MSQL : A multidatabase language

W. Litwin, A. Abdellatif, B. Nicolas, Ph. Vigier, A. Zeroual

► To cite this version:

W. Litwin, A. Abdellatif, B. Nicolas, Ph. Vigier, A. Zeroual. MSQL : A multidatabase language. RR-0695, INRIA. 1987. inria-00075858

HAL Id: inria-00075858

<https://hal.inria.fr/inria-00075858>

Submitted on 24 May 2006

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

UNITÉ DE RECHERCHE
IRIA-ROCQUENCOURT

Institut National
de Recherche
en Informatique
et en Automatique

Domaine de Voluceau
Rocquencourt

BP 105

78153 Le Chesnay Cedex
France

Tél. (1) 39 63 55 11

Rapports de Recherche

N°695

MSQL : A MULTIDATABASE LANGUAGE

Witold LITWIN
Abdellaziz ABDELLATIF
Bertrand NICOLAS
Philippe VIGIER
Abdelmalek ZEROUAL

JUIN 1987

MSQL: A Multidatabase Language

W. Litwin, A. Abdellatif, B. Nicolas, Ph. Vigier and A. Zeroual

INRIA, BP 105, 78153 Le-Chesnay, France

ABSTRACT

Many databases are now accessible through computer networks and users frequently have data of interest in different databases. Manipulations of such data call for functions that do not exist in classical manipulation languages and database systems. A new type of data manipulation language is needed that was called multidatabase manipulation languages. We describe the multidatabase language MSQL designed for manipulations of data from different relational databases. MSQL is an extension of SQL language that is became the standard for relational database systems.

MSQL: Un Langage Multibase

W. Litwin, A. Abdellatif, B. Nicolas, Ph. Vigier and A. Zeroual

INRIA, BP 105, 78153 Le-Chesnay, France

RESUME

De nombreuses bases de données sont actuellement accessibles à travers des réseaux d'ordinateurs et les données intéressant un usager sont souvent dans plusieurs bases. La manipulation de telles données nécessitent des fonctionnalités inconnues des langages et des systèmes de bases de données classiques. Un nouveau type de langage de manipulation de données est nécessaire, qu'on a appelé langage de manipulation multibase. Nous présentons le langage multibase MSQL conçu pour la manipulation de données provenant de différentes bases de données relationnelles. MSQL est une extension du langage SQL qui est devenu le standard pour les systèmes relationnels.

1. INTRODUCTION

Many databases now exist on various computers and are accessible through computer networks. A large mainframe with a database system like INGRES or a database server like Compuservice may support dozens of databases. A videotex system like Teletel may provide access to thousands of databases for millions of people. A major consequence is that many users now have data of interest in different databases. Such users frequently need to jointly manipulate data in different databases. A Teletel user who wishes to go out for a movie and for dinner, may wish to find in some cinema guide database and in some restaurant guide database, the cinemas and restaurants that are in the same street. A bank may have several branch databases and a manager of company may wish the balances of the accounts that the company has at different branches. If this manager uses Teletel, he may further dream about such a query to all the banks his enterprise uses (most French banks are accessible through Teletel). Then, a user of Teletel who wants to find the cheapest way to travel somewhere, may need to query several airline databases and maybe the SNCF (French railways) database. Furthermore, if a street changes name, it may be necessary to update all the corresponding databases, etc.

Manipulations of such data require functions that do not exist in classical manipulation languages and database systems /LIT86b/, designed for manipulations of data integrated into one database. One reason is that data in different databases are not usually mutually integrated, as, for various reasons, the administrators wish to remain independent. Data about the same real things (accounts, restaurants, etc.) present then semantic differences. The differences consist of partial redundancy, differences in naming, data structures, types, scales, independence of keys etc. This type of heterogeneity is independent of data model which means that it exists, even if all databases present the common data model and manipulation language for cooperative usage. The differences among similar data result from different perceptions of the same real universe, from local needs, from the competition for customers (restaurant guides), from political incompatibility (French socialist party database and French National Front database), etc. Usually, in such a case, there is no way to force the integration through a global schema, since no one would agree on the global administrator. The databases shown in the appendix, modeling actual databases, show some of differences that may occur.

The system for the manipulation of data in autonomous databases was called *multidatabase system* (MBS) /LIT82/. By the same token, the corresponding manipulation language was called *multidatabase manipulation language* (MML). Databases that may be manipulated together without global integration, were called *interoperable* /LIT86b/. Relational databases are a particularly

important subject for interoperability, since most future databases will present a relational interface, no matter what data model and implementation issues are used internally. /LIT86b/ describes the prototype relational MBS termed MRDSM and its MML, termed MDSL. MRDSM renders interoperable databases managed by MRDS system that is the relational DBMS of Multics system (Multics Relational Data Store /MRD82/). Most of the new functions of MDSL are still unknown to other languages.

MDSL is an extension of the MRDS manipulation language DSL. This language is particular to the MRDS system, although its syntax is rather typical of relational languages, it is somewhere between QUEL and SQL. Nevertheless, it is clear today that the future standard language for a relational database will be SQL. This language is therefore at present the best basis for multidatabase interoperability.

Below, we present the multidatabase extension of SQL, termed MSQL (Multidatabase SQL) intended for this usage. Any function of SQL is by definition the function of MSQL. New functions are designed for non-procedural manipulation of data in different and basically mutually non-integrated SQL databases. This means that the user wish (informal query) should usually become a single MSQL statement. The overall design of MSQL is based on that of MDSL, with adaptations to SQL syntax and semantics. The definition of SQL used is mainly that of ISO /SQL86/ and of DB2 dialect /DAT83/. We assume these definitions known and the main properties of SQL that result from them. The new possibilities that MSQL statements provide for interoperability are as follows :

- the single statement table creation or alteration in any number of databases and the import of the database, table or column schemas;
- the retrieval or the modification involving the joining of data in different database schemas;
- the broadcasting of a retrieval or of a modification over any number of databases where data with similar meanings have the same or different naming rules, decompositions into relations or value types.
- the dynamic transformation of actual attribute meanings, units of measure etc., into user-defined value types that may be retrieved or used for update expression (some restrictions apply to updates);
- the interdatabase queries for data flow between databases;
- the dynamic aggregation of data from different databases using various new standard (built-in) functions.
- the creation of multidatabase views and of virtual databases over such views;
- the creation of auxiliary objects like triggers, stored queries and transactions (procedures).

Section 2 presents MSQL data definition statements. Section 3 deals with functions for multidatabase retrieval and update. Section 4 presents the views. Interdatabase queries are discussed in Section 5 and auxiliary objects in Section 6. Section 7 concludes the MSQL description.

2. DATA DEFINITION

2.1 Introduction

As /ANS75/, in what follows the term *database* represents a (logical) model of a certain company (also called universe, real universe, real world, or world in some implementations of Prolog /MAR86/). It should not be confused with a meaningless physical storage structure called (physical) database in some implementations, (DB2 in particular /DAT83/) or with the physical concept of a site (network node). The databases are usually meaningfully named upon their universes, like database AIR FRANCE and the databases in the Appendix and unlike "physical" databases and sites (ex. the site name GCAM of AIR FRANCE database is meaningless for database users). Finally, the databases are defined by their schemas and data models, no matter how they are implemented internally. For instance, a database implemented using IMS, but where the user perceives data as tables manipulated using SQL is a relational database (note : this is in particular the approach in /DAT83/). Below, we deal with relational and SQL databases only.

The relational database design follows the idea of integration. The basic data structure for the integration is a table that is a set of named columns whose values are atomic and of the same value type. Data about all (real) objects (entities or relationships) of the same type are supposed to be integrated into the same table. The objects of type Suppliers are for instance supposed to constitute a single table usually named S. The price to pay is that each object has to be modelled basically through exactly the same column names and value types, unless null values are used with all the problems they create. The reward is the simplicity of relational languages for interactive usage that disappears otherwise.

Ex. 2.1 If all suppliers are described by the same table S, then the wish to select all data about all suppliers leads to the simplest possible SQL query :

```
SELECT *  
FROM S
```

If however the suppliers were in $i > 1$ tables, for instance one per department the suppliers deal with, then the expression of the same wish will need i queries like the above one in any relational

database language. If i is large, the usage of SQL becomes laborious since it is highly procedural, if i is unknown then the wish cannot be expressed at all.

The queries using column names may in addition differ in formulation if table schemas are particularized. Aggregate operations like **GROUP BY** or **SUM** etc. may further be impossible to express if tuples to aggregate are in different tables. While the departments of Ex. 2.1 may therefore find it necessary to be integrated into a single table, they may remain unhappy with the compromises they had to adopt to get the common schema. This is in fact the situation for many actual applications.

In a multidatabase environment the integration principle is naturally applied up to the level of a database. Afterwards, as data are in several databases that are naturally autonomous, the objects of the same type may be modelled by several tables. These tables may bear the same or different names and may have the same or different columns. If the users use the same natural languages and the administrators work in some cooperation, as is usual in an enterprise or even when organizations compete for customers, then the usual situation is as follows :

- some tables and columns designate the same things using the same name and value types (ex. price is called **price** in all tables and value type is FFr). This unicity may result from some cooperative standards adopted by the autonomous administrators.
- some tables have in contrast particular names and/or bear columns particular to the database. These columns may have a particular name (ex. **prices**) or a particular value type (ex. price with VAT included) or even may have no semantic equivalent in other databases (no notion of price at all).
- some real objects may be modelled by only one table, while some other may correspond to tuples in several tables (ex. a supplier working for several departments). The values may be duplicated, or may differ in precision, or may have no objective (univocal) relationship or may clearly contradict.

This "right to be different" seems highly appreciated by users and appears to be the **motivation** for the current trend towards multiple autonomous databases, instead of a single integrated one (more important than physical performance considerations). Though different tables may then be sometimes presented through a view as a single table, this approach does not seem the natural issue (both conceptually and technically because of update problems). The natural goal seems rather to be **an extension of relational principles and statements to sets of tables** that may appear subject to such needs in the multidatabase environment. The set of tables **implicitly** or **explicitly** subject to a common statement will be called *multitable*. In Ex. 2.1, the multitable would consist of all the supplier tables of all departments and could of course span over multiple

databases. The application of the relational principles implies in particular that a manipulation of a multitable should be expressible through a single statement, at least whenever the same statement would apply if data in a multitable were integrated into a table. In this sense the manipulation of multitable would become as simple as that of tables. The tables of a multitable become interoperable providing the user both relational simplicity and better recognition of his autonomy.

MSQL is designed particularly with this goal in mind, as it was also implicitly the goal of MDSL /LIT86b/. SQL possibilities are extended to multitable whenever it appears useful to do so. In particular, they make it possible to create multitable at data definition level and to form multitable corresponding to various wishes dynamically in queries. The corresponding functions and constructs of MSQL will now be progressively introduced.

2.2 Statements

As in SQL, the basic construct for data manipulation and definition in MSQL is a *statement*. By definition, any SQL statement is an MSQL statement. In what follows, we present only the statements and clauses relative to the logical level of data. The details relative to storage space, physical clustering, network access are assumed system-specific.

The data definition statements of MSQL that are extended with respect to SQL are as follows :

CREATE TABLE	CREATE DATABASE	CREATE MULTIDATABASE
CREATE VIEW		
ALTER TABLE	ALTER VIEW	
ALTER MULTIDATABASE		
DROP TABLE	DROP DATABASE	DROP MULTIDATABASE
DROP VIEW		

We defer the discussion of the statements relative to views to Section 5, as they require the knowledge of MSQL data manipulation statements. Some statements for the definition of auxiliary database objects will also be shown in Section 6.

2.3 Naming rules

SQL makes possible to name two data types that are *columns* (attributes), also called fields and *tables*. MSQL allows us to name the *domains*, the *databases* and the *multidatabases*. To refer to data or variable names, statements contain *designators*. Data designators in SQL, are column, table

or view identifiers. They are unique (unambiguous) identifiers that are column, table or view names, or unambiguous aliases or the qualified column names that consist of a table name and of the column name separated by a period. MSQL provides in addition the following types of designators :

- the database name,
- the multidatabase name,
- the virtual database name,
- the multiple identifiers,
- the domain name,
- the semantic variables,
- the dynamic column name.

The last two types of designators are used only in data manipulation statements and will be discussed in the corresponding sections. The database name is the name given to a relational database. The multidatabase name is the name given to a set of interoperable relational databases. The virtual database name is used to designate a set of tables or views of one or several databases. The multiple identifier is the name shared by several objects in a statement scope, especially tables in different databases or columns in different tables. The usage of all these types of designators will be explained during the discussion of the details of the statements.

The notion of domain, though fundamental to the relational model, does not exist in SQL, in contrast to DSL for instance /MRD82/. This state-of-the art is regrettable from both a theoretical and a practical point of view /DAT86/. In particular, since the existence of domains may usually avoid the formulation of equijoins /LIT85/. As will be shown later and is also shown in detail in /LIT85/, this possibility is seen to be even more important in the multidatabase environment, as it allows the query formulation to be largely independent of data structure heterogeneity (differences in decompositions to relations). That is why, unlike SQL, MSQL supports domains.

One way to define domains could be that called the extended SQL in /DAT86/. Another way is to consider that any column name is either the domain name or a qualified domain name if it is of form <column prefix>_<domain name>, as in /COD71/. The latter approach requires more discipline while defining the columns, as each column definition over the same domain should contain the same data type. However it is compatible with actual SQL implementations. In what follows we retain the second approach, although the corresponding features of MSQL would remain for the former approach as well.

2.4 Database creation

The following statement creates the database, either empty or with tables, or views and the corresponding privileges.

```
CREATE DATABASE <database name>  
    [<table definition>|<view definition>|<privilege definition>...] |  
FROM <database identifier>
```

<table definition>::= see §2.8

<view definition>::= see §4.1

<privilege definition>::= see §2.5

This statement must precede the creation of any table. As usual, square brackets in syntactic definitions of this and coming statements indicate that the material enclosed is optional. A '|' indicates 'or'. Ellipses '...' indicate that the material enclosed may be repeated zero or more times. The form {<a> | } means that one of the elements is mandatory. Finally, material in <> must be replaced by specific values chosen by the user.

The semantic of the statement includes that of the SQL statement, new features with respect to table or view definitions will be shown in the corresponding sections. The **FROM** clause is new with respect to SQL and results from the multidatabase environment. It makes it possible to import (copy) the entire schema of an existing database.

Ex.2.2 Create a new branch database **boulogne** upon the database **bnp**.

```
CREATE boulogne FROM bnp
```

The SQL statement would be much more procedural and subject to more errors, since it would require the whole schema of **bnp** to be retyped.

2.5 Privilege definition

The following statement defines the user's privileges or imports them. It may be used in the **CREATE DATABASE** statement or separately.

```
GRANT <privileges> [ON [<object type>] <object designators>] TO <grantees>...  
    [WITH GRANT OPTION] [FROM [<grantee> ON] <object type><object identifier> ]
```

```

<privileges> ::= ALL [PRIVILEGES] | <action> [,<action>...]
<action> ::= <MSQL statement> [(grant column list)]
<grant column list> ::= <column designator> [,<column designator>...]
<grantee> ::= PUBLIC | <authorization identifier>

```

The statement semantic is basically that of SQL GRANT statement, see for instance /DAT83/ for deeper discussion. The difference is that MSQL has more statements and object types. In particular the statement may concern one or more multitable of tables sharing the name, if it is applied independently of the CREATE DATABASE statement. The corresponding table designator(s) should then be multiple identifier(s). Furthermore, the statement may import all privileges the grantee(s) had on the object identified in the FROM clause. The object identifier in the clause should then be a unique identifier, qualified if necessary. Also, the object type in this clause and those concerned by the statement should be the same.

Ex. 2.3 The following statements authorize Nicolas and Abdellatif to select data from any of four client tables in the databases in the Appendix and import to account table in etoile database their respective and maybe different privileges on account table in bnp. The last statement imports for Nicolas the privileges of Zeroual.

```

USE bnp sg cic etoile
GRANT SELECT ON client TO Nicolas Abdellatif
GRANT ALL ON etoile.account TO Nicolas Abdellatif FROM bnp.account
GRANT ALL ON etoile.account TO Nicolas FROM Zeroual ON bnp.account

```

The USE clause is optional and discussed in Section 3.2. The designator **client** is the multiple identifier that defines dynamically the corresponding multitable. The corresponding statement replaces four SQL GRANT statements. The utility of the other two statements is obvious.

2.6 Privilege revoking

Revoking privileges is done by the following statement, also similar to the corresponding one of SQL :

```

REVOKE <privileges> ON <designator>
      FROM <grantee>...

```

2.7 Database removal

The statement :

DROP DATABASE <database name>

removes the designated database. It is particularly useful for temporary databases, like a frequently refreshed database on a PC containing a part of mainframe database etc.

2.8 Table creation

The SQL statement **CREATE TABLE** creates a table schema in the database that is known implicitly. In a multidatabase environment, it may be necessary to indicate the concerned database. Furthermore, one may need to create simultaneously several tables in cooperating databases. Finally, it may be useful to import schemas of existing tables or columns. The **MSQL CREATE TABLE** statement includes the SQL statement and responds to the new needs.

[**USE** < scope specifications>]

CREATE TABLE <table designator>

(<column definition>[, < column definition >]...) |

FROM <source table identifier>

< column definition >::=

<column name > <data type>[**NON NULL** [**UNIQUE**]] |

CONSTRAINTS <unique constraint definition>... |

FROM <source column identifier>

<unique constraint definition>::= as in /SQL86/

The **USE** clause designates the names of (multi)databases the statement applies to (the scope of the statement). It may be omitted if the current scope is used. The **FROM** clause identifies the table or the column whose schema is imported. The statement creates the tables named upon the <table designator> in all databases in the considered scope. As in SQL /SQL86/, the **UNIQUE** statement defines the keys.

Ex. 2.4 The example banks have decided to cooperate on loans. They nominated a committee for the basic definition of the corresponding table, with the **CREATE TABLE** privilege on all databases. The committee decided that each loan will be a tuple in the following table in each database :

loan (acc#, cl#, amount, rate, beg_date, duration)

The following CREATE TABLE could be used:

```
CREATE TABLE
  loan (acc# INT NON NULL,
        cl# INT NON NULL,
        CONSTRAINTS UNIQUE (acc#, cl#),
        amount INT,
        beg_date FROM bnp.open_date,
        duration INT)
```

The common format and constraints adopted for **beg_date** are upon those of **open_date** in **bnp**. The statement replaces as many SQL statements as there are databases in the current scope. If afterwards **loan** should be imported to a new database, like **boulogne**, the following statement would do it :

```
CREATE TABLE boulogne.loan FROM bnp.loan
```

2.9 Table expansion

MSQL extends the SQL ALTER TABLE statement as follows :

```
[USE < scope specifications> ]
```

```
ALTER TABLE <table designator> [ <table designator >...]
```

```
  ADD <column designator> {<data type> | FROM <column identifier> }
```

```
  MODIFY <column designator > {<data type> | FROM <column identifier>}
```

The statement adds a column to all tables designated in the databases that are in the scope or modify the data type of a column. The column schema may be imported. Designators may be unique or the multiple identifiers. They may also contain the SQL generic characters: '_', '%' or '*', designating then all tables in the scope whose names conform to the resulting pattern. This possibility as well as the new possibility of the list of table designators, are provided for the case that should be very frequent, where similar tables are named to some extent differently .

Ex. 2.5 The example banks decided to add **manager_name** column to their tables describing the branch offices. They may use the following single statement :

```
USE bnp cic sg
ALTER TABLE br% ADD manager_name CHAR(80)
```

The statement is equivalent to three SQL **ALTER TABLE** statements. An alternative statement for the same scope could be :

```
ALTER TABLE br branch ADD manager_name CHAR(80)
```

To import afterwards the column to *etoile*, assuming *br* created, one could use the statement :

```
ALTER TABLE etoile.br ADD manager_name FROM bnp.br
```

2.10 Table removal

MSQL **DROP TABLE** statement is as follows :

```
[USE < scope specifications > ]
```

```
DROP TABLE <table designator> [ <table designator >...]
```

As previously, unlike in SQL, one may use multiple identifiers and generic characters for table designation (the possibility of the list is new as well). One may thus simultaneously destroy several tables with the same or different names.

2.11 Multidatabase creation

A multidatabase is a set of interoperable databases or multidatabases. Some databases may be virtual, in the sense that their conceptual schemas are external schemas in the ANSI/SPARC sense. To declare a multidatabase means to name the set. The *multidatabase name* may then be used as the qualifier or as the single designator of the query scope. The declaration of a multidatabase is optional, but may obviously be useful. It may name databases of some federation, like cooperating banks, airlines etc. It may also designate some or all databases or multidatabases of a server etc.

The multidatabase creation statement has the following syntax :

```
CREATE MULTIDATABASE <multidatabase name>
```

```
( <(multi)database identifier > [, <(multi)database identifier >] ... )
```

```
<(multi)database identifier > ::= <database identifier > | <multidatabase identifier >
```

```
<database identifier > ::= <actual database identifier > | <virtual database identifier >
```

Identifiers are unique identifiers. Databases have to exist unlike in the **CREATE DATABASE** statement. A database or a multidatabase may belong to several multidatabases. As a table name, a (multi)database name may be qualified with the name of the multidatabase it belongs to. Different databases may indeed share a name, as already sometimes happens for databases of different servers.

Ex. 2.6 The bank databases may be declared as the multidatabase **Banks** :

```
CREATE MULTIDATABASE Banks (bnp cic sg )
```

The order has no importance.

2.12 Multidatabase modification

MSQL makes it possible to dynamically modify the multidatabase schemas through the following statement :

```
ALTER MULTIDATABASE <multidatabase identifier >  
    INCLUDE <(multi)database identifier > [, <(multi)database identifier >]... |  
    REMOVE <(multi)database identifier > [, <(multi)database identifier >]...
```

The 'or' is not exclusive, i.e. one may include and remove databases in the same **ALTER** statement. The database to be included must exist. The removed database is not affected with respect to its data or schema. It remains interoperable with databases in the multidatabase, but is no more in the scope of the multidatabase name.

Ex. 2.7 The bank Vernes wishes to join the **Banks** multidatabase with its database (or maybe its multidatabase), while **cic** wishes to leave it.

```
ALTER Banks  
    INCLUDE Vernes  
    REMOVE cic
```

The next **USE Banks** statement will open **Vernes**, but not **cic**.

3 DATA RETRIEVAL AND UPDATE

3.1 Statements

The general form of MSQL data manipulation statement is as follows :

<USE statement >

<auxiliary statement >

<SELECT statement > | <INSERT statement> | <DELETE statement> | <UPDATE statement> |

<REPLACE statement> | <STORE statement> | < COPY statement>

The <USE statement> opens (multi)databases, prior to any manipulation, except the creation. It defines the scope of the query. The auxiliary statements are introduced specifically for the multidatabase environment. The statements : SELECT,..., COPY are main statements.

3.2 USE statement

We call *queries* requests for retrievals or updates. SQL is designed for queries to a default database. MSQL is in contrast designed for manipulations of tables that may be explicitly in different databases. The SQL query scope is implicitly the default database. The MSQL query scope may be several databases. The USE statement defines this scope through the (multi)databases to be opened. In addition this statement makes it possible to define the aliases for database names, simplifying the qualification of table names.

USE [CURRENT][([< (multi)database name > [<alias>])] [([< (multi)database name > [<alias>])]]...

<alias>::=<character string>

A multidatabase name concerns all the corresponding databases. The existing scope is disregarded, unless the keyword CURRENT is used. In this case, the scope is only extended with the (multi)databases in this statement.

Ex. 3.1 The statements :

USE Banks

...

USE CURRENT (etoile e)

...

USE cic

opens the databases **bnp**, **cic**, **sg**, then adds **etoile** with **e** as alias, then restricts the scope to **cic**. If the multidatabase **Banks** was not declared, one should enumerate the databases.

3.3 Elementary queries

An MSQL query is an *elementary query* if all designators are unique identifiers of data types within the scope. The result of an elementary query is a relation, or an update to a database relation. SQL queries, like that in Ex. 2.1, as well as queries formulated using known relational DBSs are all *elementary monodatabase queries*. An *elementary multidatabase query* differs only in that it designates relations in different schemas. The **WHERE** statement of any elementary multidatabase query involves then interdatabase joins (if no implicit join introduced in Section 3.10 is allowed).

In the multidatabase environment, it may in particular occur that relations in different databases bear the same name. The basic new feature of an elementary MSQL query is that one may use then database names as prefixes for unique identification of such relations. Also the **FROM** clause may be omitted if no ambiguity results from it. Other new possibilities will be shown progressively.

Ex. 3.2 Retrieve all branches of **bnp** and **sg** that are in a same street.

```
USE Banks
SELECT brname, braname, street
FROM bnp.br, branch
WHERE bnp.br.street = branch.street
```

The table **br** is qualified to distinguish it from **cic.br**.

Ex. 3.3 Add 500 FFr bonus to clients in **bnp** who keep more money in this bank than in **sg**.

```
USE (bnp b) sg
UPDATE account
SET account.balance = account.balance +500
WHERE account.balance > acc.balance
AND b.client.cname = sg.client.cname AND b.client.street = sg.client.street
```

3.4 Multiple Queries

Multiple queries are queries extending the basic semantics of relational statements to multitables. They are essentially formulated as elementary queries, except that designators refer to multitables, instead of tables. The multitables are defined dynamically in the statement through multiple

identifiers or semantic variables /ABD83/, depending whether data names are the same or different. The result of a multiple query is a set of tables, unless a built-in function transforms it into a single table. The idea in multiple queries is to express in a single statement a manipulation broadcast to several databases, like, for instance the retrieval of all **Banks** branches in a given street or of some clients etc. Such queries have been also called *diffusion query* or *broadcast query* /LYN83/.

3.4.1 Multiple Identifiers

A (multiple) query with multiple identifiers is an equivalent of the set of pertinent (elementary) subqueries resulting from all the substitutions of the multiple identifiers by the corresponding unique ones. The substitutions are in the order of qualification : (multi)database names are chosen first, then the identifiers of tables in the chosen databases are set for the **FROM** clause if any, finally the names of columns of the identified tables are determined for the **SELECT** clause. A subquery is *pertinent* when it may be evaluated for the given schema(s). An MSQL subquery is considered pertinent if it designates existing data objects. In particular, the resulting **SELECT** clause must contain only column names of tables identified for the **FROM** clause (this is an implicit assumption for SQL as well). Other queries may sometimes be rendered pertinent through less or more complex transformations to equivalent pertinent queries /LIT82/.

Ex. 3.4 Retrieve all clients of **Banks**.

```
USE Banks
SELECT *
FROM client
```

The designator **client** is a multiple identifier of three tables. The query extends the semantic of the SQL **SELECT** statement to a multitable formed from tables bearing the same name, but not necessarily **union compatible**, as in the example. It replaces three (elementary) SQL queries whose execution order is immaterial. The result is the set of three working tables, each identified by the scope of the subquery (identifier(s) of the database(s) that the tuples come from).

Note that a client may be duplicated, up to three times, but that each tuple bears then obviously a different semantic information even if values are the same.

Ex. 3.5 Update street name 'Etoile' to 'Charles de Gaulle' in **Banks** and **etoile**.

```
USE Banks etoile
UPDATE *
SET street = 'Charles de Gaulle'
WHERE street = 'Etoile'
```

This statement updates all tables that have column **street**. In occurrence, the statement replaces seven SQL **UPDATE** statements. It could replace any number, provided that the cooperating banks agree on the common column name.

3.4.2 Semantic variables

A *semantic variable* is a variable whose domain are data names in the scope of the query, i.e names of columns, of tables or of (multi)databases. The domain may be explicit, which means that names are enumerated in the query. It may also be implicit, which means that it is computed from the variable name through some rules.

The aim of this concept is to enable the user to broadcast his intention over data named differently. A query may invoke several semantic variables, together with multiple identifiers. Each semantic variable means that the query concerns all the names in its domain. The names may be unique or multiple identifiers of some data names. The query is equivalent to the set of pertinent subqueries resulting from possible substitutions of semantic variables and multiple identifiers by unique identifiers. As above, the variables are substituted in the descending order of qualification.

The explicit semantic variables are introduced by one or more following clauses :

LET <semantic variable> **BE** <designator> [<designator>]...

<semantic variable> ::= <single variable> | <compound variable>

<single variable> ::= <character string>

<compound variable> ::= <single variable> . <single variable> [. <single variable>]....

<designator> ::= <data designator> [. <data designator>]

The data designators may include the SQL generic characters '%' for any sequence of $n \geq 0$ characters and '_' for any single character. The designator of the form (D.*) means "all columns of table designated by **D** or all tables of database **D** etc.". The semantic variable and the designators following it have to be of the same arity. The substitution rules are as follows :

LET **x** **BE** **city town**

x designates **city** or **town**.

```
LET x.y BE br.t# branch.tel
```

The couple (x, y) designates either (br, t#) or (branch, tel). Thus the substitution, for instance x <-- br and y <-- tel is not allowed. A similar rule holds for triplets, etc.

The implicit semantic variables are assumed to be only character strings with generic characters (note that identifiers may also be considered as trivial implicit semantic variables). Thus the rule for the domain determination is very simple. However, in general, it may be arbitrarily complex, including for instance the recognition that X = 'Ville' designates all columns in the scope named 'City' or 'Town', since the user is French. Implicit variables are used directly in the clauses as are all the other identifiers.

Ex. 3.6 Assuming the scope Banks, retrieve all branches at Av. Champs Elysées.

```
LET x BE br branch
SELECT *
FROM x
WHERE street = 'Av. Champs Elysées'
```

This query will be first decomposed into two queries by substituting to x respectively br and branch. The query with branch will be elementary, but the query with br will still be multiple, since br is a multiple identifier. This query will therefore further lead to two elementary queries. The result will consist of three elementary queries, to be executed in arbitrary order.

```
USE bnp
SELECT *
FROM br
WHERE street = 'Av. Champs Elysées'
```

```
USE sg
SELECT *
FROM branch
WHERE street = 'Av. Champs Elysées'
```

```
USE cic
SELECT *
FROM br
WHERE street = 'Av. Champs Elysées'
```

Using the implicit variable, the same query could be formulated as follows :

```
SELECT *
FROM br%
WHERE street = 'Av. Champs Elysées'
```

Both queries extend the semantic of the SQL SELECT statement to multitables of tables named differently which may also have different columns.

Ex. 3.7 Retrieve from **Banks** the names phone numbers and addresses of branches in department Hauts de Seine (Zip code starting with 92).

```
USE Banks
LET x.y.z.v BE br.tel.street#.city branch.t#.s#.town
SELECT %name, street, z, v, zip%
FROM x
WHERE zip% LIKE '92%'
```

Note the difference between the applications of '%', one to the column name evaluation and the other, classical to SQL, to the column value evaluation. The query generates three SQL queries.

Ex. 3.8 Note in **cltype** : "Special care" for any **bnp** client who is also a **sg** or **cic** client.

```
USE banks
LET b.t BE sg.ct# cic.cttel
UPDATE bnp.client
SET bnp.client.cltype = 'Special care'
WHERE bnp.client.clname = b.client.%name
      AND bnp.client.tel = b.client.t
```

The **WHERE** clause addresses a table and a multitable.

3.5 Options

SQL assumes implicitly that all the columns in the **SELECT** statement are in tables in **FROM** clause and, in general, in the schema of the addressed database. It is useful to relax this assumption in the multidatabase environment. The concept of *options* is intended for this purpose. The corresponding syntax is as follows.

Let **d** be a column designator within the **SELECT** statement. Let **q** be a subquery resulting from some substitutions and **a** the unique identifier corresponding to **d** in **q**.

- if **d** is preceded by a space, as is usual in SQL, then **q** is pertinent only if there is column **a** in its scope. Thus, by default **a** is mandatory.

- **d** written '~d' means that **q** may be pertinent without a column named **a** in the scope. **q** is then considered as equivalent to a query formulated like **q** without **a** in the **SELECT** list. The column **a** is thus optional.

- a list $d_1|d_2| \dots |d_n$ means that the pertinent form of q should contain one and only one a_i . The choice follows the list order. A list preceded with `~` means that the whole list is optional.

Ex. 3.9 Create a new account for Durand D. in the branch database **etoile** and in **bnp**.

```
USE bnp etoile
INSERT
INTO a% (a%, c%, ba%, ~br%#, ~open_date)
VALUES (123, 456, 789, 010, 11/04/87)
```

The value of **br#** will be inserted only to **bnp**, this of **open_date** only to **etoile**. If these columns were not optional, none of the insertions would be made since both subqueries would not be pertinent. The implicit semantic variables were useful only to shorten the query expression.

Ex. 3.10 Select from **Banks** the account number, client number, balance and opening date, if any, of all accounts.

```
USE Banks
LET x BE acc%
SELECT ac%#, c%#, balance, ~open_date
FROM x
```

Since the column **open_date** is optional, all three databases will be addressed. If **open_date** was mandatory, the tuples would be retrieved only from **bnp** and **sg** database.

Ex. 3.11 Assume that **sg** database does not have the column **t#** in the table **branch**. Retrieve from **Banks** branch names and either only the phone numbers if available or the corresponding streets.

```
USE Banks
LET x BE br%
SELECT %name, t%|street
FROM x
```

The query will provide the phone number from **bnp** and **cic**, and the street from **sg**.

Options deal with the existence of column names in schemas and not with null values within rows. However, one may extend this concept to null values as well.

3.6 Outer Joins

Inner joins eliminate the unmatched rows. MSQL makes it possible to keep such rows through **outer joins**, like also /SYB86/. This possibility is fundamental in the multidatabase environment /DAY84/. The outer join is symbolized with either '?' or '*' symbol for **natural outer join**. The symbol preceeds or follows a usual comparison operator, let it be β . Thus :

- left outer join is denoted : $\beta?$ or $*\beta$,
- right outer join is denoted : $\beta?$ or $\beta*$,
- and (full) outer join is denoted : $\beta\beta?$ or $*\beta*$

Note that the order of the outer join clauses in the WHERE statement is important if more than two tables are joined together, except for the natural outer join.

Ex. 3.12 Select name, phone number and balance of all **bnp** clients if their balance exceeds 1,000,000 FFr and their **sg** balance if any.

```
USE (bnp b) sg
SELECT cname, cltel, b.client.balance sg.client.balance
FROM b.client, account sg.client
WHERE b.client.cl# = account.cl# AND b.client.cname ** sg.client.cname
      AND b.client.cltel ** sg.client.ct#
      AND account.balance > 1 000 000
```

3.7 Column labels

As in most SQL dialects, a column designator in the MSQL SELECT statement may be followed by a label. The label heads then the corresponding column(s) values in the query result. This possibility is in particular useful for homogenizing column names as a result of a multiple query. If there is no label, the actual column name is displayed

Ex. 3.13 Select names, phone and street number of all branches at 'Champs Elysées'. Present them with unified headers : **branch_name**, **tel#** and **street#**.

```
USE Banks
LET t BE tel t#
SELECT %name branch_name, t tel#, s%# street#
FROM br%
WHERE street = 'Champs Elysées'
```

3.8 Union

The UNION operator eliminates duplicates from to union compatible tables and makes them a single

table. The corresponding syntax is as follows :

```
<SELECT statement > [UNION <SELECT statement >]... |  
<SELECT statement > UNION *
```

The SELECT statements in the first statement must define tables. The latter statement applies to a single multiple query.

Ex. 3.14 Select from **bnp**, **sg** and **cic** the names, phone numbers and addresses of all clients.

```
USE Banks  
LET x.y.z BE cname.cltel.street# cname.ct#.s#  
SELECT x client_name, y client_tel#, street client_street z street_  
FROM client  
UNION *
```

The result will be a single table with the labels as column headers.

3.9 Built-in functions

3.9.1 Extended SQL functions

SQL has some built-in functions that are COUNT, AVG, MAX, MIN and SUM. MSQl extends the use of these functions to elementary multidatabase queries and to multitables. In the former case, the function is used under its SQL name. In the latter case two situations occur :

- the function aggregates each table of the multitable. It then keeps its original SQL name.
- the function aggregates the whole multitable. It then has the same syntax, except that its name is prefixed with the letter M. For instance, AVG becomes MAVG.

In particular, one can use either DISTINCT or MDISTINCT.

Ex. 3.15 Retrieve the average balance per bank(i) and then the average balance of all banks(ii).

```
USE Banks  
SELECT AVG(ba%) (i)  
FROM a%
```

```
SELECT MAVG(ba%) (ii)  
FROM a%
```

Ex. 3.16 To find the concentration of millionaires near the Place de l'Etoile area, retrieve the number of millionaires per **bnp** and **etoile** (i) and then the total number of millionaires (ii).

```
USE bnp etoile (i)
LET x BE bnp etoile
SELECT COUNT(DISTINCT client.cl#)
FROM x.client x.account
WHERE x. client.cl# = x.account.cl#
AND x.account.ba% > 1 000 000
```

```
SELECT MCOUNT (MDISTINCT client.cl#) (ii)
FROM x.client x.account
WHERE x. client.cl# = x.account.cl#
AND x.account.ba% > 1 000 000
```

The keywords **DISTINCT** and **MDISTINCT** are needed, as a millionaire may have more than one large account.

Ex. 3.17 Retrieve the list of clients respectively from **bnp** and **sg** banks grouped (i) by the branch they belong to and (ii) by the client type only.

```
USE bnp sg (i)
SELECT c%name
FROM client x, br% y
GROUP BY br%name
HAVING x.c%# = acc%.c%#
AND acc%.br%# = y.br%#
```

```
SELECT c%name (ii)
FROM client
MGROUP BY c%type
```

3.9.2 Specific functions

The analysis showed also the need for specific new built-in functions /LIT86b/. The following functions should be found particularly useful.

MERGE

This function, called **NORM** in /LIT86b/ merges into one row all rows corresponding to the same real object in different relations of a multiple query. The merge uses the natural outer join over the designated columns. Its syntax is:

```
MERGE ON <column designator>[, <column designator>]...
```

Ex. 3.18 Select from all banks the accounts of all millionaires. If the client has an account in more

than one bank, then merge the corresponding data.

```
USE Banks
LET x.y BE cname.ctel cname.ct#
LET z BE Banks.*
SELECT *
FROM z.a%
WHERE z.a%.c%# = z.client.c%#
AND z.a%.balance > 1 000 000
MERGE ON x y
```

NAME function

Let **n** be a designator. Then **NAME(n)** provides the name of data designated by **n**, **NAME(.n)** provides the name of the container of data (table for a column, etc.) designated by **n**, **NAME(..n)** refers to **NAME(.NAME(.n))**, etc..

This function results from the need for relational operations not only on data values, but also on data names. It may be applied instead of a column name within **SELECT** and **WHERE** statements. The result is then considered as if the value of the column was the corresponding character string.

Ex. 3.19 Retrieve the names of branches in Nice and of banks they belong to.

```
USE Banks
LET x.y BE br.city branch.town
SELECT %name, NAME(x)
FROM x
WHERE y = 'Nice'
```

One sees that database names may bear logical information in the multidatabase environment.

UPTO function

This function limits the duplication of information that may come from several databases. For instance, a query to ten bank databases may at most ask for two references of a client. In particular the user may give priority to databases he trusts more than others. The function syntax is as follows:

```
UPTO (n (<A>) [<B>])
<A> ::= <column list>
<B> ::= <database list>
```

The function provides at most $n \geq 1$ tuples sharing the values of columns designated in the list **A**. Priorities correspond to the order of the list **B** that designates database names. **A**, **n** and **B** are optional. If **A** is not specified, the query processing stops after a non null response from **n** databases. The default value of **n** is 1. Finally, empty **B** means that the user has no preference.

Ex. 3.20 Retrieve five branches at 'Champs Elysées', preferably **bnp** branches if not **cic**.

```
USE Banks
UPTO (5 (br%) [bnp cic])
SELECT *
WHERE br%.street = 'Champs Elysées'
```

The result will come from **cic** database only if there are less than five **bnp** branches at Champs Elysées. If both banks have less than five branches at this location, then **sg** will be searched.

3.10 Incomplete queries

While formulating MSQL queries, the user may avoid specifying some equijoins. Basically, one may avoid equijoins linking primary or foreign keys, that share a domain. Such queries are called *incomplete queries*. A subquery of a multiple query may in particular be an incomplete query. Omitted joins are called *implicit joins* [LIT85]. They are deduced by the system from database schemas. The result is called the *complete query*. The completion algorithm is described in detail in [LIT85]. It is shown that this process leads to the intuitively expected result in more cases than the present algorithms for the universal relation interface [ULL83]. A major consequence is also that updates may be performed.

The goals of this function are (i) to further simplify query formulation, and (ii) to make the expression of a multiple query independent of differences in decomposition into tables. Otherwise, there is sometimes no way of formulating a single statement, if one has to indicate all equijoins corresponding to different decompositions.

Ex. 3.21 Retrieve from **bnp** the addresses of all branches where 'Dupond' has an account.

The incomplete query could be :

```
USE bnp
SELECT street
WHERE cname = 'Dupond' (Q1)
```

The complete query would be :

```

USE bnp
SELECT street
FROM br x, account y, client z
WHERE z.clname = 'Dupond'
AND z.cl# = y.cl#
AND y.br# = x.br#

```

(Q1')

Ex. 3.22 Consider now that instead of three tables, *cic* contains only one (universal) table with all columns in *cic*. Assume further that the user wishes to broadcast the query about "Durand" to both *bnp* and *cic*. The formulation (Q1) will then remain valid, provided both databases are open. The clauses will however define a multiple query. The query will be the equivalent of two subqueries differing by equijoins. These are (Q1') and the query :

```

USE cic
SELECT street
WHERE clname = 'Durand'

```

See also Ex. 5.4.

3.13. DYNAMIC COLUMNS

Similar data in autonomous databases may differ with respect to value types or meaning and may have a highly fluctuating relationship. The balance may be expressed in US\$ in one database and in FFr in another database, with the exchange rate varying constantly. Nevertheless, one may need to specify a join on such data, while the comparison between values in US\$ and in FFr is meaningless or one may need to convert all the selected values to the same currency etc. The multidatabase manipulation languages should therefore provide functions for the corresponding needs.

The concept of *Dynamic Attribute* is intended for such situations [LIT86a]. To match SQL terms, we speak in MSQL about *dynamic columns*. The user defines a dynamic column, let it be *D*, in a query. He then may use *D* as an actual column, except sometimes for updates. Unlike an actual or a virtual column, *D* is however unknown to the database schema or to any view schema and disappears with the query or at the end of the session. To define *D*, one specifies the *retrieval mapping*, let it be $M : C \rightarrow D$, for the conversion of the actual values of *source columns* *C*, to the dynamic ones, and/or an inverse *update mapping* $M' : D \rightarrow C$, if *D* encounters updates. The system may be able to automatically determine one of the mappings given the other, as does MRDSM, [VIG87]. See [LIT86a] for details of possibilities the concept of Dynamic Attribute opens. This concept responds also to some of the proposals in [DAT83].

A dynamic column is defined by D-COLUMN clause whose syntax is as follows :

D-COLUMN [HOLD] <definition>

<definition>::= <auto-conversion> | <user-specification>

<auto-conversion>::= (<Source column>) <D-name> FROM <Source-Unit> TO <Target-Unit>

<user-specif> ::= (< Source column >[,< Source column >]...) <D-name> [= <retrieval mapping>] [AND
[<Update column>] [= <update mapping>]]

<mapping> ::= <Arithmetical formula> | <SQL String function> |
 <SQL Date function> | <Table>

<Table> ::= (<Source value>[,< Source value>]...,<Target value>[, <Target value>]...) [(<Source value>[,<
Source value>]...,<Target value>[, <Target value>]...)...

The D-name is the name given to the dynamic column, let it be **D**. The keyword **HOLD** means that **D** definition holds for further queries, otherwise **D** may be referred to within the query only. The auxiliary commands **LIST HOLD** lists the set of hold names and **DROP HOLD <D-name>**, drops the specified column from the corresponding catalog. See /VIG87/ for deeper discussion of **HOLD** option.

The <auto-conversion> defines automatically the unit conversion mappings, when the system supports it /VIG87/. Otherwise, <user-specif> is the user definition of **D**. The update (retrieval) mapping may be omitted if one performs a retrieval (an update) only or when the system may deduce it from the other mapping /VIG87/. The update column has to be indicated when the retrieval mapping maps several source columns to **D** and does not allow by itself the deduction of the source values the **D** update should modify. Consider for instance the mapping :

month_sal = nb_days * day_sal.

Ex. 3.23 Consider that **citybank** has a database like **bnp**, except that the balance is in US\$. For **bnp** clients whose **bnp** amount is greater than **citybank** amount, retrieve client names and balances in to US\$, according to the exchange rate of the hour that is 6.5 FFr/US\$. Order the result by the converted value. Then, add to each **bnp** client the fidelity bonus of 10 US\$.

```
USE (citybank c) (bnp b)                                     (3.23.1)
D-COLUMN HOLD (b.balance)
    balance = b.balance / 6.5 AND balance = bal$ * 6.5
SELECT b.cname balance
FROM b.account, b.client
WHERE balance > c.balance
    AND b.client.cl# = b.account.cl# AND b.client.cname = c.client.cname
    AND c.client.cl# = c.account.cl#
ORDER BY balance
```

```

UPDATE account
SET balance = balance + 10
WHERE balance > c.balance
      AND b.client.cl# = b.account.cl# AND b.client.clname = c.client.clname
      AND c.client.cl# = c.account.cl#

```

(3.23.2)

The query (3.23.2) uses the scope and the dynamic column of (3.23.1). The update mapping may be omitted, if the system is able to deduce it. The dynamic column name has priority in the case of name conflict, as in these queries. The query (3.23.1) may be expressed in SQL using value expression a view or a derived table, but each time in much more procedural form. The query (3.23.2) cannot be expressed.

Ex. 3.24 Assume that as MRDSM, the system is able to find inverse mappings if the retrieval mapping is a polynomial formula. Consider that the client address is one column with client town, street etc. Retrieve the name, the town and balance in US\$ of clients that do not live in Paris nor Marseilles and the value of a variable $x = 2 \cdot \text{balance}^2 + 2 \cdot \text{balance}$, if $x = 24$. Order the result by town and ascending balance. Then set $x <- 40$.

Using the dynamic columns, the query may be expressed as follows :

```

USE bnp
D-COLUMN HOLD (addr) town =
      LTRIM (RTRIM (SUBSTR (addr, (INSTR( addr, " "))), " "), " ")
D-COLUMN (balance) x = 2*balance2 + 2*balance
SELECT town, clname, balance
FROM account, client
WHERE ( town <> "PARIS") AND ( town <> "MARSEILLES")
      AND ( x = 24 ) AND account.cl# = client.cl#
ORDER BY town, balance ASC

```

(3.23.3)

```

UPDATE account
SET x = 40
WHERE ( town <> "PARIS") AND ( town <> "MARSEILLES")
      AND ( x = 24 ) AND account.cl# = client.cl#

```

(3.23.4)

It is instructive to find SQL expression of (3.23.3). The query (3.23.4) illustrates the interesting problem that the above retrieval mapping has potentially two inverse mappings corresponding to the solutions of the equation :

$$2 \cdot \text{balance}^2 + 2 \cdot \text{balance} - 24 = 0$$

MRDSM chooses the inverse formula that corresponds to the actual value *before the update*. Thus if this value is 3 the update will lead to **balance = 4** else **balance = -5**.

4. VIEWS

4.1 View creation

A view in MSQL is a virtual table defined by the **CREATE VIEW** statement. As in SQL, this statement simply names and stores a selection expression. The query involving the view is processed through the query modification principle, extended to the multidatabase environment. To create a view may be useful when a query is frequent and complex or for security reasons. A *multidatabase view* is a view whose selection expression refers to several databases.

The statement form is as follows :

```
CREATE VIEW <view designator> [(<view column list>)]  
    AS < select statement>  
    [WITH CHECK OPTION]
```

```
<view designator > ::= [<database designator>.] <view name>  
<view column list> ::= <column name >[, <column name >...]
```

Any view belongs to some database which may be virtual. The statement creates view(s) named upon the <view name> either in each designated database or in each database in the current scope. To designate several databases one should use a semantic variable ranging over desired database names in the scope. If the view column list is used, then i-th view column name renames i-th column in the selection expression (if '*' is used then the order is this of the corresponding **CREATE TABLE** statement). Otherwise, the view columns are named upon these in the selection expression. The selection expression may include all the new functions, provided the result is a table (and not a multitable). It may also refer to other views. If the **WITH CHECK OPTION** is specified, then the created views may be updated.

In SQL, a view is implicitly a view of the database used and in this database. These notions should be carefully distinguished while using MSQL, where the view of one or more database may be created in different and explicitly named one or more databases. To create a view in a database **B**, basically means that the view is stored in **B** catalogs. The view is created in **B** only if **B** accepts it. An actual database may refuse a user view if it is a public database or system catalogs refuse references to tables in other databases. Such views, as well as any view, may be included into a virtual database, especially that of the view creator. The concept of the virtual database will be discussed in the next section.

Ex. 4.1 Create in a virtual database **my_bank** the view **same_street_branches** of **bnp** and **sg** branches at the same street.

```
USE my_bank bnp sg
CREATE VIEW my_bank.same_street_branches
      (bnp_name, bnp_s#, sg_name, sg_s#, street, city)
AS      SELECT brname, street#, braname, s#, street, city
      FROM br, branch
      WHERE br.street = branch.street
      AND br.city = branch.town
```

This view is a multidatabase view.

Ex. 4.2 Create in **bnp** and **sg** views **account\$** of selected data about client accounts in **Banks**, with homogenized naming and balance updatable and expressed in US\$ instead of FFr.

```
USE Banks
LET x BE bnp sg
CREATE VIEW x.account$ (cl#, name, street, acc#, balance$)
AS LET y BE Banks.*
      D_COLUMN (balance) balance$ = balance/5.95
      SELECT c%#, c%name, street, ac%#, balance$
      FROM y.a%, y.c%
      WHERE y. a%.ac%# = y.c%.ac%#
```

4.2 View removal

MSQL makes it possible to drop one or several views using **DROP VIEW** statement, similar to **DROP TABLE** statement.

```
DROP VIEW <view designator> [<view designator >...]
```

4.3 Virtual databases

A *virtual database* is a named set of views. One can use it as an actual database, except for eventual restrictions on updates. It may be useful for the following purpose :

- as an external schema (in ANSI/SPARC) on an actual database, defining the conceptual schema for the multidatabase usage, also called the *import schema* in /HEI85/. It may hide tables one wishes to keep private, rename tables and columns to the common naming schemes etc. on it almost as exactly as any other actual database.
- as the user perception of one or more databases, through mono or multidatabase views. The database is then the import schema in the sense of /HEI85/ and global schema in the sense of

Ingres/Star /ING86/. In particular, one may use then only SQL to manipulate the virtual data.

A virtual database is created, altered or dropped using the corresponding **CREATE DATABASE**, **ALTER DATABASE** and **DROP DATABASE** statements.

5. INTERDATABASE QUERIES

5.1 Introduction

Traditional queries basically transfer data between a database and the user workspace. An *interdatabase query* transfers data **between databases** /NIC86a/, /NIC86b/. The corresponding MSQL statements copy or move data from one or more *source tables* to one or more *target tables*. Both kinds of tables may be in different databases. The transfer may expand the target table by imported column schemas and values. The need for interdatabase queries often exists now, especially when mainframe databases cooperate with those on workstations. MSQL provides for this purpose new possibilities for the **INSERT** and **UPDATE** statements and specific statements termed **STORE**, **REPLACE** and **COPY**. For all statements, source data are selected using a selection expression including the possibilities discussed up to now and specific ones discussed below.

Different statements express different interactions between incoming data and existing ones, especially when the keys match. One may wish indeed to preserve the existing tuple, while another application may need the incoming tuple to replace the existing one, or one may even wish to replace the content of the target table as the whole. Furthermore, the statements provide new facilities for data conversion, as incoming data will usually need to fit a preexisting and different schema and may need to go to different schemas simultaneously. The conversion may concern (i) column value representation, units of measure or meaning (price without taxes to prices with taxes), (ii) column order or names and (iii) table names. MSQL assumes the value representation conversion to be automatic. Value expressions, dynamic attributes and built-in functions are available for units and meanings conversions. Functions for conversions (ii) and (iii) are described below.

5.2 INSERT statement

This statement adds rows to target tables and may add columns. The converted row is inserted if its key does not match the existing one (if no key is indicated all existing target columns are considered). Otherwise, only the values of the new columns are transferred. For other existing rows, these values are set to null. If the selection expression is multiple, the union of all

sub-queries results must provide only one row by key value. Otherwise, one cannot not know which row should be inserted.

The general form of **INSERT** statement is as follows:

```
[ < USE statement > ]                                     (5.1)
[ < LET statement > ]
INSERT [INTO] [< target table designator >] | [( <target column list> )] [WITH CONSTRAINTS]
    < source selection statement >

< source selection statement > :=
    [ < USE statement > ] [ < LET statement > ]
    SELECT | MOVE < source column designators >
    [<FROM clause >] [<WHERE clause>]
```

The target may be a multitable. The **MOVE** deletes the source tuples the selected values come from. The possibilities for column name and order matching are basically of two kinds :

- *by order*. (default option). The i-th selected column either in the source list or in the **CREATE** statement of the source table if the source list is '*', is mapped on i-th target column in either target column list or in the **CREATE** statement of the target table. Each list may include characters '\$', meaning that the corresponding source or target column is skipped.
- *by name*. Each selected column is mapped on target column with the same name or with the name of the source column label, if used. A label is mandatory with a source value expression. This possibility requires either '* *' or '*=*' in the target specifications as shown in the examples below, or at least one label in the source list.

The mapping by order (without the '\$' possibility) is already used by the SQL **INSERT** statement. For multiple target tables, it provides the independence with respect to target column names. The mapping by name is new and provide independence with respect to target column order. It further provides independence with respect to decomposition of incoming columns into target tables. This type of mapping is closer to the spirit of the relational model, since columns basically have no order. In the multidatabase environment, it is also frequently more useful. In particular, it limits mapping specification errors since column lists are shorter and not redundant.

The options may be used in both source and target lists. In addition to those already defined, the option '+a' means that if the target column a does not exist then it is created. If the **WITH CONSTRAINTS** option is used, then a inherits any constraint attached to its source column, unless

it contradicts the target data. If there is no '+a' option, then, for any target table, the only source values effectively mapped are those whose target columns exist. The choice may depend on the table if the tables designated by the target table designator differ with respect to columns. The '+' option may be nested with others, but only as the last element (ex.: a|b|c|+a) .

The column lists may contain elements of form **a : b**. They mean that all columns between **a** and **b** are implicitly involved in the order of the **CREATE TABLE** statement, unless a column is explicitly designated elsewhere in the statement. By default, **a** is the first column and **b** the last one.

Ex. 5.1 Add into **bnp** new clients from **etoile**.

```
USE bnp etoile
INSERT bnp.client ( :cltel, street: )
  SELECT *
  FROM etoile.client
```

The mapping is by order. In all new rows, **cltype** is set to null value.

Ex. 5.2 Move to **bnp.account** data in **open_date** column and all new accounts in **etoile**, **opera** and **nation** branches.

```
USE bnp etoile opera nation
INSERT INTO bnp.account ( :balance, +open_date )
  LET x BE etoile opera nation
  MOVE *
  FROM x.acc%
```

The column **br#** will be set to null in new target rows. The query is equivalent to one **ALTER TABLE** statement and three elementary interdatabase **INSERT** statements.

5.3 STORE statement

The only difference with respect to the **INSERT** statement, is that **STORE** refreshes all existing columns corresponding to incoming ones when the keys match. The target columns that do not correspond to incoming ones remain unchanged. The general form of the **STORE** statement is as (5.1) except that the command **STORE** replaces **INSERT**.

Ex. 5.3 Refresh in **bnp** new balances of all branches and add all new accounts. Then assuming some processing of the **bnp** accounts, in real life usually done overnight, send back all processed accounts.

```

USE bnp etoile opera nation
STORE bnp.account
  LET x BE etoile opera nation
  SELECT y.:balance, br.br#
  FROM x.acc% y, br
  WHERE br.brname = name(x)

```

```

USE bnp etoile opera nation
LET y BE etoile nation opera
STORE y.acc%
  USE bnp
  SELECT acc# %c#, balance
  FROM account br
  WHERE account.br# = br.br# and br.brname = name(y)

```

The target `open_date` values will be preserved, since no incoming column is mapped onto them. Note the use of '%' in the label, one could also apply the option `acc#|nac#`.

5.4 REPLACE statement

This statement erases all existing target rows and then inserts the incoming ones. The general form of the statement is as (5.1) except for the command itself.

Ex. 5.4 Replace all clients and accounts of `etoile` and `nation` by these of `bnp`.

```

USE etoile nation
LET db BE nation etoile
REPLACE (* *)
  USE bnp
  SELECT y.*, x.:, cltel tel
  FROM client x, account y, br
  WHERE x.cl# = y.cl# and y.br# = br.br# and br.brname = NAME(db)

```

The label is necessary for `cltel` mapping since the mapping by name is used. The source columns `cltype` and `br#` are not mapped, despite being selected. The query is a multiple interdatabase query leading to four elementary queries. The target tables are selected through the column name matching, regardless of table names. When target key columns do not match, the target table is not involved.

5.5 UPDATE statement

This statement does not create any new row, but only modifies existing columns using values of some source columns. The mapping results from the selection expression that has the `FROM`

clause, unknown to the SQL UPDATE statement, except for /SYB86/. The statement may copy or move the incoming values.

Ex. 5.5 Update **bnp.client.ctel** with the values from **etoile**, **nation** and **opera** DBs.

```
USE bnp etoile opera nation
UPDATE bnp.client
  LET x BE etoile nation opera
  SET ctel = %tel
  FROM x.c%
  WHERE bnp.account.cl# = x. c%.%cl#
```

5.6 COPY statement

The COPY statement duplicates selected source rows and schemas. It also duplicates the corresponding source constraints if WITH CONSTRAINT option is used and the constraints are preserved by the selection expression. The selection may involve value conversion and may change column names or order. The general form of this statement is as (5.1) except for the statement command.

Ex. 5.6 Copy to a new branch database **versailles** all corresponding accounts in **bnp**, creating the table **acc (nacc#, ncl#, balance)**.

```
USE (versailles v)
COPY v.acc
  USE bnp
  SELECT a.acc# nacc#, a.cl# ncl#, a.balance
  FROM account a, br
  WHERE a.br# = br.br# and br.city = 'versailles'
```

Ex. 5.7 Save **bnp** tables **account** and **client** into tables **a** and **c** in **save_bnp**.

```
USE bnp save_bnp
LET y.x BE a.account c.client
COPY y
  SELECT *
  FROM x
```

6. AUXILIARY DATA OBJECTS

6.1 Manipulation dependencies

Manipulation dependencies specify manipulations of some tables triggered by a given manipulation

of a given table. They are manipulated using the well known concept of triggers. Triggers are defined using the following statement, close to that of /SYB86/ :

```
CREATE TRIGGER < trigger name >
ON < table designator > <trigger definition >

<trigger definition > ::= FROM < trigger identifier > |
FOR { INSERT | UPDATE | DELETE },...
AS < program >
```

One may define any number of triggers on a table, provided their names differ. The table designator may designate a multitable, provided the tables are in different databases. The statement creates then the same trigger on each table. The program may be an arbitrary program or MSQL statements with, possibly, flow control statements **if**, **else**, etc. It may include the following statement that tests the update of the given column(s) :

```
IF UPDATE (<column designator> ) [AND | OR <column designator>],... <program>
```

Finally the program may refer to two dynamic tables **inserted** and **deleted** /SYB86/. The tuples of **inserted** are copies of those inserted or of these already updated by the statement that fired the trigger. Tuples in **deleted** are copies of those deleted from the table or of those updated, but with the before-update values.

The table designator may in particular refer to a table with a trigger that may trigger a manipulation of another trigger table etc. This possibility should be used carefully, as it may lead to loops or infinite executions. An implementation may also put limitations on the statement capabilities /SYB86/. The trigger is removed through the statement :

```
DROP TRIGGER < trigger identifier >
```

Ex. 6.1 The banks have decided to create in each database a table where data are diffused to prevent fraud :

```
crooks (bank, cname, acc# )
```

It was further decided that any insertion should go to any table, in order to allow any selection to use only the local table (thus **crooks** s tables are fully duplicated). The following trigger makes it

possible to designate the local table only, for both insertions and selections.

```
USE banks
CREATE TRIGGER warning FOR INSERT
AS
INSERT crooks FROM inserted
```

The table **inserted** is then automatically always that of the database whose **bad_clients** is designated in the **INSERT** statement that fired **warning**.

6.2 Equivalence dependencies

These dependencies declare the columns whose equal values identify the same real object. They are particularly useful for implicit joins, where they explicit the implicit interdatabase joins /LIT85/. The corresponding statement is as follows :

```
CREATE LINK < link name > ( < column designator > = < column designator >
[AND < column designator > = < column designators > ]...
[, < column designator > = < column designators >...])
```

The link is created for the databases in the scope. The **AND** clause indicates columns to be considered together, for instance client name and phone number, all columns on the same side of '=' being then in the same table. The commas separate parts of the dependency, for instance the equivalence between databases **etoile** or **bnp** and **nation**, and then that between **nation** and **opera** (see the example below). Parts and separate links may lead to transitivities, if they share columns.

The link may be altered or dropped using the corresponding statements with obvious syntax.

Ex. 6.2 Declare the equivalences of keys of clients and of accounts in BNP databases.

```
USE bnp opera etoile nation
CREATE LINK same_client (cl# = %cl#)
CREATE LINK same_account (acc# = nation.acc#, nation.acc# = opera.nac#)
```

The link **same_account** allows for instance the query asking for selection of **bnp.br#** and of **etoile.open_date** to omit in **WHERE** clause the interdatabase join on **acc#**.

6.3 Stored queries and transactions

Any MSQL query or transaction may be named, stored in source compiled form and then recalled for execution. One should use for this purpose the **PROCEDURE** statement whose syntax is :

```
PROCEDURE < procedure name > < parameters >  
< MSQL statements >
```

The semantic of the statement and the usage of procedures are as in SQL in /SQL86/, except that (i) MSQL statements may be used and (ii) the procedure name may be a multiple identifier. In the latter case, the statement duplicates the procedure in all the databases in the scope, providing the appropriate privileges.

8 CONCLUSION

As databases become easily accessible on computers and networks, users face multiple and autonomous databases. New functions for data manipulation languages are then needed, as the present languages were designed for manipulations of a single integrated database. MSQL offers such functions for relational databases. Numerous examples showed that the proposed functions should prove useful for a large variety of user needs. This is due to their flexibility and open nature with respect to the accommodation of autonomous names, values and structures. Unlike the classical database languages designed for a single truth, they are intended to face multiple different truths, disagreements, contradictions, but also cooperations, that all together are the most exciting aspects of the real life. Most of these functions are not yet available in other languages and systems.

MSQL extends the possibilities of SQL to a multidatabase environment and the proposed functions are designed for this environment. Several functions have nevertheless proved useful in a single database context. Thus, the concept of the multiple identifier may help in preserving referential integrity. Implicit joins simplify the formulation of most relational queries. Dynamic attributes are useful for applications where subjective or frequently changing value mapping rules render the traditional concept of view too static. It should thus be worthwhile to incorporate similar functions to any relational system.

The proposed functions lead to many open problems at the implementation level. In particular , the approach discussed in /CER84/ should be an interesting basis for multiple queries. Also, further functions may be added. Knowledge processing techniques should be investigated, as they seem particularly interesting in the multidatabase environment /DEY84/, /MAR86/. In particular, they should enlarge the class of intentions expressible as a single query and should make it possible to further simplify the expression of some queries.

References

- /ABD83/: A. Abdellatif, Multiple queries in the multidatabase system MRDSM (in French). Th., University of Tunis. INRIA, June 1983, 80p.
- /ANS75/: ANSI-SPARC Interim report on database management systems. Doc. 7514TS01, Washington DC, Feb. 1975
- /CER84/ S.Ceri and G. Pelagatti, Distributed databases principles and systems, Mc Graw-Hill, New-York, 1984
- /COD71/: E. F. Codd, A database sublanguage founded on the relational calculus ACM SIGFIDET, Nov. 1971, p.35-68
- /DAT83/: C. J. Date, A guide to DB2, Addison-Wesley publishing company, 1983
- /DAT84/: C. J. Date, A critique of the SQL database language. ACM SIGMOD Record 14, N3, Nov. 1984
- /DAT86/: C. J. Date, An introduction to database systems, Addison-Wesley publishing company, 4th ed, 1986
- /DAY84/: U. Dayal, M.G.Gouda, Using semiouter joins to process queries in multidatabase systems . ACM-PODS, Waterloo (Canada), Apr. 1984, pp. 153-162
- /DEY84/ Deyi, Li. A prolog database system. Research STudy Press, England, 1984.
- /HEI85/: D. Heimbigner and D. McLeod, Federated architecture for information management. ACM-TOIS, vol. 3, N3, 1985, pp. 253-278
- /ING86/: INGRES-STAR User's guide. Relational Technology Incorporation, 1986
- /LIT82/: W. Litwin, and al., Sirius systems for distributed data management. Distributed Databases, ed. H. Schneider, North-Holland, New York, 1982, pp. 311-366
- /LIT85/: W. Litwin, "Implicit joins in the multidatabase system MRDSM" IEEE-COMPSAC, Chicago, Oct.1985
- /LIT86a/ W. Litwin and P. Vigier, Dynamic attributes in the multidatabase system MRDSM, COMPDEC Los Angeles, Feb. 1986
- /LIT86b/: W. Litwin and A. Abdellatif, Multidatabase interoperability, Computer, Vol. 19, Nu. 12, Dec 1986, p. 10-18
- /LIT87/: W.Litwin and A. Abdellatif, An overview of the multidatabase manipulation language MDSL. To appear in IEEE special issue of the distributed database systems. May 1987, 22p.
- /LYN83/: P.Lyngbaek and D. LcLeod, "An approach to object sharing in distributed databases," VLDB 83, Florence, Oct. 1983, pp.364-376
- /MAR86/ Marcus, C. Prolog programming. Addison-Wesley. 1986. 325p.
- /MRD82/ MRDS Multics Relational Data Store, Reference manual, CII HB, January 1982
- /NIC86a/ B.Nicolas, A. Zeroual and W. Litwin, Interdatabase queries in multidatabase systems (in French), Journées Internationales de l'Informatique et de l'Automatisme - JIIA 86, ed. JIIA, Paris, June 17-20, 1986, p.233-243
- /NIC86b/ B. Nicolas and W. Litwin, Interdatabase queries in multidatabase systems, IEEE Computer Society's Workshop On Design Principles For Experimental Distributed Systems, Purdue University, West Lafayette, Indiana, October 16-17, 1986
- /SQL86/: Information processing systems - Database language SQL. Draft international standard ISO/DIS 9075, 114p.
- /SYB86/: Sybase database management. Transact-SQL User's guide. Doc. 3231-2.0 Sept.1986
- /ULL83/: J. D. Ullman, Principles of database systems, Second ed., Computer Science Press, Rockville, Md., 1983
- /VIG87/: Vigier Ph. and W. Litwin, New Functions for dynamic attributes in MRDSM

Example Databases

The databases used throughout the examples are supposed to be databases of the French banks BNP, Société Générale and CIC. The databases **bnp**, **sg** and **cic** are the main databases of the corresponding banks, on mainframes in real life. The **etoile**, **opera** and **nation** are databases of BNP branches, named upon their location inside Paris and used for local processing. They contain only data related to their customers and exchange them back and forth with the main database, usually during the night time. Some of their clients or accounts may be unknown to **bnp**, at least temporarily.

DB **bnp** :

br (br#, brname, street, street#, city, zipcode, tel)
account (acc#, cl#, balance, br#)
client(cl#, clname, cltel, cltype, street, street#, city, zipcode)
spe-acc (acc#, br#, cl#, balance, cur)

DB **sg** :

branch (bra#, braname, street, s#, town, zip, t#, class)
acc (acc#, bra#, c#, balance)
client (c#, cname, ct#, ctype, street, s#, town, zip)

DB **cic** :

br (br#, brname, street, street#, city, zipcode, tel)
account (ac#, br#, cl#, balance, open_date)
client(cl#, clname, cltel, cltype, street, street#, city, zipcode)

DB **etoile**:

account(acc#, cl#, balance, open_date)
client(cl#, clname, tel, street, street#, city, zipcode)

DB **nation**

acc(acc#, cl#, balance, open_date)
client(cl#, clname, tel, street, street#, city, zipcode)

DB **opera**

account(nac#, ncl#, balance, open_date)
customer(ncl#, cname, ctel, street, street#, city, zipcode ,ctype)

The tables within the main databases model respectively branches, clients and accounts in FFr. The table **spe-acc** in **bnp** contain accounts in foreign currencies, defined in **cur** column. The underlined columns are keys. The databases are to some extent semantically heterogenous with respect to names, structure and value types despite being all relational databases. This is because they are autonomous, as banks and branches compete for clients. Nevertheless, the banks and

branches also cooperate, the cooperation being naturally stronger between branches of the same bank. That is why data in different databases have also some similarities which are stronger inside the same bank. The whole situation leads to the following assumptions for the example :

1. Banks partly disagree upon the names modeling the same concepts.
2. A client may be in several databases.
3. Primary key values in databases of different banks are independent, despite sometimes the same column names.
4. However, they are the same for the same client, account or branch, inside BNP databases.
5. Furthermore, the banks consider that the same name and address or phone number identify the same client in any database.