)}, 4),
	((sname → string, class → string), 5),	({("Phil", "junior"), ("Linda", "senior"), ("Ralph", "senior")}, 5),
(SNAPSHOT, 8, –)	((ssn → integer, class → string), 8)	({(250861414, "junior"), (147894290, "senior"),
⟩	⟩	(459326889, "senior")}, 8) ⟩

The relation shown here was defined to be a rollback relation by transaction 2 and remained a rollback relation through transaction 6. While the relation was a rollback relation, all changes to its signature and state were recorded; its state was changed by transaction 4 and both its signature and its state were changed by transaction 5. Transaction 7 redefined the relation's class as SNAPSHOT and the relation was last updated as a snapshot relation by transaction 8. Only when a relation's current class is either rollback or temporal is the relation treated as an append-only relation. In all other cases, updates cause outdated information to be discarded. Hence, the lack of information about the relation's class, signature, and state before transaction 2 and at transaction 7 implies that the relation was either undefined or a snapshot or historical relation at those times. Note that this relation can be rolled back only to transactions 2 through 6. Also note that the last element in the class sequence defines the relation to be a snapshot relation from transaction 8 to the present. □

13

$$\mathcal{DATABASE\ STATE} = \mathcal{IDENTIFIER} \rightarrow \mathcal{RELATION}$$

A database state is a function that maps each identifier onto a relation. If an identifier I is mapped onto a relation whose current class is UNDEFINED, then I denotes an undefined relation. In the empty database state, all identifiers map onto undefined relations (i.e., ($\langle(\text{UNDEFINED},\ 0,\ -)\rangle,\ \langle\ \rangle,\ \langle\ \rangle$)).

$$\mathcal{DATABASE} = \mathcal{DATABASE\ STATE} \times \mathcal{TRANSACTION\ NUMBER}$$

A database is an ordered pair consisting of a database state and the transaction number assigned to the most recently committed transaction on the database state (i.e., the last transaction to cause a change to the database state).

2.3 A Semantic Type System for Expressions

Before specifying the semantics of the expressions defined syntactically in Section 2.1, we introduce a *semantic type system* for expressions. All syntactically correct expressions are not necessarily semantically correct. An expression is semantically correct, with respect to a database state and a command, only if its evaluation on the database state during the command's execution produces either a snapshot or a historical state. Also, if the expression contains a rollback operator, it must be consistent with the class and signature of the relation being rolled backed at the time of the transaction to which the relation is rolled back. Because the class and signature, as well as the state, of a relation are allowed to change over time, the semantic correctness of expressions also can vary over time. Hence, expressions that are semantically correct on a database state when one command is executed may not be semantically correct on the same database state when a subsequent command is executed (although the correctness of rollback operations to existing states will be unaffected by subsequent commands).

The semantic type system defined here allows us to do expression type-checking independent of expression evaluation. In Section 2.4, where we define the semantics of expressions, we will use the type system to restrict evaluation of expressions to semantically correct expressions only. Hence, any future implementation of the language can avoid the unnecessary cost associated with attempted evaluation of semantically incorrect expressions. The type system will also be used to define the semantics of commands so that commands whose execution would result in an incompatibility among a relation's class, signature, and state will never be executed. Also, separation of semantic type-checking and evaluation of expressions simplifies the formal definitions of the semantics of both expressions and commands. Note that while semantic type-checking and evaluation of some expressions (i.e., those expressions involving only constant expressions and rollback operators that roll back a relation prior to the query analysis time) can be done when a query is analyzed, most semantic type-checking and expression evaluation will have to be done when the query is executed.

Semantically correct expressions evaluate to either a single snapshot state or a single historical state. We define a snapshot state's type to be an ordered pair whose first component is SNAPSHOT

and whose second component is the state's signature. Similarly, we define a historical state's type to be an ordered pair whose first component is HISTORICAL and whose second component is the state's signature. A semantically correct expression's type is therefore the class and signature of the relation state resulting from the expression's evaluation and two expressions are said to be of the same type if and only if they evaluate to either snapshot or historical states on the attributes of the same signature.

We use the semantic function $\mathbf{T}$ to specify an expression's type. A semantic function is simply a function that maps a language construct onto its denotation or meaning. $\mathbf{T}$ defines an expression as a function that maps a database state and a transaction number onto either an ordered pair or TYPEERROR, depending on whether the expression is a semantically correct expression on the database state when a command in the transaction assigned the transaction number is executed. The ordered pair will have as its first component either SNAPSHOT or HISTORICAL and as its second component the signature of the relation state that the expression represents. Hence, $\mathbf{T}$ defines the type denotation of expressions.

$$\mathbf{T} : \mathcal{EXPRESSION} \rightarrow [\,[\,\mathcal{DATABASE\ STATE} \times \mathcal{TRANSACTION\ NUMBER}\,] \rightarrow$$
$$[\,[\,\{\text{SNAPSHOT, HISTORICAL}\} \times$$
$$\mathcal{RELATION\ SIGNATURE}\,] + \{\text{TYPEERROR}\}\,]\,]$$

The result of type-checking a syntactically correct expression is the class and signature of the relation state that the expression represents if the expression is semantically correct and an error if the expression is semantically incorrect. An expression's type may depend on a database state's contents. The type of an expression involving a rollback operator also depends on the transaction number of the transaction in which the command containing the expression occurs. Hence, a database state and transaction number together define the environment in which type-checking is performed.

Before defining the semantic function $\mathbf{T}$, we describe informally several functions used in its definition. Formal definitions for these auxiliary functions appear elsewhere [McKenzie 1988].

$\mathbf{N}$ is a semantic function that maps the syntactic category $\mathcal{NUMERAL}$ of decimal numerals into the semantic domain $\mathcal{INTEGER}$ of integers.

$\mathbf{S}$ is a semantic function that maps each alphanumeric representation of a snapshot state in the syntactic category $\mathcal{S\text{--}STATE}$ onto its corresponding snapshot state in the semantic domain $\mathcal{SNAPSHOT\ STATE}$, if it denotes a valid snapshot state on a given signature. Otherwise, $\mathbf{S}$ maps the snapshot state onto ERROR.

$\mathbf{VALIDF}$ is a semantic function that maps the alphanumeric representation of a boolean predicate in the syntactic category $\mathcal{SIGMA\ EXPRESSION}$ onto the boolean value TRUE or FALSE, to indicate whether the predicate is a valid boolean predicate for the selection operator σ (or $\hat{\sigma}$) and a given signature.

VALIDX is a semantic function that maps the alphanumeric representation of a list of identifiers in the syntactic category $\mathcal{IDENTIFIER}$ $\mathcal{LIST}$ onto the boolean value TRUE or FALSE, to indicate whether the identifiers denote a valid subset of the attributes in a given signature.

X is a semantic function that maps the alphanumeric representation of a list of identifiers in the syntactic category $\mathcal{IDENTIFIER}$ $\mathcal{LIST}$ onto an element in $\wp\,(\mathcal{INDENTIFIER})$, the power set of $\mathcal{IDENTIFIER}$, if the identifiers denote a valid subset of the attributes in a given signature. Otherwise, **X** maps the list onto ERROR.

Y is a semantic function that maps each character string in the syntactic category $\mathcal{CLASS}$ onto the relation class that it denotes in the semantic domain $\mathcal{RELATION}$ $\mathcal{CLASS}$.

Z is a semantic function that maps each alphanumeric representation of a relational signature in the syntactic category $\mathcal{SIGNATURE}$ onto its corresponding relational signature in the semantic domain $\mathcal{RELATION}$ $\mathcal{SIGNATURE}$.

FindClass maps a relation onto the class component of the element in the relation's class sequence whose first transaction-number component is less than or equal to a given transaction number and whose second transaction-number component is greater than or equal to the transaction number. If no such element exists in the sequence, then *FindClass* returns ERROR.

FindSignature maps a relation onto the signature component of the element in the relation's signature sequence having the largest transaction-number component less than or equal to a given transaction number, if *FindClass* does not return an error for the same transaction number. If *FindClass* returns an error or no such element exists in the sequence, then *FindSignature* returns ERROR.

LastClass maps a relation onto the class component of the last element in the relation's class sequence. If the sequence is empty, *LastClass* returns ERROR.

LastSignature maps a relation onto the signature component of the last element in the relation's signature sequence. If the relation's signature sequence is empty, *LastSignature* returns ERROR.

We now define formally the semantic function **T** for each kind of expression. For this and later definitions of semantic functions, let e be the number of value domains $\mathcal{D}_u$, $1 \leq u \leq e$, and let

d range over the domain $\mathcal{DATABASE}$ $\mathcal{STATE}$,
z, z_1, and z_2 range over the domain $\mathcal{RELATION}$ $\mathcal{SIGNATURE}$, and
tn range over the domain $\mathcal{TRANSACTION}$ $\mathcal{NUMBER}$.

$$\mathbf{T}[\![\,[\texttt{snapshot},\, Z,\, S\,]\,]\!]\,(d,\, tn) \;=\; \begin{array}{ll} \text{if} & (\mathbf{Z}[\![Z]\!] \neq \textsc{error} \,\wedge\, \mathbf{S}[\![S]\!]\,\mathbf{Z}[\![Z]\!] \neq \textsc{error}) \\[4pt] \text{then} & (\textsc{snapshot},\, \mathbf{Z}[\![Z]\!]) \\[4pt] \text{else} & \textsc{typeerror} \end{array}$$

If a constant expression represents a snapshot state on a signature, the expression's type is the ordered pair whose first component is SNAPSHOT and whose second component is the snapshot state's signature. Otherwise, evaluation of the expression's type results in an error.

EXAMPLE. For this and later examples in Section 2, assume that we are given the database (DS, 8) where the database state DS maps the identifier R1 onto the relation shown in the example on page 13.

$$\begin{aligned} \mathbf{T}[\![\,[\texttt{snapshot, (sname:string, class:string), (sname:"Phil", class:"junior"),} \\ \texttt{(sname:"Linda", class:"senior"),} \\ \texttt{(sname:"Ralph", class:"senior")]} \\ ]\!]\,(\text{DS}, 9) = (\textsc{snapshot},\, (\texttt{sname} \rightarrow \text{string},\, \texttt{class} \rightarrow \text{string})) \end{aligned}$$

Here we assume that type-checking is being performed as part of transaction 9. Note, however, that the database state is not consulted to determine the constant expression's type; the expression's type is independent of the database state. Actually, the only expressions whose type depends directly on the database state are identifiers and expressions involving the rollback operators. □

Evaluation of a snapshot constant's type produces an error if and only if the expression does not represent a snapshot state on a signature. As we will see in Section 2.4, evaluation of a constant expression's type produces an error under exactly the same conditions that evaluation of the expression produces an error. This relationship between a constant expression's type and value is both a necessary and a sufficient condition to ensure that the evaluation of any expression will result in an error when evaluation of the expression's type results in an error.

17

$$\mathbf{T}[\![I]\!]\,(d,\,tn) \;=\; \text{if} \qquad (LastClass(d(I)) = \textsc{snapshot}$$

$$\lor\, LastClass(d(I)) = \textsc{rollback})$$

$$\text{then} \quad (\textsc{snapshot},\, LastSignature(d(I)))$$

$$\text{else if} \;\; (LastClass(d(I)) = \textsc{historical}$$

$$\lor\, LastClass(d(I)) = \textsc{temporal})$$

$$\text{then} \quad (\textsc{historical},\, LastSignature(d(I)))$$

$$\text{else} \quad \textsc{typeerror}$$

where the notation $d(I)$ stands for the relation denoted by the identifier I in the database state d. The type of an expression I is the ordered pair whose first component is $\textsc{snapshot}$ if I's current class is either snapshot or rollback and $\textsc{historical}$ if its current class is either historical or temporal. The ordered pair's second component is always I's current signature. An error occurs if the relation is currently undefined.

EXAMPLE.

$$\mathbf{T}[\![\texttt{R1}]\!]\,(\text{DS}, 9) \;=\; (\textsc{snapshot},\, (\texttt{ssn} \to \text{integer},\, \texttt{class} \to \text{string}))$$

$\square$

$$\mathbf{T}[\![E_1 \cup E_2]\!]\,(d,\,tn) \;=\; \text{if} \quad \mathbf{T}[\![E_1]\!]\,(d,\,tn) = \mathbf{T}[\![E_2]\!]\,(d,\,tn) = (\textsc{snapshot},\, z)$$

$$\text{then} \;\; \mathbf{T}[\![E_1]\!]\,(d,\,tn)$$

$$\text{else} \quad \textsc{typeerror}$$

$$\mathbf{T}[\![E_1 - E_2]\!]\,(d,\,tn) \;=\; \text{if} \quad \mathbf{T}[\![E_1]\!]\,(d,\,tn) = \mathbf{T}[\![E_2]\!]\,(d,\,tn) = (\textsc{snapshot},\, z)$$

$$\text{then} \;\; \mathbf{T}[\![E_1]\!]\,(d,\,tn)$$

$$\text{else} \quad \textsc{typeerror}$$

$\mathbf{T}[\![E_1 \times E_2]\!](d,\ tn)\ =$

if $\quad (\mathbf{T}[\![E_1]\!](d,\ tn) = (\textsc{snapshot},\ z_1)\ \wedge\ \mathbf{T}[\![E_2]\!](d,\ tn) = (\textsc{snapshot},\ z_2)$

$\quad\quad \wedge\ \forall I,\ I \in \mathcal{IDENTIFIER},\ (z_1(I) = \textsc{unbound} \vee z_2(I) = \textsc{unbound}))$

then $\quad (\textsc{snapshot}, \{(I,\ \mathcal{D}_u) \mid 1 \leq u \leq e \wedge ((I,\ \mathcal{D}_u) \in z_1 \vee (I,\ \mathcal{D}_u) \in z_2)\}$

$\quad\quad\quad \cup \{(I,\ \textsc{unbound}) \mid I \in \mathcal{IDENTIFIER} \wedge (I,\ \textsc{unbound}) \in z_1$

$\quad\quad\quad\quad\quad \wedge (I,\ \textsc{unbound}) \in z_2\})$

else $\quad \textsc{typeerror}$

$\mathbf{T}[\![\boldsymbol{\pi}\, X\, (E)]\!](d,\ tn)\ =$

if $\quad (\mathbf{T}[\![E]\!](d,\ tn) = (\textsc{snapshot},\ z)\ \wedge\ \mathbf{VALIDX}[\![X]\!]\, z)$

then $\quad (\textsc{snapshot},\ \{(I,\ \mathcal{D}_u) \mid I \in \mathbf{X}[\![X]\!]\, z \wedge 1 \leq u \leq e \wedge (I,\ \mathcal{D}_u) \in z\}$

$\quad\quad\quad \cup \{(I,\ \textsc{unbound}) \mid I \notin \mathbf{X}[\![X]\!]\, z \wedge I \in \mathcal{IDENTIFIER}\})$

else $\quad \textsc{typeerror}$

$\mathbf{T}[\![\boldsymbol{\sigma}\, F\, (E)]\!](d,\ tn)\ =$ if $\quad (\mathbf{T}[\![E]\!](d,\ tn) = (\textsc{snapshot},\ z)\ \wedge\ \mathbf{VALIDF}[\![F]\!]\, z)$

$\quad\quad\quad\quad\quad\quad\quad$ then $\quad \mathbf{T}[\![E]\!](d,\ tn)$

$\quad\quad\quad\quad\quad\quad\quad$ else $\quad \textsc{typeerror}$

The type of an expression involving one of the five basic snapshot operators is an ordered pair whose first component is SNAPSHOT and whose second component is the signature of the relation state produced when the expression is evaluated, if two conditions are met. The first component of the type of all subexpressions must be SNAPSHOT and the second component of the type of all subexpressions must be a signature satisfying any restrictions placed on the signatures of relation states in corresponding expressions in the snapshot algebra. For example, our definitions of union and difference require that the signatures for E_1 and E_2 be identical while our definition of cartesian product requires that the attributes defined by the signatures for E_1 and E_2 be disjoint. (Note that we can eliminate this last restriction and effectively allow the cartesian product of snapshot states on arbitrary signatures through the introduction of a simple attribute renaming operator [Maier 1983] into the language.) If either condition is not met, evaluation of the expression's type results in an error.

$$\mathbf{T}[\![\boldsymbol{\rho}(I\,,\,N)]\!]\,(d,\,tn) \;=\; \text{if} \quad \mathbf{N}[\![N]\!] < tn \,\wedge\, FindClass(d(I),\,\mathbf{N}[\![N]\!]) = \text{ROLLBACK}$$

$$\text{then} \quad (\text{SNAPSHOT},\, FindSignature(d(I),\,\mathbf{N}[\![N]\!]))$$

$$\text{else} \quad \text{TYPEERROR}$$

A rollback expression's type is the ordered pair whose first component is SNAPSHOT and whose second component is the signature of the relation denoted by I when transaction $\mathbf{N}[\![N]\!]$ was processed, if the relation was a rollback relation at that time. Otherwise, evaluation of the expression's type results in an error. Because we assume sequential transaction processing, tn is the transaction number of the one active transaction and all committed transactions have transaction numbers less than tn. Hence, we allow rollback only to committed transactions.

EXAMPLES.

$$\mathbf{T}[\![\boldsymbol{\rho}(\texttt{R1},\texttt{4})]\!]\,(\text{DS},9) \;=\; (\text{SNAPSHOT},\, (\textbf{sname} \rightarrow \text{string},\, \textbf{ssn} \rightarrow \text{integer}))$$

$$\mathbf{T}[\![\boldsymbol{\pi}(\texttt{sname})(\boldsymbol{\rho}(\texttt{R1},\texttt{4}))]\!]\,(\text{DS},9) \;=\; (\text{SNAPSHOT},\, (\textbf{sname} \rightarrow \text{string}))$$

$$\mathbf{T}[\![\boldsymbol{\pi}(\texttt{sname})(\boldsymbol{\rho}(\texttt{R1},\texttt{4}))\times[\texttt{snapshot},\,(\texttt{course:string}),\,(\texttt{course:"English"})]$$
$$]\!]\,(\text{DS},9) = (\text{SNAPSHOT},\, (\textbf{sname} \rightarrow \text{string},\, \textbf{course} \rightarrow \text{string}))$$

$\square$

The semantic function $\mathbf{T}$ for expressions involving historical operators follows directly. The type denotations for these expressions are identical to those for expressions involving snapshot operators, except that HISTORICAL and TEMPORAL are substituted for SNAPSHOT and ROLLBACK, respectively.

Finally, we present the definition of the semantic function $\mathbf{T}$ for the last expression construct, which is used to group subexpressions.

$$\mathbf{T}[\![(E)]\!]\,(d,\,tn) \;=\; \mathbf{T}[\![E]\!]\,(d,\,tn)$$

2.4   Expressions

The semantic function $\mathbf{E}$ defines the denotation of expressions. $\mathbf{E}$ defines an expression as a function that maps a database state and a transaction number onto either a snapshot state (i.e., an element of the $\mathcal{SNAPSHOT\ STATE}$ semantic domain), a historical state (i.e., an element of the $\mathcal{HISTORICAL\ STATE}$ semantic domain), or ERROR.

$$\mathbf{E} : \mathcal{EXPRESSION} \rightarrow [\,[\,\mathcal{DATABASE\ STATE} \times \mathcal{TRANSACTION\ NUMBER}\,] \rightarrow$$
$$[\,\mathcal{SNAPSHOT\ STATE} + \mathcal{HISTORICAL\ STATE} + \{\text{ERROR}\}\,]\,]$$

If an expression is a semantically correct expression on a database state, expression evaluation on the database state produces either a snapshot state or a historical state. Otherwise, expression evaluation produces an error. The environment for expression evaluation, a database state and the transaction number of the active transaction, is the same as that for expression type-checking. Note that expression evaluation has no side-effect; it leaves the database state unchanged.

Before defining the semantic function $\mathbf{E}$, we describe informally additional auxiliary functions used in $\mathbf{E}$'s definition. Formal definitions appear elsewhere [McKenzie 1988].

FindState maps a relation onto the state component of the element in the relation's state sequence having the largest transaction-number component less than or equal to a given transaction number, if *FindClass* does not return an error for the same transaction number. If *FindClass* returns an error or no such element exists in the sequence, then *FindState* returns ERROR.

LastState maps a relation onto the state component of the last element in the relation's state sequence. If the relation's state sequence is empty, *LastState* returns ERROR.

We now define formally the semantic function $\mathbf{E}$ for each kind of expression allowed in the language.

$$\mathbf{E}[\![\,[\text{snapshot}, Z, S\,]\,]\!]\,(d, tn) \;=\; \text{if} \quad \mathbf{T}[\![\,[\text{snapshot}, Z, S\,]\,]\!]\,(d, tn) \neq \text{TYPEERROR}$$
$$\text{then} \quad \mathbf{S}[\![S]\!]\,\mathbf{Z}[\![Z]\!]$$
$$\text{else} \quad \text{ERROR}$$

EXAMPLE.

$\mathbf{E}[\![\,$[snapshot, (sname:string, class:string), (sname:"Phil", class:"junior"),
$\qquad\qquad\qquad\qquad\qquad$(sname:"Linda", class:"senior"),
$\qquad\qquad\qquad\qquad\qquad$(sname:"Ralph", class:"senior")]
$\quad]\!]\,(\text{DS}, 9) = \{\,(\text{"Phil"}, \text{"junior"}), (\text{"Linda"}, \text{"senior"}), (\text{"Ralph"}, \text{"senior"})\,\}$

$\square$

$$\mathbf{E}[\![I]\!]\,(d, tn) \;=\; \text{if } \mathbf{T}[\![I]\!]\,(d, tn) \neq \text{TYPEERROR then } \textit{LastState}(d(I)) \text{ else ERROR}$$

An identifier expression, if semantically correct, always evaluates to the current state of the relation denoted by I.

EXAMPLE.

$\mathbf{E}[\![\texttt{R1}]\!] \, (\text{DS}, 9) \; = \{ \, (250861414, \text{ "junior"}), \, (147894290, \text{ "senior"}), \, (459326889, \text{ "senior"}) \, \}$

$\square$

$$
\begin{aligned}
\mathbf{E}[\![E_1 \cup E_2]\!] \, (d, \, tn) \; = \; \text{if} \quad & \mathbf{T}[\![E_1 \cup E_2]\!] \, (d, \, tn) \neq \text{TYPEERROR} \\
\text{then} \quad & \mathbf{E}[\![E_1]\!] \, (d, \, tn) \cup \mathbf{E}[\![E_2]\!] \, (d, \, tn) \\
\text{else} \quad & \text{ERROR}
\end{aligned}
$$

The definitions of $\mathbf{E}$ for the other four snapshot operators are analogous to that for the union operator. For each of these operators, the denotation of a semantically correct expression containing the operator is defined as the standard snapshot operator over the denotation of the argument(s) to that operator.

$$
\begin{aligned}
\mathbf{E}[\![\rho(I, \, N)]\!] \, (d, \, tn) \; = \; \text{if} \quad & \mathbf{T}[\![\rho(I, \, N)]\!] \, (d, \, tn) \neq \text{TYPEERROR} \\
\text{then} \quad & FindState(d(I), \, \mathbf{N}[\![N]\!]) \\
\text{else} \quad & \text{ERROR}
\end{aligned}
$$

A semantically correct rollback expression evaluates to the snapshot state of the relation denoted by I at the time of transaction $\mathbf{N}[\![N]\!]$. The rollback operator always rolls a relation backward, but never forward, in time. Because transactions always update the database as they are executed, it is impossible to roll a relation forward in time. Although relations can't be rolled forward in time, our orthogonal treatment of valid and transaction time provides support for both *retroactive* changes and *postactive* changes (i.e., changes that will occur in the future) [Snodgrass & Ahn 1985]. Recall from the definition of the semantic function $\mathbf{T}$ that a rollback expression is semantically correct only if the relation was a rollback relation when the transaction was processed.

EXAMPLES.

$\mathbf{E}[\![\rho(\texttt{R1,4})]\!](\text{DS}, 9) =$

     { ("Phil", 250861414), ("Linda", 147894290), ("Ralph", 459326889) }

$\mathbf{E}[\![\pi(\texttt{sname})(\rho(\texttt{R1,4}))]\!](\text{DS}, 9) =$ { ("Phil"), ("Linda"), ("Ralph") }

$\mathbf{E}[\![\pi(\texttt{sname})(\rho(\texttt{R1,4})) \times [\texttt{ snapshot, (course:string), (course:"English")}]$
     $]\!](\text{DS}, 9) =$ { ("Phil", "English"), ("Linda", "English"), ("Ralph", "English") }

                                                          □

The semantic function **E** for expressions involving historical operators follows directly. The denotations for these expressions are identical to those for expressions involving snapshot operators, except that HISTORICAL and TEMPORAL are substituted for SNAPSHOT and ROLLBACK, respectively.

We now present the definition of the semantic function **E** for the expression construct that groups subexpressions.

$\mathbf{E}[\![(E)]\!](d, tn) = \mathbf{E}[\![E]\!](d, tn)$

2.5 Commands

The semantic function **C** defines the denotation of commands defined syntactically in Section 2.1. **C** defines a command as a function that maps a database state and a transaction number onto a database state and a status code. Execution of a semantically correct command produces a new database state and the status code OK, indicating that the command was successfully executed. Execution of a semantically incorrect command produces the original database state unchanged and the status code ERROR, indicating that the command could not be executed.

$\mathbf{C} : \mathcal{COMMAND} \to [[\mathcal{DATABASE\ STATE} \times \mathcal{TRANSACTION\ NUMBER}] \to$

$[\mathcal{DATABASE\ STATE} \times \{\text{OK, ERROR}\}]]$

The environment for command execution is the same as that for expression type-checking and evaluation, a database state and the transaction number of the active transaction (i.e., the transaction in which the command being executed occurs). A command produces a new database state from the given database state by changing a relation.

We use semantic type-checking of expressions in the definition of **C** to restrict evaluation of expressions to semantically correct expressions only. We also incorporate error-checking, based on the type system for expressions, into **C**'s definition to guarantee consistency among a relation's class, signature, and state following update. Error-checking ensures that commands actually change relations only when the change would result in a relation with compatible class, signature, and state. Commands whose execution would result in an inconsistency among a relation's class, signature, and state are effectively ignored (i.e, they do not alter the database state).

Before defining the semantic function **C**, we describe informally several additional auxiliary functions used in its definition. The formal definition for $MSoT$ appears in Appendix B. Formal definitions for the other functions appear in [McKenzie 1988].

$\mathbf{Y}'$ is the same as the semantic function **Y** with the exception that it maps the special symbol $*$ onto a relation's current class.

$\mathbf{Z}'$ is the same as the semantic function **Z** with the exception that it maps the special symbol $*$ onto a relation's current signature.

Consistent is a boolean function that determines whether a class and signature are consistent with an expression's type.

$MSoT$ (*M*odified *S*tart *o*f *T*ransaction) is a function that maps a relation and a transaction number onto the history of the relation as a rollback or temporal relation prior to the start of the transaction assigned the transaction number. We refer to this history as the relation's MSoT for that transaction. The significance of $MSoT$ will become apparent when we discuss multiple-command transactions.

EXAMPLE. Again assume, as in earlier examples, that we are given the database (DS, 8) where the database state component maps the identifier R1 onto the relation shown in the example on page 13.

$MSoT(\text{R1}, 9) =$

class	signature	state
$\langle(\text{ROLLBACK}, 2, 6)$	$\langle((\texttt{sname} \rightarrow \text{string},$ $\quad \texttt{ssn} \rightarrow \text{integer}), 2),$	$\langle(\emptyset, 2),$
		$(\{(\text{``Phil''}, 250861414),$   $\quad (\text{``Linda''}, 147894290),$ $\quad (\text{``Ralph''}, 459326889))\}, 4),$
	$((\texttt{sname} \rightarrow \text{integer},$ $\quad \texttt{class} \rightarrow \text{string}), 5)$	$(\{(\text{``Phil''}, \text{``junior''}),$ $\quad (\text{``Linda''}, \text{``senior''}),$
$\rangle$	$\rangle$	$\quad (\text{``Ralph''}, \text{``senior''})\}, 5) \quad \rangle$

In this example, $MSoT$ retains R1's history as a rollback relation prior to transaction 9. Although

`R1`'s current class, signature, and state were recorded before the start of transaction 9, they have been discarded because they are not part of `R1`'s history as a rollback relation. If, however, the last element in `R1`'s class sequence had been (ROLLBACK, 8, –), then `R1`'s current class, signature, and state also would have been retained. In this case, *MSoT* simply would have changed the second transaction-number component of the last element in `R1`'s class sequence to 8 to indicate that the resulting relation only records `R1`'s history as a rollback relation through transaction 8. If `R1` had never been a rollback or temporal relation, then *MSoT* would have mapped `R1` onto $(\langle\,\rangle, \langle\,\rangle, \langle\,\rangle)$. □

Expand replaces the second transaction-number component in the last element of a relation's MSoT class sequence with the special element "–". *Expand* has the effect of making the length of the interval for the class component of this element dynamic, extending to the present.

NewSignature maps a relation's MSoT and a (signature, transaction number) pair onto the empty sequence, if the signature in the last element of the relation's MSoT signature sequence is equal to the signature in the (signature, transaction number) pair, or a one-element sequence containing the (signature, transaction number) pair, otherwise.

NewState maps a relation's MSoT, a (relation state, transaction number) pair, and a (class, signature) pair onto the empty sequence, if the class and signature in the last elements of the relation's MSoT class and signature sequences are consistent with the (class, signature) pair and the state in the last element of the relation's MSoT state sequence is equal to the relation state in the (relation state, transaction number) pair, or a one-element sequence containing the (relation state, transaction number) pair, otherwise.

We define formally the semantics of commands using the same approach we used to define the semantics of expressions. We define the semantic function **C** for each kind of command allowed in the language. In each of the following definitions, the predicate specifies the conditions under which the command is executed. If these conditions hold, a new database state is produced and the status code OK is returned; otherwise, the database state is left unchanged and the status code ERROR is returned. The conditions specified in each definition are both necessary and sufficient to ensure that only semantically correct expressions are evaluated and that the class, signature, and state of each relation in the database state following execution of the command are consistent. In all five definitions we assume that if a_1, a_2, a_3, b_1, b_2, and b_3 are all sequences, then $(a_1, a_2, a_3) \,\|_3\, (b_1, b_2, b_3)$ denotes the triple $(a_1 \,\|\, b_1, a_2 \,\|\, b_2, a_3 \,\|\, b_3)$, where "$\|$" is the concatenation operator on sequences. Also, the notation $d\,[\,r\,/\,I\,]$ stands for a new database state that differs from the database state d only in that it maps the identifier I onto the relation r.

2.5.1   Defining a Relation

The `define_relation` command assigns to a relation, whose current class is UNDEFINED, a new class and signature and the empty relation state consistent with the new class. The assignment becomes effective when the transaction in which the command occurs is committed. The changes that the command makes to the relation to effect this assignment depend on the relation's current class; the last class, signature, and state, if any, in the relation's MSoT for the transaction in

Table 1: Define Relation Command

which the command occurs; and whether the new class is a single-state class (i.e., SNAPSHOT or HISTORICAL) or a multi-state class (i.e., ROLLBACK or TEMPORAL). We hereafter refer to the last class, signature, and state in a relation's MSoT, if present, as the relation's MSoT class, signature, and state, respectively. The actions performed by the `define_relation` command, for all possible combinations of these variables, can be reduced to the three cases shown in Table 1.

If the relation's current class is UNDEFINED, the `define_relation` command replaces the relation with its MSoT, augmented to include the new class, signature, and state. If the new class represents a non-disjoint extension of the relation's MSoT class, the interval assigned the MSoT class is extended (i.e., made into a dynamically expanding interval by changing the second transaction-number component to "–") to include the transaction in which the command occurs. This case is limited to `define_relation` commands in multiple-command transactions, which we discuss in Section 2.5.5. Otherwise, the new class is appended to the MSoT class sequence. In either case, a new signature (state) is added to the MSoT signature (state) sequence only if it differs from the MSoT signature (state). If the relation's current class is other than UNDEFINED, the command encounters an error condition and leaves the relation unchanged.

The formal definition of `define_relation` follows directly from Table 1.

$\mathbf{C}[\![\texttt{define\_relation}(I\texttt{,}\ Y\texttt{,}\ Z)]\!](d,\ tn)\ =$

$\qquad$ if $\qquad (M = MSoT(d(I),\ tn) \wedge LastClass(d(I)) = \text{UNDEFINED}$

$\qquad\qquad\qquad \wedge\ \mathbf{Y}[\![Y]\!] \neq \text{ERROR} \wedge \mathbf{Z}[\![Z]\!] \neq \text{ERROR})$

$\qquad$ then if $\qquad FindClass\,(M,\ tn-1) = \mathbf{Y}[\![Y]\!]$

$\qquad\qquad$ then $(d\,[\,(Expand(M) \,\|_3\, (\langle\ \rangle,\ NewSignature(M,\ (\mathbf{Z}[\![Z]\!],\ tn)),$

$\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad NewState(M,\ (\emptyset,\ tn),\ (\mathbf{Y}[\![Y]\!],\ \mathbf{Z}[\![Z]\!])))$

$\qquad\qquad\qquad\qquad)/I\,],\ \text{OK})$

$\qquad\qquad$ else $(d\,[\,(M \,\|_3\, (\langle(\mathbf{Y}[\![Y]\!],\ tn,\ -)\rangle,\ NewSignature(M,\ (\mathbf{Z}[\![Z]\!],\ tn)),$

$\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad NewState(M,\ (\emptyset,\ tn),\ (\mathbf{Y}[\![Y]\!],\ \mathbf{Z}[\![Z]\!])))$

$\qquad\qquad\qquad\qquad)/I\,],\ \text{OK})$

$\qquad$ else $\quad (d,\ \text{ERROR})$

where M ranges over the domain $\mathcal{RELATION} + \{\,(\langle\ \rangle,\ \langle\ \rangle,\ \langle\ \rangle)\}.$

EXAMPLES. In these, and later examples, we show the result of executing a sequence of commands, starting with the database (DS, 8). We assume that each command corresponds to a single-command transaction that commits. For simplicity, we always refer to the current database state as DS, although it changes with each command's execution (i.e., transaction's commitment). We also restrict the commands to the relations denoted by the identifiers R1, R2, and R3 and show only the portion of the database state changed by each command's execution. We assume that DS maps the identifiers R2 and R3 onto the following relations.

	class	*signature*	*state*
R2→	$\langle$(ROLLBACK, 1, 5),	$\langle$(($\textbf{ename}\to$ string, $\textbf{ssn}\to$ integer), 1)	$\langle$($\emptyset$, 1), ({("Phil", 250861414), ("Linda", 147894290), ("Ralph", 459326889)}, 3)
	(UNDEFINED, 6, $-$)$\rangle$	$\rangle$	$\rangle$

	class	*signature*	*state*
R3→	$\langle$(UNDEFINED, 0, $-$)$\rangle$	$\langle\rangle$	$\langle\rangle$

Note that a relation whose current class is UNDEFINED has neither a current signature nor a current state. The relation denoted by R2 has a MSoT signature (state), but not a current signature (state). The relation denoted by R3 has neither a MSoT signature (state) nor a current signature (state).

$\textbf{C}[\![$define_relation(R2, rollback, (ename:string, ssn:integer))$]\!]$(DS, 9)

	class	*signature*	*state*
R2→	$\langle$(ROLLBACK, 1, 5),	$\langle$(($\textbf{ename}\to$ string, $\textbf{ssn}\to$ integer), 1)	$\langle$($\emptyset$, 1), ({("Phil", 250861414), ("Linda", 147894290), ("Ralph", 459326889)}, 3),
	(ROLLBACK, 9, $-$) $\rangle$	$\rangle$	($\emptyset$, 9) $\rangle$

$\mathbf{C}[\![\texttt{define\_relation(R3, snapshot, (sname:string, class:string))}]\!](\text{DS}, 10)$

	class	signature	state
R3→	$\langle(\text{SNAPSHOT}, 10, -)$ $\rangle$	$\langle((\texttt{sname} \to \text{string},$ $\texttt{class} \to \text{string}), 10)\rangle$	$\langle(\emptyset, 10)$ $\rangle$

The first command makes the relation denoted by R2 a rollback relation over the attributes ename and ssn, effective when transaction 9 commits. Although the new class and the relation's MSoT class are equal, the intervals associated with the two are disjoint. Hence, the new class is appended to the relation's MSoT class sequence. The new signature is not appended to the relation's MSoT signature sequence because it is the same as the relation's MSoT signature. The new state, the empty set, differs from the relation's MSoT state. Hence, it is added to the relation's MSoT state sequence. The second command makes the relation denoted by R3 a snapshot relation over the attributes sname and class, effective when transaction 10 commits. Because the relation's MSoT at transaction 10 is $(\langle\,\rangle, \langle\,\rangle, \langle\,\rangle)$, the command transforms the relation's class, signature, and state sequences into single-element sequences containing the new class, signature, and state. Note that information about both relations when they were undefined has been discarded as it is not needed for rollback. □

2.5.2 Modifying a Relation

The modify_relation command assigns to a relation, whose current class is other than UNDEFINED, a new class, signature, and relation state. The assignment becomes effective when the transaction in which the command occurs is committed. The modify_relation command differs from the define_relation command in only three respects. First, the modify_relation command only updates a relation if its current class is not UNDEFINED, whereas the define_relation command does just the opposite. Second, the modify_relation command, unlike the define_relation command, allows the new class (signature) to be the relation's current class (signature). Third, the modify_relation command allows the new relation state to be the value of any semantically correct expression consistent with the new class and signature, whereas the define_relation command requires that the new state be the empty state consistent with the new class. Otherwise, the semantics of the two commands is the same. The actions performed by the modify_relation command are summarized in Table 2.

The formal definition of modify_relation follows directly from the above description of the command and Table 2.

Table 2: Modify Relation Command

$\mathbf{C}[\![\text{modify\_relation}(I,\, Y',\, Z',\, E)]\!]\,(d,\, tn)\ =$

    if $(M = MSoT(d(I),\, tn) \,\wedge\, \mathbf{T}[\![E]\!]\,(d,\, tn) \neq \text{ERROR} \,\wedge\, LastClass(d(I)) \neq \text{UNDEFINED}$

        $\wedge\, Consistent(\mathbf{Y}'[\![Y']\!]\,(d(I)),\, \mathbf{Z}'[\![Z']\!]\,(d(I)),\, \mathbf{T}[\![E]\!]\,(d,\, tn)))$

  then if $FindClass\,(M,\, tn - 1) = \mathbf{Y}'[\![Y']\!]\,(d(I))$

      then $(d\,[\,(Expand(M) \,\|_3\, (\langle\ \rangle,\ NewSignature(M,\, (\mathbf{Z}'[\![Z']\!]\,(d(I)),\, tn)),$

                             $NewState(M,\, (\mathbf{E}[\![E]\!]\,(d,\, tn),\, tn),\, \mathbf{T}[\![E]\!]\,(d,\, tn)))$

          $)/I\,],\ \text{OK})$

      else $(d\,[\,(M \,\|_3\, (\langle(\mathbf{Y}'[\![Y']\!]\,(d(I)),\, tn,\, -)\rangle,\ NewSignature(M,\, (\mathbf{Z}'[\![Z']\!]\,(d(I)),\, tn)),$

                             $NewState(M,\, (\mathbf{E}[\![E]\!]\,(d,\, tn),\, tn),\, \mathbf{T}[\![E]\!]\,(d,\, tn)))$

          $)/I\,],\ \text{OK})$

  else $(d,\ \text{ERROR})$

If a relation's current class is other than UNDEFINED, the `modify_relation` command replaces the relation with its MSoT, augmented to include the new class, signature, and state. If the relation's current class is UNDEFINED, the command encounters an error and leaves the relation unchanged.

EXAMPLES.

$\mathbf{C}[\![\texttt{modify\_relation(R2, *, *, } \boldsymbol{\rho}\texttt{(R2,5) - } \boldsymbol{\sigma}\,\texttt{ename="Ralph"}\,(\boldsymbol{\rho}\texttt{(R2,5)))}]\!]\,(\text{DS, 11})$

	class	signature	state
R2→	$\langle(\textsc{rollback}, 1, 5),$	$\langle((\texttt{ename} \rightarrow \text{string},$	$\langle(\emptyset, 1),$
		$\texttt{ssn} \rightarrow \text{integer}), 1)$	
			$(\{(\text{"Phil"}, 250861414),$
			$(\text{"Linda"}, 147894290),$
			$(\text{"Ralph"}, 459326889)\}, 3),$
	$(\textsc{rollback}, 9, -)$		$(\emptyset, 9),$
			$(\{(\text{"Phil"}, 250861414),$
	$\rangle$	$\rangle$	$(\text{"Linda"}, 147894290)\}, 11)\rangle$

$\mathbf{C}[\![\texttt{modify\_relation(R3, *, *, } \boldsymbol{\rho}\texttt{(R1,5))}]\!]\,(\text{DS, 12})$

	class	signature	state
R3→	$\langle(\textsc{snapshot}, 12, -)$	$\langle((\texttt{sname} \rightarrow \text{string},$	$\langle(\{(\text{"Phil"}, \text{"junior"}),$
		$\texttt{class} \rightarrow \text{string}), 12)$	$(\text{"Linda"}, \text{"senior"}),$
	$\rangle$	$\rangle$	$(\text{"Ralph"}, \text{"senior"})\}, 12)\rangle$

The first command changes the state of the relation denoted by R2 while the second command changes the state of the relation denoted by R3. The commands, however, do not change the class or signature of either relation. For the first command, the new class (i.e., R2's current class) is a non-disjoint extension of R2's MSoT class. Hence, the interval for R2's MSoT class is made into a dynamically expanding interval that includes transaction 11, but no new element is added to R2's MSoT class sequence. The new signature (i.e., R2's current signature) is the same as R2's MSoT signature, hence it is not added to R2's MSoT signature sequence. The new state differs from R2's MSoT state, hence it is appended to R2's MSoT state sequence. Because R3's MSoT at transaction 12 is still $(\langle\,\rangle, \langle\,\rangle, \langle\,\rangle)$, the second command transforms R3's class, signature, and state sequences into single-element sequences containing the new class (i.e., R3's current class), signature (i.e., R3's current signature), and state. Note that R2's state at transaction 9 through transaction 10 has been retained and remains accessible via the rollback operator ρ, but R3's state before transaction 12 has been discarded (i.e., physically deleted from the database state).

$$\mathbf{C}[\![\texttt{modify\_relation(R3, *, (sname:string, course:string),}$$
$$\texttt{π(sname)(R3)$\times$[snapshot, (course:string),}$$
$$\texttt{(course:"English")])}]\!](\text{DS}, 13)$$

	class	*signature*	*state*
R3→	$\langle(\textsc{snapshot}, 13, -)$ $\rangle$	$\langle((\textsf{sname} \rightarrow \text{string},$ $\textsf{course} \rightarrow \text{string}), 13)$ $\rangle$	$\langle(\{(\text{"Phil"}, \text{"English"}),$ $(\text{"Linda"}, \text{"English"}),$ $(\text{"Ralph"}, \text{"English"})\}, 13)\rangle$

This command changes R3's signature and state but leaves the relation's class unchanged. It illustrates two possible changes to a relation's signature, deletion of one attribute and addition of another attribute. Deletion of an attribute is usually expressed as a projection over the remaining attributes. Addition of an attribute requires that a value for the new attribute be determined for each tuple in the relation. Often, as in this example, a single default value is specified, which is then appended to each tuple. Note again that R3's state before transaction 13 has been discarded. $\quad\square$

The modify_relation command has several noteworthy properties. First, the command supports all update operations on a relation's state. Append is accommodated by an expression E, generally containing a union operator, that evaluates to a snapshot or historical state containing all the tuples in a relation's current state plus one or more tuples not in the relation's current state. Delete is accommodated by an expression E, generally containing a difference operator, that evaluates to a snapshot or historical state containing only a proper subset of the tuples in a relation's current state. Replace is accommodated by an expression E that evaluates to a snapshot or historical state that differs from a relation's current state only in the attribute values of one or more tuples.

Second, the modify_relation command ensures that a relation's class, signature, and state are consistent following update. The command changes a relation's state only if the new state is consistent with the relation's class and signature. Whenever the command changes a relation's signature, it also changes the relation's state to ensure consistency among the relation's class, signature, and state [Navathe & Fry 1976]. Likewise, whenever the command changes a relation's class, it also updates the relation's state, if necessary, to ensure consistency among the relation's class, signature, and state.

Finally, the modify_relation command always treats a relation's signature (state) sequence as an *append-only* sequence when the relation's current class is either rollback or temporal, but it does not automatically discard a relation's current signature (state) on update even if the relation's current class is snapshot or historical. If a relation's current class is a single-state class, the command discards the relation's current signature (state) on update only if the signature (state) is not part of the relation's history as a rollback or temporal relation.

2.5.3 Deleting a Relation

The command `destroy` assigns to a relation, whose current class is other than UNDEFINED, the new class UNDEFINED. It also deletes, either logically or physically, the relation's current signature and state.

$$\mathbf{C}[\![\texttt{destroy}(I)]\!](d,\,tn)\,=$$

$$\text{if} \quad M = MSoT(d(I),\,tn) \,\wedge\, LastClass(d(I)) \neq \textsc{undefined}$$

$$\text{then} \quad (d\,[\,(M\,\|_3\,(\langle(\textsc{undefined},\,tn,\,-)\rangle,\,\langle\,\rangle,\,\langle\,\rangle))/I\,],\,\textsc{ok})$$

$$\text{else} \quad (d,\,\textsc{error})$$

If the identifier I denotes a relation whose current class is other than UNDEFINED, the command simply appends the new class UNDEFINED to the relation's MSoT for the transaction in which the command occurs.

EXAMPLES.

$\mathbf{C}[\![\texttt{destroy(R2)}]\!](\text{DS},\,14)$

	class	*signature*	*state*
R2→	$\langle(\textsc{rollback},\,1,\,5),$	$\langle((\texttt{ename} \to \text{string},$	$\langle(\emptyset,\,1),$
		$\texttt{ssn} \to \text{integer}),\,1)$	
			$(\{(\text{"Phil"},\,250861414),$
			$(\text{"Linda"},\,147894290),$
			$(\text{"Ralph"},\,459326889)\},\,3),$
	$(\textsc{rollback},\,9,\,13),$		$(\emptyset,\,9),$
			$(\{(\text{"Phil"},\,250861414),$
			$(\text{"Linda"},\,147894290)\},\,11)$
	$(\textsc{undefined},\,14,\,-)\rangle$	$\rangle$	$\rangle$

$\mathbf{C}[\![\texttt{destroy(R3)}]\!](\text{DS},\,15)$

	class	*signature*	*state*
R3→	$\langle(\textsc{undefined},\,15,\,-)\rangle$	$\langle\,\rangle$	$\langle\,\rangle$

Because R2 denotes a relation whose current class is ROLLBACK, the first command uses the function *MSoT* to "close" the interval associated with the relation's current class. It then appends the element (UNDEFINED, 14, −) to R2's class sequence. These actions together have the effect of logically deleting R2's current signature and state when transaction 14 commits. Note, however, that this signature and state information is still accessible via the rollback operator ρ. The second command uses the function *MSoT* to physically delete R3's current class, signature, and state. No record of R3 as a snapshot relation is retained. □

It is important to observe from these, and previous, examples that signature and state information associated with a relation when its class was either snapshot or historical was transient. It was physically removed when it became outdated. Hence, the language is consistent with conventional relational DBMS's that discard out-of-date signature and state information (relation R3 illustrates this). However, signature and state information associated with a relation when its class was rollback or temporal is retained, ensuring later access to past states via the rollback operator. Definition of the rollback operator assumes access to a complete record of a relation's signature and state during intervals when the relation's class was either rollback or temporal.

2.5.4   Renaming a Relation

The command `rename_relation` binds a relation's current class, signature, and state to a new identifier.

$$\mathbf{C}[\![\texttt{rename\_relation}(I_1,\, I_2)]\!]\,(d,\, tn)\ =$$

$$\quad \text{if} \quad (LastClass(d(I_1)) \neq \text{UNDEFINED} \,\wedge\, LastClass(d(I_2)) = \text{UNDEFINED}$$

$$\quad\quad\quad \wedge\, \mathbf{Y}[\![Y]\!] = LastClass(d(I_1)) \,\wedge\, \mathbf{Z}[\![Z]\!] = LastSignature(d(I_1))$$

$$\quad\quad\quad \wedge\, \mathbf{C}[\![\texttt{define\_relation}(I_2,\, Y,\, Z)]\!]\,(d,\, tn)\ = (d',\, \text{OK})$$

$$\quad\quad\quad \wedge\, \mathbf{C}[\![\texttt{modify\_relation}(I_2,\texttt{*},\texttt{*}, I_1)]\!]\,(d',\, tn)\ = (d'',\, \text{OK})$$

$$\quad\quad\quad \wedge\, \mathbf{C}[\![\texttt{destroy}(I_1)]\!]\,(d'',\, tn)\ = (d''',\, \text{OK}))$$

$$\quad \text{then}\quad (d''',\, \text{OK})$$

$$\quad \text{else}\quad (d,\, \text{ERROR})$$

The `rename_relation` first assigns to the relation denoted by I_2 the current class and signature of the relation denoted by I_1. It then assigns to I_2 the current state of I_1. Finally, it assigns the class UNDEFINED to I_1 and deletes, either logically or physically, I_1's current signature and state. Note that the execution environments for `rename_relation`'s three subordinate commands, while containing different database states, contain the same transaction number. Hence, the changes to both I_1 and I_2 become effective when a single transaction commits.

EXAMPLE. Recall that R1 is the relation shown on page 13.

$\mathbf{C}[\![\texttt{rename\_relation(R1, R3)}]\!](\text{DS},\,16)$

	class	*signature*	*state*
R1→	$\langle(\text{ROLLBACK},\,2,\,6),$	$\langle((\texttt{sname} \to \text{string},$	$\langle(\emptyset,\,2),$
		$\quad\texttt{ssn} \to \text{integer}),\,2),$	
			$(\{(\text{``Phil''},\,250861414),$
			$\quad(\text{``Linda''},\,147894290),$
			$\quad(\text{``Ralph''},\,459326889)\},\,4),$
		$((\texttt{sname} \to \text{string},$	$(\{(\text{``Phil''},\,\text{``junior''}),$
		$\quad\texttt{class} \to \text{string}),\,5)$	$\quad(\text{``Linda''},\,\text{``senior''}),$
			$\quad(\text{``Ralph''},\,\text{``senior''})\},\,5)$
	$(\text{UNDEFINED},\,16,\,-)\rangle$	$\rangle$	$\rangle$

	class	*signature*	*state*
R3→	$\langle(\text{SNAPSHOT},\,16,\,-)$	$\langle((\texttt{ssn} \to \text{integer},$	$\langle(\{(250861414,\,\text{``junior''}),$
		$\quad\texttt{class} \to \text{string}),\,16)$	$\quad(147894290,\,\text{``senior''}),$
	$\rangle$	$\rangle$	$\quad(459326889,\,\text{``senior''})\},\,16)\rangle$

This command binds the current class, signature, and state of the relation denoted by R1 to the identifier R3. Hence, R3 becomes a snapshot relation when transaction 16 commits. The command also transforms R1 into an undefined relation, effective when transaction 16 commits. Because R1's current class, signature, and state are not part of the relation's history as either a rollback or temporal relation, they are physically deleted. $\qquad\square$

2.5.5   A Sequence of Commands

If two or more commands appear in sequence, the commands are executed sequentially. If a command executes without error, the next command is executed using the database state resulting from the previous command's execution. If all the commands execute without error, the commands are mapped onto the final database state and the status code OK. If, however, any command's execution causes an error, the remaining commands are not executed and the status code ERROR is returned.

$$\mathbf{C}[\![C_1,\,C_2]\!](d,\,tn) \;=\; \text{if } \mathbf{C}[\![C_1]\!](d,\,tn) = (d',\,\text{OK}) \text{ then } \mathbf{C}[\![C_2]\!](d',\,tn) \text{ else } (d,\,\text{ERROR})$$

Two or more commands appearing in sequence are all commands in the same transaction. Their execution environments have different database states but the same transaction number. Hence, if

the commands change the same relation only the last changes to the relation's class, signature, and state are recorded in the final database state. Recall that while a relation's new class, signature, and state may depend on its current class, signature, and state, all commands define the resulting relation in terms of the relation's modified start of transaction. Also, if the commands change several relations, all the changes become effective when the transaction commits.

EXAMPLES. In the previous examples, we assumed that the commands were all taken from single-command transactions. We now show the result of executing multiple commands from the same transaction. Recall from page 33 that R2 is currently undefined.

C⟦define_relation(R2, rollback, (ename:string, ssn:integer)),

  modify_relation(R2, *, *, ρ(R2,5)),

  modify_relation(R2, *, *, R2 - σ ename="Linda" (R2))⟧ (DS, 17)

	class	*signature*	*state*
R2→	⟨(ROLLBACK, 1, 5),	⟨((ename → string,	⟨(∅, 1),
		ssn → integer), 1)	
			({("Phil", 250861414),
			("Linda", 147894290),
			("Ralph", 459326889)}, 3),
	(ROLLBACK, 9, 13),		(∅, 9),
			({("Phil", 250861414),
			("Linda", 147894290)}, 11),
	(ROLLBACK, 17, −)		({("Phil", 250861414),
	⟩	⟩	("Ralph", 459326889)}, 17) ⟩

C⟦destroy(R2), destroy(R3)⟧ (DS, 18)

	class	signature	state
R2→	⟨(ROLLBACK, 1, 5),	⟨((ename → string, ssn → integer), 1)	⟨(∅, 1),
			({("Phil", 250861414), ("Linda", 147894290), ("Ralph", 459326889)}, 3),
	(ROLLBACK, 9, 13),		(∅, 9),
			({("Phil", 250861414), ("Linda", 147894290)}, 11),
	(ROLLBACK, 17, 17)		({("Phil", 250861414), ("Ralph", 459326889)}, 17)
	(UNDEFINED, 18, −)⟩	⟩	⟩

	class	signature	state
R3→	⟨(UNDEFINED, 18, −)⟩	⟨⟩	⟨⟩

In the first example, all three commands change R2. Yet, only the last changes to the relation's class, signature, and state are recorded in the database state. Although the first command defined R2 as a rollback relation and the other commands changed R2's state, only the final change in state is recorded. Hence, all the commands in a single transaction that change the same relation are treated as an atomic update operation. Note that temporary relations can be defined, modified, and then deleted within a transaction without their creation being recorded. In the second example, both R2 and R3 are deleted when transaction 18 commits. □

2.6   Programs

The semantic function **P** defines the denotation of programs in our language, where a program is a sequence of one or more transactions. Transactions, in turn, may be either single-command or multiple-command transactions. **P** defines a program as a function that maps a database onto a database and a status code. A program is the only language construct that changes a database. Execution of a transaction that commits produces a new database and the status code OK, while execution of a transaction that aborts produces the original database unchanged and the status code ERROR.

$$\mathbf{P} : \mathcal{PROGRAM} \to [\,\mathcal{DATABASE} \to [\,\mathcal{DATABASE} \times \{\text{OK, ERROR}\}\,]\,]$$

Note that the environments for command and program execution, although similar, are different. The environment for command execution is a database state and the transaction number of the active transaction. In contrast, the environment for program execution is a database, which is

an ordered pair consisting of a database state and the transaction number of the most recently committed transaction on that database state.

We now define formally the semantic function **P** for each kind of program.

$$\textbf{P}[\![\texttt{begin\_transaction}\ C\ \texttt{commit\_transaction}]\!]\,(d,\,tn)\ =$$
$$\quad \text{if}\ \textbf{C}[\![C]\!]\,(d,\,tn+1) = (d',\,\textsc{ok})\ \text{then}\ ((d',\,tn+1),\,\textsc{ok})\ \text{else}\ ((d,\,tn),\,\textsc{error})$$

Committed transactions represent transactions that commit if their commands all execute without error. If all the commands in a transaction execute without error, the transaction is committed. The database's database-state component is updated to record the changes that the commands make to relations, the database's transaction-number component is incremented to record the transaction number of this most recently committed transaction, and the status code $\textsc{ok}$ is produced. If any command's execution produces an error, the transaction is aborted. The database is left unchanged and the status code $\textsc{error}$ is produced. The database is valid independent of the status code.

$$\textbf{P}[\![\texttt{begin\_transaction}\ C\ \texttt{abort\_transaction}]\!]\,(d,\,tn)\ =\ ((d,\,tn),\,\textsc{ok})$$

Aborted transactions are transactions, which the user initiates, that for some reason, dictated either by the user or by the system, abort rather than commit. They do not change the database.

$$\textbf{P}[\![P_1;\,P_2]\!]\,(d,\,tn)\ =$$
$$\quad \text{if}\ \textbf{P}[\![P_1]\!]\,(d,\,tn) = ((d',\,tn'),\,\textsc{ok})\ \text{then}\ \textbf{P}[\![P_2]\!]\,(d',\,tn')\ \text{else}\ \textbf{P}[\![P_2]\!]\,(d,\,tn)$$

If a program contains multiple transactions, they are processed in sequence. If the first transaction commits and produces a new database, the second transaction is processed using the new database. Otherwise, the second transaction is processed using the original database.

Finally, we require that each arbitrary sequence of transactions representing a program map onto the database resulting from the execution of the transactions, in order, starting with the empty database. The empty database, (**EMPTY**, 0), is defined using the semantic function **EMPTY** : $\mathcal{IDENTIFIER} \rightarrow (\,\langle(\textsc{undefined},\,0,\,-)\rangle,\,\langle\,\rangle,\,\langle\,\rangle\,)$. Hence, the database-state component of the empty database is defined to be the function that maps all identifiers onto undefined relations; the transaction-number component of the empty database is defined to be 0. This requirement is both necessary and sufficient to ensure that the transaction-number components of elements in the class, signature, and state sequences of each relation in the database are strictly increasing. A database will always be the cumulative result of all the transactions that have been performed on it since it was created.

We now define the semantic function $\mathbf{P}'$ that maps a program onto the database resulting from the execution of the program's transactions, starting with the empty database.

$$\mathbf{P}' : \mathcal{PROGRAM} \rightarrow \mathcal{DATABASE}$$

$$\mathbf{P}'[\![P]\!] \;=\; \mathit{First}(\mathbf{P}[\![P]\!](\mathbf{EMPTY}, 0))$$

where *First* is the function that maps an ordered pair onto the first component of the ordered pair.

2.7 Language Properties

We now state, as theorems, four properties of our algebraic language for database query and update. Informal proofs of these theorems are given in Appendix A. The first property was stated initially as an objective in Section 1.

Theorem 1 *The language is a natural extension of the relational algebra for database query and update.*

By natural extension, we mean that our semantics subsumes the expressive power of the relational algebra for database query and update. Expressions in the language are a strict superset of those in the relational algebra. Also, if we restrict the class of all relations to UNDEFINED and SNAPSHOT, then a natural extension implies that (a) the signature and state sequences of a defined relation will have exactly one element each: the relation's current signature and state; (b) a new state always will be a function of the current signature and state of defined relations via the relational algebra semantics; and (c) deletion will correspond to physical deletion.

The second property argues that the semantics is minimal, in a specific sense. Other definitions of minimality, such as minimal redundancy or minimal space requirements, are more appropriate for the physical level, where actual data structures are implemented, than for the algebraic level.

Theorem 2 *The semantics of the language minimizes the number of elements in a relation's class, signature, and state sequence needed to record the relation's current class, signature, and state and its history as a rollback or temporal relation.*

The third property ensures that the language accommodates implementations that use WORM optical disk to store non-current class, signature, and state information, another objective of our extensions.

Theorem 3 *Transactions change only a relation's class, signature, and state current at the start of the transaction.*

3   Additional Aspects of the Rollback Operators

The rollback operators in our language are more powerful than suggested in the previous section, in several ways. First, the rollback operators, as defined, are restricted to the retrieval of a single snapshot or historical state from a named relation current at the time of a specified transaction. In reality, however, the rollback operators derive a single snapshot or historical state from one or more of the named relation's stored states rather than simply retrieving a single state. The rollback operators actually roll back a relation to the subsequence of the relation's state sequence corresponding to an interval of time of arbitrary length, if the relation's class and signature remained constant over that interval of time. The rollback operators return the single state composed of tuples from all the states in the specified subsequence of relation states (effectively, a relational union, either snapshot or historical, is performed). The rollback operators thus take two transaction times as arguments:

$$E \quad ::= \quad \rho(I, N, N) \mid \hat{\rho}(I, N, N)$$

Second, the rollback operators do not simply retrieve a snapshot or historical state from a named relation but rather an augmented version of that state. To the state's explicit attributes, defined in its signature, the rollback operators add new explicit attributes corresponding to the state's implicit time attributes (i.e., transaction times for snapshot states, transaction and valid times for historical states). The rollback operators' addition of these new attributes to the state's existing explicit attributes allows the user to display the values of the state's implicit time attributes without allowing direct access to the attributes themselves. These explicit values are considered to be in the domain of user-defined time. This behavior requires that the semantic function **T** compute a relational signature containing these additional attributes.

Third, the rollback operator ρ can be applied to temporal relations as well as rollback relations. If ρ rolls back a relation to a time when the relation's class was TEMPORAL, ρ will convert the relation's historical state current at that time into a corresponding snapshot state and return this new snapshot state. Likewise, the rollback operator $\hat{\rho}$ can be applied to rollback relations as well as temporal relations. If $\hat{\rho}$ rolls back a relation to a time when the relation's class was ROLLBACK, $\hat{\rho}$ will convert the relation's snapshot state current at that time into a corresponding historical state and return this new historical state.

While these extensions are conceptually straightforward, the notation required to define them formally is cumbersome and will not be presented.

4   Summary and Related Work

In summary, this paper has defined an algebraic language for database query and update that subsumes the relational algebra, can accommodate an arbitrary historical algebra, and supports both snapshot and historical rollback. The language also has a simple semantics and supports

scheme evolution. Only two additional operators, ρ and $\hat{\rho}$, were necessary. The additions required for transaction time did not compromise any of the useful properties of the (conventional) snapshot algebra. Type-checking was also introduced, freeing the encapsulated algebra from dealing with expressions not consistent with the (possibly time-varying) scheme.

The primary contribution is an algebraic means of supporting schema evolution in the context of general support for transaction time. As an algebraic language for database query and update, our language can serve as the underlying evaluation mechanism for queries and updates in a temporal data manipulation language that supports evolution of a database's contents and schema. It can also be used as the basis for proving various physical implementations of temporal database management systems correct. Our language also is compatible with efforts to add transaction time to the relational data model at both the user-interface and physical levels. At least three temporal query languages have been proposed that support rollback operations [Ariav 1986, Ben-Zvi 1982, Snodgrass 1987] and several studies have investigated efficient storage and access strategies for temporal databases [Ahn 1986A, Ahn 1986B, Ahn & Snodgrass 1986, Ahn & Snodgrass 1988, Lum et al. 1984, Rotem & Segev 1987, Shoshani & Kawagoe 1986, Thirumalai & Krishna 1988].

There have been two other attempts to incorporate both valid time and transaction time in an algebra. In BenZvi's proposal, valid time and transaction time were supported through the addition of implicit time attributes to each tuple in a relation [Ben-Zvi 1982]. The algebra was extended with the *Time-View* algebraic operator which takes a relation and two times as arguments and produces the subset of tuples in the relation valid at the first time (the valid time) as of the second time (the transaction time). The Time-View operator thus rolls back a relation to a transaction time but returns only a subset of the tuples in the relation at that transaction time (i.e., those tuples valid at some specified time). This restricted definition of the Time-View operator is tied inextricably to his particular handling of valid time. Our approach is compatible with any historical algebra. Gadia represents valid time and transaction time as two symmetrical dimmensions in a boolean algebra of multidimensional time stamps [Gadia & Yeung 1988]. He allows rollback operations on transaction time through a generalized restriction operator, which may be applied to any of a relation's time dimensions. He does not, however, address the problems of database update or schema evolution.

While a few authors have envisaged the benefits of a time-varying schema [Ariav 1986, Ben-Zvi 1982, Shiftan 1986, Woelk et al. 1986], only one other extension of the relational algebra, that proposed by Ben-Zvi, includes support for schema versioning. Ben-Zvi proposes that a temporal relation's schema itself be represented as a temporal relation, thus providing a uniform treatment for evolution of a relation and its schema [Ben-Zvi 1982]. He does not, however, provide formal semantics for schema evolution in the context of general support for transaction time. Martin proposes a non-algebraic solution to the problem of an evolving schema in temporal databases using modal temporal logic [Martin et al. 1987]. A schema temporal logic is proposed to deal with changes in schema. A set of schema temporal logic formulae are associated with a schema to describe its evolution and temporal queries are interpreted in the context of these formulae. This approach, unlike ours, forces synchronization between valid time and schema changes. Again, formal semantics are not provided. Finally, Adiba, in describing mechanisms for the storage and manipulation of historical multi-media data, advocates, like Ben-Zvi, that the history notion used to model changes in a database's contents also be used to model changes in the database's schema [Adiba & Bui Quang 1986].

While there has been significant interest in database reorganization and evolution (also termed restructuring) [Banerjee et al. 1987, Kim & Korth 1989, Markowitz & Makowsky 1987, Navathe & Fry 1976, Navathe 1980, Roussopoulos & Mark 1985, Shu et al. 1977, Shu 1987, Sockut & Goldberg 1979], such approaches have assumed that the schema (and hence the contents) of the entire database will be modified during restructuring, ensuring that only one schema is in force. These approaches address schema evolution but not schema versioning. Since we formalize the schema as a sequence ordered by transaction time, several schemas can be in force, selectable through the rollback operator. A second difference is that we focus solely on algebraic support for schema evolution, while the other papers considered the related issues of determining what changes to the schema are necessary and what those changes imply regarding the new state to be calculated. Certainly, all these issues must be addressed before a comprehensive solution to schema evolution is developed.

In contrast to these previous approaches, the WAND system did permit several generations of schemas to be simultaneously present [Gerritsen & Morgan 1976]. This system differs from our approach in two respects. First, the WAND system was based on the network model, whereas our approach is based on the relational model. More significantly, schema evolution was supported in the WAND system to allow dynamic restructuring of the database. While data in the WAND system could also be associated with one of several generations of schemas, the data were always restructured to match the most recent schema as they were referenced. Multiple generations were introduced to achieve concurrency between restructuring and execution of application programs. Hence, the underlying model did not support transaction time or rollback. The WAND system was effectively a snapshot DBMS that permitted applications to access and change the database while a global restructuring was being performed.

ORION, a prototype object-oriented database system being developed at MCC, takes a similar approach [Banerjee et al. 1987]. An important difference is that when the schema in ORION is modified, no disk-resident data instances need be updated. Instead, when an instance is referenced by an application program and fetched into memory, it is transformed into an instance conforming to the schema currently in effect. Again, only one schema is ever in effect; the implementation places the burden of updating the data across a schema change on subsequent retrievals.

Schema versioning has also been investigated in the context of object-oriented database models [Kim & Korth 1989, Kim & Chou 1988]. However, their versioning differs significantly from that proposed here. First, their versions are identified by user-supplied names, rather than by transaction numbers, as in our model. Each object update operation names a schema version to which it is to be applied, and an update to one version of an object could affect the data stored in other versions. In our model, only the most recent version can be updated, and updates do not affect other versions.

In the ENCORE object-oriented DBMS [Skarra & Zdonik 1986, Skarra & Zdonik 1987], versions are identified by transaction numbers, as in our model, but updates can be made to any version, as in the other object-oriented models. However, in ENCORE, an update applied to one version does not affect the other versions. At the risk of oversimplifying, object-oriented schema versions resemble traditional relational views [Chamberlin et al. 1975] and do not involve transaction time as defined in this paper.

Several researchers have used denotational semantics to define formally the semantics of databases, DBMS's, and query languages. Subieta proposes an approach for defining query languages formally using denotational semantics [Subieta 1987]. This approach allows powerful query languages with precise semantics to be defined for most database models. Rishe proposes that denotational semantics be used to provide a uniform treatment of database semantics at different information levels based on hierarchies of domains of mappings from "less semantic" representations of information into "more semantic" representations [Rishe 1985]. Neither Subieta nor Rishe, however, include in their approaches any facilities for dealing with transaction time or an evolving schema. Lee proposes a denotational semantics for administrative databases, where databases are regarded as a collection of logical assertions [Lee 1985]. Here, the denotation of an expression in a first-order predicate calculus is based, in part, on its evaluation in a time dimension, analogous to valid time, in a possible world, analogous to a cross-section of a database state at a transaction.

An obvious next step would be to implement an evolving schema that fully supports the rollback operator. One approach we are considering converts the system relations in our prototype [Ahn 1986A] to be rollback relations, rather than snapshot relations as they are now. Changes to the semantic analysis portions of the query analysis would be required, but it appears that changes to the backend of the DBMS would be minimal.

Another step would be to investigate extensions to the language. A straightforward extension of the language would introduce algebraic operators that map between the domain of snapshot states and the domain of historical states directly. The introduction of such operators into the snapshot and historical algebras would render the algebras multisorted. Because the two algebras, without these operators, are unsorted and because we wish to retain this property for the algebras, we have elected not to introduce such conversion operators into our language.

A second extension would introduce an algebra of signatures, analogous to the algebras of snapshot and historical states, to remove the restriction that signature specifications in the commands `define_relation` and `modify_relation` be a relation's current signature or a constant. This extension would support signature changes dependent on both the current and past signatures of relations in the database.

A third extension would remove the requirement of a relation's schema being constant over the transaction interval specified in the rollback operation. The major problem is in calculating the schema for the resulting relation. A general but simple approach has not yet been found.

Finally, the effect of schema evolution on applications programs accessing the database should be considered [Gerritsen & Morgan 1976]. Maintaining consistency between such programs and the database schema becomes more difficult. Similarly, query pre-compilation, such as performed in System R [Chamberlin et al. 1981], may or may not be effective, depending on whether the timestamps provided to the rollback operators are constants or are values supplied by the application program. However, it appears that techniques similar to those employed by the WAND system, those appearing in ORION, and those proposed [Kim & Korth 1989] could serve to amortize the cost of schema changes.

5 Acknowledgements

6 Bibliography

[Adiba & Bui Quang 1986] Adiba, M.E. and N. Bui Quang. *Historical Multi-media Databases*, in *Proceedings of the Conference on Very Large Databases*. Ed. Y. Kambayashi. Kyoto, Japan: Aug. 1986, pp. 63-70.

[Ahn 1986A] Ahn, I. *Towards an Implementation of Database Management Systems with Temporal Support*, in *Proceeding of the International Conference on Data Engineering*. IEEE Computer Society. Los Angeles, CA: IEEE Computer Society Press, Feb. 1986, pp. 374-381.

[Ahn 1986B] Ahn, I. *Performance Modeling and Access Methods for Temporal Database Management Systems*. PhD. Diss. Computer Science Department, University of North Carolina at Chapel Hill, July 1986.

[Ahn & Snodgrass 1986] Ahn, I. and R. Snodgrass. *Performance Evaluation of a Temporal Database Management System*, in *Proceedings of ACM SIGMOD International Conference on Management of Data*. Ed. C. Zaniolo. Association for Computing Machinery. Washington, DC: May 1986, pp. 96-107.

[Ahn & Snodgrass 1988] Ahn, I. and R. Snodgrass. *Performance Analysis of Temporal Queries (to appear). Information Sciences*, (1988).

[Ariav 1986] Ariav, G. *A Temporally Oriented Data Model. ACM Transactions on Database Systems*, 11, No. 4, Dec. 1986, pp. 499-527.

[Banerjee et al. 1987] Banerjee, J., W. Kim, H.-J. Kim and H.F. Korth. *Semantics and Implementation of Schema Evolution in Object-Oriented Databases*, in *Proceedings of ACM SIGMOD International Conference on Management of Data*. Ed. U. Dayal and I. Traiger. Association for Computing Machinery. San Francisco, CA: 1987, pp. 311-322.

[Ben-Zvi 1982] Ben-Zvi, J. *The Time Relational Model*. PhD. Diss. Computer Science Department, UCLA, 1982.

[Bontempo 1983] Bontempo, C. J. *Feature Analysis of Query-By-Example*, in Relational Database Systems. New York: Springer-Verlag, 1983. pp. 409-433.

[Chamberlin et al. 1975] Chamberlin, D.D., J.N. Gray and I.L. Traiger. *Views, Authorization, and Locking in a Relational Data Base System*, in *AFIPS Conference Proceedings*. AFIPS. Anaheim, CA: 1975, pp. 425-430.

[Chamberlin et al. 1981] Chamberlin, D.D., M.M. Astrahan, W.F. King, R.A. Lorie, J.W. Mehl, T.G. Price, M. Schkolnick, P. Selinger Griffiths, D.R. Slutz, B.W. Wade and R.A. Yost. *Support for Repetitive Transactions and Ad Hoc Queries in System R. ACM Transactions on Database Systems*, 6, No. 1, Mar. 1981, pp. 70-94.

[Clifford & Croker 1987] Clifford, J. and A. Croker. *The Historical Relational Data Model (HRDM) and Algebra Based on Lifespans*, in *Proceedings of the International Conference on Data Engineering*. IEEE Computer Society. Los Angeles, CA: IEEE Computer Society Press, Feb. 1987, pp. 528-537.

[Codd 1970] Codd, E.F. *A Relational Model of Data for Large Shared Data Banks. Communications of the Association of Computing Machinery*, 13, No. 6, June 1970, pp. 377-387.

[Date 1976] Date, C. J. *An Introduction to Database Systems*. Systems Programming Series. Reading, MA: Addison-Wesley Publishing Company, 1976.

[Gadia 1984] Gadia, S.K. *A Homogeneous Relational Model and Query Languages for Temporal Databases*. 1984. (Unpublished paper.)

[Gadia 1986] Gadia, S.K. *Toward a Multihomogeneous Model for a Temporal Database*, in *Proceedings of the International Conference on Data Engineering*. IEEE Computer Society. Los Angeles, CA: IEEE Computer Society Press, Feb. 1986, pp. 390-397.

[Gadia & Yeung 1988] Gadia, S.K. and C.S. Yeung. *A Generalized Model for a Relational Temporal Database*, in *Proceedings of ACM SIGMOD International Conference on Management of Data*. Association for Computing Machinery. Chicago, IL: June 1988, pp. 251-259.

[Gerritsen & Morgan 1976] Gerritsen, R. and H.L. Morgan. *Dynamic Restructuring of Databases with Generation Data Structures*, in *Proceedings of the ACM Annual Conference*. Association for Computing Machinery. Houston, TX: Oct. 1976, pp. 281-286.

[Gordon 1979] Gordon, M.J.C. *The Denotational Description of Programming Languages*. New York-Heidelberg-Berlin: Springer-Verlag, 1979.

[Jones et al. 1979] Jones, S., P. Mason and R. Stamper. *LEGOL 2.0: A Relational Specification Language for Complex Rules. Information Systems*, 4, No. 4, Nov. 1979, pp. 293-305.

[Kim & Korth 1989] Kim, H.-J. and H.F. Korth. *Schema Versions and DAG Rearrangement Views in Object-Oriented Databases*. Technical Report. ICS; Department of Computer Science. Mar. 1989.

[Kim & Chou 1988] Kim, W. and H.T. Chou. *Versions of Schema in OODB*, in *VLDB88*. Long Beach, CA: 1988.

[Lee 1985] Lee, R.M. *A Denotational Semantics for Administrative Databases*, in *Proceedings of*

the IFIP WG 2.6 Working Conference on Data Semantics (DS-1). Ed. T.B. Steel and R. Meersman. IFIP. Hasselt, Belgium: Jan. 1985, pp. 83-120.

[Lorentzos & Johnson 1987] Lorentzos, N.A. and R.G. Johnson. *TRA: A Model for a Temporal Relational Algebra*, in *Proceedings of the Conference on Temporal Aspects in Information Systems*. AFCET. France: May 1987, pp. 99-112.

[Lum et al. 1984] Lum, V., P. Dadam, R. Erbe, J. Guenauer, P. Pistor, G. Walch, H. Werner and J. Woodfill. *Designing DBMS Support for the Temporal Dimension*, in *Proceedings of ACM SIGMOD International Conference on Management of Data*. Ed. B Yormark. Association for Computing Machinery. Boston, MA: June 1984, pp. 115-130.

[Maier 1983] Maier, D. *The Theory of Relational Databases*. Rockville, MD: Computer Science Press, 1983.

[Markowitz & Makowsky 1987] Markowitz, V.M. and J.A. Makowsky. *Incremental Reorganization of Relational Databases*, in *Proceedings of the Conference on Very Large Databases*. Ed. P. Hammersley. Brighton, England: Sep. 1987, pp. 127-135.

[Martin et al. 1987] Martin, N.G., S.B. Navathe and R. Ahmed. *Dealing with Temporal Schema Anomalies in History Databases*, in *Proceedings of the Conference on Very Large Databases*. Ed. P. Hammersley. Brighton, England: Sep. 1987, pp. 177-184.

[McKenzie 1986] McKenzie, E. *Bibliography: Temporal Databases. ACM SIGMOD Record*, 15, No. 4, Dec. 1986, pp. 40-52.

[McKenzie & Snodgrass 1987A] McKenzie, E. and R. Snodgrass. *Extending the Relational Algebra to Support Transaction Time*, in *Proceedings of ACM SIGMOD International Conference on Management of Data*. Ed. U. Dayal and I. Traiger. Association for Computing Machinery. San Francisco, CA: May 1987, pp. 467-478.

[McKenzie & Snodgrass 1987B] McKenzie, E. and R. Snodgrass. *An Evaluation of Historical Algebras*. Technical Report TR87-020. Computer Science Department, University of North Carolina at Chapel Hill. Oct. 1987.

[McKenzie 1988] McKenzie, E. *An Algebraic Language for Query and Update of Temporal Databases*. PhD. Diss. Computer Science Department, University of North Carolina at Chapel Hill, Sep. 1988.

[Navathe & Fry 1976] Navathe, S.B. and J.P. Fry. *Restructuring for Large Databases: Three Levels of Abstraction. ACM Transactions on Database Systems*, 1, No. 2, June 1976, pp. 138-158.

[Navathe 1980] Navathe, S.B. *Schema Analysis for Database Restructuring. ACM Transactions on Database Systems*, 5, No. 2, June 1980, pp. 157-184.

[Navathe & Ahmed 1987] Navathe, S.B. and R. Ahmed. *TSQL-A Language Interface for History Databases*, in *Proceedings of the Conference on Temporal Aspects in Information Systems*. AFCET. France: May 1987, pp. 113-128.

[Overmyer & Stonebraker 1982] Overmyer, R. and M. Stonebraker. *Implementation of a Time Expert in a Database System. ACM SIGMOD Record*, 12, No. 3, Apr. 1982, pp. 51-59.

[Rishe 1985] Rishe, N. *On Denotational Semantics of Data Bases*, in *Proceedings of the International Conference on Mathematical Foundations of Programming Semantics*. Ed. A. Melton. Manhattan, KA: Springer-Verlag, Apr. 1985, pp. 249-274.

[Rotem & Segev 1987] Rotem, D. and A. Segev. *Physical Organization of Temporal Databases*, in *Proceedings of the International Conference on Data Engineering*. IEEE Computer Society. Los Angeles, CA: IEEE Computer Society Press, Feb. 1987, pp. 547-553.

[Roussopoulos & Mark 1985] Roussopoulos, N. and L. Mark. *Schema Manipulation in Self-Describing and Self-Documenting Data Models. International Journal of Computer and Information Sciences*, 14, No. 1, Jan. 1985, pp. 1-28.

[Schmidt 1986] Schmidt, D.A. *Denotational Semantics, A Methodology for Language Development*. Newton, Massachusetts: Allyn and Bacon, 1986.

[Scott 1976] Scott, D.S. *Data Types as Lattices. SIAM Journal of Computing*, 5, No. 3, Sep. 1976, pp. 522-587.

[Shiftan 1986] Shiftan, J. *An Assessment of the Temporal Differentiation of Attributes in the Implementation of a Temporally Oriented DBMS*. PhD. Diss. Information Systems Area, Graduate School of Business Administration, New York University, Aug. 1986.

[Shoshani & Kawagoe 1986] Shoshani, A. and K. Kawagoe. *Temporal Data Management*, in *Proceedings of the Conference on Very Large Databases*. Ed. Y. Kambayashi. Kyoto, Japan: Aug. 1986, pp. 79-88.

[Shu et al. 1977] Shu, N.C., B.C. Housel, R.W. Taylor, S.P. Ghosh and V.Y. Lum. *EXPRESS: A Data EXtraction, Processing, and REStructuring System. ACM Transactions on Database Systems*, 2, No. 2, June 1977, pp. 134-174.

[Shu 1987] Shu, N.C. *Automatic Data Transformation and Restructuring*, in *Proceedings of the International Conference on Data Engineering*. IEEE Computer Society. Los Angeles, CA: IEEE Computer Society Press, Feb. 1987, pp. 173-180.

[Skarra & Zdonik 1986] Skarra, A.H. and S.B. Zdonik. *The Management of Changing Types in an Object-Oriented Database*, in *Proceedings of the Conference on Object-Oriented Programming Systems, Languages and Applications*. Ed. N. Meyrowitz. Association for Computing Machinery. Portland, OR: Nov. 1986, pp. 483-495.

[Skarra & Zdonik 1987] Skarra, A.H. and S.B. Zdonik. *Type Evolution in an Object-Oriented Database*, in Research Directions in Object-Oriented Programming. of Computer Systems Series. Cambridge, MA: MIT Press, 1987. Chap. Part 3. pp. 393-415.

[Snodgrass & Ahn 1985] Snodgrass, R. and I. Ahn. *A Taxonomy of Time in Databases*, in *Proceedings of ACM SIGMOD International Conference on Management of Data*. Ed. S. Navathe. Association for Computing Machinery. Austin, TX: May 1985, pp. 236-246.

[Snodgrass & Ahn 1986] Snodgrass, R. and I. Ahn. *Temporal Databases. IEEE Computer*, 19, No. 9, Sep. 1986, pp. 35-42.

[Snodgrass 1987] Snodgrass, R. *The Temporal Query Language TQuel. ACM Transactions on Database Systems*, 12, No. 2, June 1987, pp. 247-298.

[Sockut & Goldberg 1979] Sockut, G.H. and R.P. Goldberg. *Database Reorganization - Principles and Practice. ACM Computing Surveys*, 11, No. 4, Dec. 1979, pp. 371-395.

[Stoy 1977] Stoy, Joseph E. *Denotational Semantics: The Scott-Strachey Approach to Programming Language Theory.* The MIT Series in Computer Science. The MIT Press, 1977.

[Strachey 1966] Strachey, C. *Towards a Formal Semantics*, in Formal Language Description Languages for Computer Programming. North Holland, 1966. pp. 198-220.

[Subieta 1987] Subieta, K. *Denotational Semantics of Query Languages. Information Systems*, 12, No. 1 (1987), pp. 69-82.

[Tandem 1983] Tandem Computers, Inc. *ENFORM Reference Manual.* Cupertino, CA, 1983.

[Tansel 1986] Tansel, A.U. *Adding Time Dimension to Relational Model and Extending Relational Algebra. Information Systems*, 11, No. 4 (1986), pp. 343-355.

[Thirumalai & Krishna 1988] Thirumalai, S. and S. Krishna. *Data Organization for Temporal Databases.* Technical Report. Raman Research Institute, India. 1988.

[Ullman 1982] Ullman, J.D. *Principles of Database Systems, Second Edition.* Potomac, Maryland: Computer Science Press, 1982.

[Woelk et al. 1986] Woelk, D., W. Kim and W. Luther. *An Object-Oriented Approach to Multimedia Databases*, in *Proceedings of ACM SIGMOD International Conference on Management of Data.* Ed. C. Zaniolo. Washington, DC: Association for Computing Machinery, May 1986, pp. 311-325.

[Yeung 1986] Yeung, C.S. *Query Languages for a Heterogeneous Temporal Database.* Master's Thesis, EE/CS Department, Texas Tech University, 1986.

A   Proofs of Language Properties

In this appendix, we provide informal proofs for three properties of our language, stated as theorems in Section 2.7.

Theorem 1 *The language is a natural extension of the relational algebra for database query and update.*

PROOF. First, we show that expressions in our language are a strict superset of those in the relational algebra. Suppose we only allow expressions involving constants that denote snapshot states, identifiers that denote relations whose current class is SNAPSHOT, and the five relational operators. Then, expressions in the language are exactly those allowed in the relational algebra. But expressions in our language also may involve constants that denote historical states, identifiers that denote relations whose current class is other than SNAPSHOT, and both historical and rollback operators. Hence, expressions in our language are a strict superset of those in the relational algebra.

Next, we show that our semantics reduces to the conventional semantics of database state and database update via the relational algebra. Suppose we restrict the class of all relations to UNDEFINED and SNAPSHOT. Then,

(a) The signature and state sequences of a defined relation will have exactly one element each, the relation's current signature and state. The relation can have no history as a rollback or temporal relation; hence its MSoT always will be $(\langle\ \rangle, \langle\ \rangle, \langle\ \rangle)$. Because the `define_relation` and `modify_relation` commands change a relation's signature sequence by appending no more than one element to the relation's MSoT signature sequence, these commands always will produce a relation with a single-element signature sequence. The same holds for the relation's state sequence.

(b) A new state always will be a function of the current signature and state of defined relations via the relational algebra semantics. Both the `define_relation` and `modify_relation` commands determine a new state via expression evaluation. The only semantically correct expressions are those involving constants that denote snapshot states, identifiers that denote relations whose current class is SNAPSHOT, and the five relational operators. These expressions are exactly those allowed in the relation algebra, their value depending on the current state and signature of defined relations only.

(c) Deletion will correspond to physical deletion. The `destroy` command changes a relation by appending an element to the relation's MSoT class sequence; it never adds information to the relation's signature or state sequences. The `destroy` command always will produce a relation whose signature and state sequences are empty, which corresponds to physical deletion of a relation's current signature and state. ∎

Theorem 2 *The semantics of the language minimizes the number of elements in a relation's class, signature, and state sequence needed to record the relation's current class, signature, and state and its history as a rollback or temporal relation.*

PROOF. Assume that the number of elements in a relation's class sequence exceeds the minimum needed to record the relation's current class and its history as a rollback or temporal relation. Then, (a) there are two consecutive elements in the sequence that can be combined or (b) there is an element in the sequence that can be removed. Consider case (a). Consecutive elements in the class sequence can be combined only if they record the same class over non-disjoint intervals. But the commands only append a new element to a relation's class sequence if it either differs from the relation's MSoT class or its interval is disjoint from that of the relation's MSoT class. Hence, no two consecutive elements in a relation's class sequence can have the same class but non-disjoint intervals. Now, consider case (b). Commands always produce a new relation by appending new

49

class information to a relation's MSoT class sequence. But, it can be shown that all elements in a relation's MSoT class sequence record intervals when the relation was either a rollback or temporal relation. Hence, no element can be removed. If no two elements can be combined and no element can be removed, our assumption is contradicted and the number of elements in the class sequence must be minimal. Similar arguments hold for the relation's signature and state sequences. ∎

Theorem 3 *Transactions change only a relation's class, signature, and state current at the start of the transaction.*

PROOF. This property is a consequence of the way the *MSoT* function is defined and used. We first prove the property for a relation's signature sequence and then for its class and state sequences.

A relation's current signature at the start of a transaction is the last element in the relation's signature sequence. Assume, therefore, that a transaction changes an element that is in the relation's signature sequence at the start of the transaction but is not the last element in the sequence. Such a change must occur during the execution of a command. When the first command in a transaction executes, *MSoT* discards the last element in the relation's signature sequence, if the relation's current class is either SNAPSHOT or HISTORICAL. Otherwise, it retains all the elements. When each subsequent command in the transaction is executed, *MSoT* only discards any element that the preceding command added to the sequence. Hence, *MSoT* never changes an element in a relation's signature sequence that precedes the last element in the sequence at the start of the transaction. Commands, although they may append an element to the relation's MSoT signature sequence, never change existing elements. Hence, commands never change an element in a relation's signature sequence that precedes the last element in the sequence at the start of the transaction and our assumption is contradicted. The same argument holds for the relation's state sequence.

The above argument holds for a relation's class sequence with the following provisos. When the first command in a transaction executes, *MSoT* discards the last element in the relation's class sequence if the relation's current class is UNDEFINED. Also, if the relation's current class is either ROLLBACK or TEMPORAL, *MSoT* changes the last element in the sequence to "close" the interval assigned to the relation's current class at the start of the transaction. When each subsequent command in the transaction is executed, *MSoT* "re-closes" this same interval, if extended by the preceding command. Hence, *MSoT* never changes an element in a relation's class sequence that precedes the last element in the sequence at the start of the transaction. Commands may change the last element in a relation's MSoT class sequence to "extend" the interval assigned to the class component of that element, but only if the new class and the relation's MSoT class are equal and their intervals abut. This occurs only when the last element in the relation's MSoT class sequence corresponds to the last element in the relation's class sequence at the start of the transaction (i.e., the class of the relation at the start of the transaction was either ROLLBACK or TEMPORAL). Otherwise, the intervals could not abut as there would exist an intervening interval when the relation's class was either SNAPSHOT, HISTORICAL, or UNDEFINED. Hence, commands never change an element in a relation's class sequence that precedes the last element in the sequence at the start of the transaction. ∎

B MSoT and its Subordinate Functions

In this appendix, we present a formal definition for *MSoT* (*M*odified *S*tart *o*f *T*ransaction) and descriptions of its subordinate functions. For these definitions, let

tn and tn' range over the domain [$\mathcal{TRANSACTION\ NUMBER}$ + {–}],

u and u' range over the domain [$\mathcal{RELATION\ CLASS}$ × $\mathcal{TRANSACTION\ NUMBER}$]*,

v range over the domain [$\mathcal{RELATION\ SIGNATURE}$ × $\mathcal{TRANSACTION\ NUMBER}$]*,

w range over the domain

$$[\ [\mathcal{SNAPSHOT\ STATE}\ \times\ \mathcal{TRANSACTION\ NUMBER}] +$$
$$[\mathcal{HISTORICAL\ STATE}\ \times\ \mathcal{TRANSACTION\ NUMBER}]\ ]^*,$$

MSoT maps a relation $(u,\ v,\ w)$ and a transaction number tn onto the history of the relation as a rollback or temporal relation before the start of transaction tn.

$$MSoT : [\ [\ \mathcal{RELATION} + \{\ (\langle\ \rangle, \langle\ \rangle, \langle\ \rangle)\ \}\ ]\ \times\ \mathcal{TRANSACTION\ NUMBER}\ ] \rightarrow$$
$$[\ \mathcal{RELATION} + \{\ (\langle\ \rangle, \langle\ \rangle, \langle\ \rangle)\ \}\ ]$$

$MSoT((u,\ v,\ w),\ tn) =$

    if $(u' = PrefixClasses(u,\ tn) \wedge u' \neq \langle\ \rangle \wedge tn' = LastTrNumber(u'))$

    then if $MultiStateClass(LastClass((u',\ v,\ w)))$

        then $(Close(u',\ tn - 1),\ PrefixSigs(v,\ tn),\ PrefixStates(w,\ tn))$

        else $(PrefixClasses(u,\ tn'),\ PrefixSigs(v,\ tn'),\ PrefixStates(w,\ tn'))$

    else $(\langle\ \rangle,\ \langle\ \rangle,\ \langle\ \rangle)$

Close maps a relation's class sequence u and a transaction number tn onto the subsequence recorded through transaction tn. It also sets the the second transaction-number component in the last element of the resulting sequence to tn if the component is either "–" or greater than tn.

LastTrNumber maps a relation's class sequence onto the transaction number of the transaction that appended the last element to the sequence. If the relation's class sequence is empty, *LastTrNumber* returns ERROR.

MultiStateClass is a boolean function that determines whether a class is either ROLLBACK or TEMPORAL.

PrefixClasses maps a relation's class sequence u and a transaction number tn onto the subsequence recorded before the start of transaction tn.

PrefixSigs maps a relation's signature sequence v and a transaction number tn onto the subsequence recorded before the start of transaction tn.

PrefixStates maps a relation's state sequence w and a transaction number tn onto the subsequence recorded before the start of transaction tn.