

**Outer Joins and Filters
for Instantiating Objects from
Relational Databases through Views**

Byung Suk Lee
and
Gio Wiederhold

**TECHNICAL REPORT
Number 30**

May, 1990

Stanford University

Copyright © 1990 by
Center for Integrated Facility Engineering

If you would like to contact the authors please write to:

*c/o CIFE, Civil Engineering,
Stanford University,
Terman Engineering Center
Mail Code: 4020
Stanford, CA 95305-4020*

It is desirable to make the system to generate those left outer joins and filters as needed rather than requiring that a programmer specifies it manually as part of the query for every view definition. We develop such a mechanism in this paper.

Since left outer joins are not symmetric, they inhibit a query optimizer from attempting to reorder joins for more efficient query processing. Besides, application of non-null filters is not free. It incurs the cost of evaluating the corresponding selection predicates on a base relation. We show that these two operators can be avoided without affecting the query result for the cases we will define in this paper.

We made the following contributions in the context of instantiating objects from relational databases through views.

- To introduce the two key operators – a left outer join a non-null filter – for preventing information loss and the retrieval of unwanted information.
- To develop a simple mechanism of specifying those two operators in a relational view query, given a system model we define; The system model is easily implementable in existing systems.
- To address the efficiency issue of reducing the number of the two operators without affecting query results.

2 Background Framework

2.1 Integration of Object-oriented Programs and Databases

The desire for integrating object-oriented programs with databases has been increasing recently. This integration enables applications working in object-oriented environment to have shared, concurrent access to persistent storage. Examples are the engineering applications such as computer-aided design and computer-aided software engineering. These are not well supported by conventional databases such as relational databases.

We distinguish two alternative approaches to the integration of objects and databases: the *direct object storage* approach and the *indirect base relation storage* approach. In the object storage approach, an object-oriented model is used uniformly for applications and persistent storage [2, 3, 1, 4, 5, 6]; Objects are retrieved and stored as objects. In the relation storage approach, an object-oriented model is used for the applications while a relational storage model is used for persistent storage [7, 8, 9, 10, 11, 12], and objects are retrieved by evaluating queries to databases.

The relation storage approach incurs the overhead of mapping between different models [10, 13]. This additional cost is motivated for *large* databases since the relation storage approach supports *sharing* of different user views better than the object storage approach. Direct storage of objects is simple, but inhibits sharability [10]. For example, let us assume two users define Employee objects differently as Employee(name, salary) and Employee(name, department) respectively. In the object storage approach, the two Employee objects are stored separately. To provide sharing requires a separate mechanism for identifying the owners. In the relation storage approach however, this problem does not occur because the information to support the two Employee objects are stored in a single relation Employee(name, salary, department), and their owners are distinguished by the database view mechanism.

Outer Joins and Filters for Instantiating Objects from Relational Databases through Views

Byung Suk Lee
Electrical Engineering
Stanford University
Stanford, CA 94305
blee@cs.stanford.edu

Gio Wiederhold
Computer Science
Stanford University
Stanford, CA 94305
wiederhold@cs.stanford.edu

Abstract

One of the approaches for integrating an object-oriented programs with databases is to instantiate objects from relational databases by evaluating view queries. In that approach, it is often necessary to evaluate some joins of the query by left outer joins to prevent information loss caused by the tuples discarded by inner joins. It is also necessary to filter some relations with selection conditions to prevent the retrieval of unwanted nulls.

The system should automatically prescribe joins as inner or left outer joins and generate the filters, rather than letting it be specified manually for every view definition. We develop such a mechanism in this paper. To overcome the heterogeneity of an object-oriented model and the relational model, we first develop a rigorous system model. The system model provides a well-defined context for developing a simple mechanism.

The mechanism requires only one piece of information from users: null options on an object attribute. The semantics of these options are mapped to referential integrity constraints on the query result. Then the system prescribes joins and generates filters accordingly. We also address reducing the number of left outer joins and the filters so that the query can be processed more efficiently.

1 Introduction

One of the approaches for integrating object-oriented programs with relational databases is to generate objects from relational databases through views [10, 11, 12, 7, 8, 9]. A view is defined by a relational query and a function for mapping between object attributes and relation attributes. The query is used to materialize the necessary data into a relation from database, and the function is used to restructure the materialized relation into objects. This approach provides an effective mechanism for building object-oriented applications on top of relational databases.

In generating objects, some particular conditions arise that are not so common in traditional relational database operations. First of all, as will be shown in Section 3.2.1, it is often necessary to evaluate some joins of the query by *left outer joins* to prevent information loss caused by the missing tuples discarded by inner joins. It is also necessary to *filter* some relations with selection conditions which eliminate some tuples containing null attributes to prevent the retrieval of unwanted nulls.

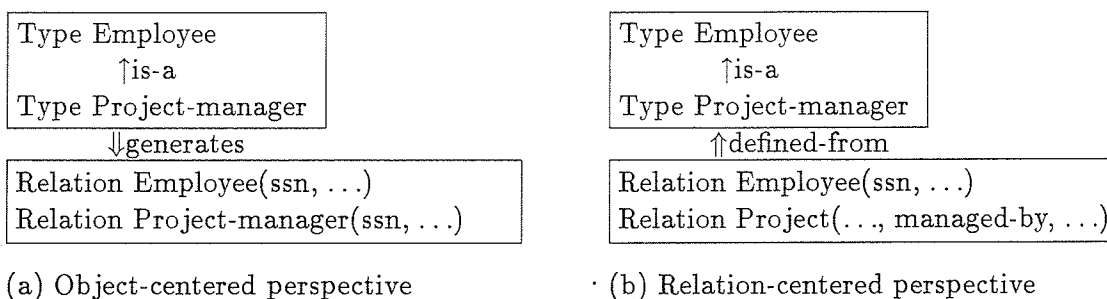


Figure 1: Two perspectives of relation storage approach

2.2 Two Perspectives of the Relational Storage Approach

We observed two different perspectives within the relation storage approach: *object-centered* [7, 8, 9] and *relation-centered* [10, 11, 12]. In object-centered perspective, relation schemas are generated from given object schemas, i.e., types and their hierarchy. Relations are the destination for storing objects, and objects are decomposed into relations using the concept of normalization. On the other hand, in *relation-centered* perspective, object schemas are defined from given relation schemas. Relations are the source for generating objects, and objects are composed from relations. The composition of objects is useful for building object-oriented applications on top of *existing* relational databases¹. The two perspectives may look like the two sides of the same coin, but they differ operationally. Figure 1 shows the two perspectives. In Figure 1a, the Project-manager type is mapped to the Project-manager relation. There exists a separate relation for each corresponding object type. In Figure 1b, there does not exist a separate Project-manager relation in the given database. Rather, the Project-manager type is defined as an *abstraction* through views, such as defining a join between the Employee relation and Project relation along the managed-by foreign key. The join retrieves only the employees who are managing one or more projects. Let us consider the Project-manager as an *abstract relation* of the Employee and Project relations. Note the abstract relation is analogous to the intensional database (IDB) relation [15, 16] used in the integration of the logic-based model and relational model [16, 17, 18]. For example, the IDB relation of the Project-manager is written as follows using the notion of Datalog [15].

$$\text{Project-manager(ssn, \dots)} \text{ :- Employee(ssn, \dots) \& Project(\dots, managed-by, \dots) \& ssn = managed-by.}$$

We use the relation-centered perspective throughout the discussion in this paper but the result is applicable to the object-centered perspective as well.

2.3 Instantiating Objects from Relations through Views

Views provide a user-defined subset of a large database. Thus, as mentioned in Section 2.1 and Section 2.2, views are used as a tool for providing sharing and abstraction in interfacing between an object-oriented model and the relational model. We also want to use the views for instantiating

¹We cannot throw away the relational data model in a decade. Remember that the IMS hierarchical data model implementation is still prevalent while we call the relational model ‘conventional’.

Database schema: /* Underlined attributes are keys. */
 Employee(ssn, e_name, salary, dept#)
 Department(dept#, d_name, manager_ssn)
 Child(ssn, c_name, sex, birth_date)

Type Employee /* [] denotes a tuple. */
 [name: string, dept: Department,
 children: [name: string, birthDate: string]]

View:

- Query expressed in relational algebra:

$$\Pi_{\{ssn, dept\#, c_name, birthDate\}} Employee \bowtie_{ssn=ssn} Child$$
- Mapping between object attributes and relation attributes:

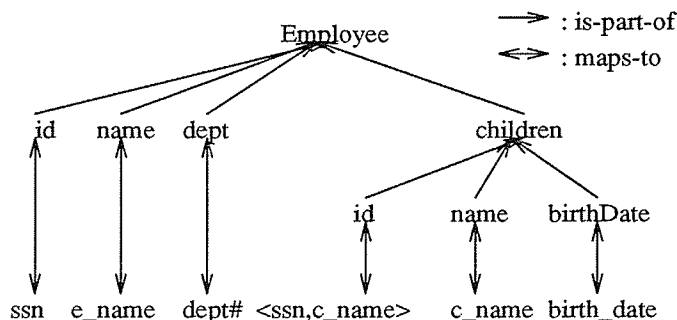


Figure 2: An example of instantiating an object type through views

objects from relations. To achieve this, views should provide mapping between heterogeneous structures of the two models. The mapping is done by linking object attributes to corresponding relation attributes. Objects have more complex structure than relations. For instance, objects support aggregation hierarchies [24] through an is-part-of relationship.² Hence objects have a nested structure, which is different from nested tuples because the type of an attribute can be a *reference* to another object. Therefore, given relation attributes, it is difficult to map the relation attributes to object attributes without explicitly specified mapping information. We thus need to extend the views by adding additional component for the mapping, that is, an *attribute mapping function*.

Figure 2 shows an example of instantiating objects through such an extended view. The object type defines the structure of objects to be retrieved from the database. The query part of the view specifies how to materialize the objects from the relational database. The join between the Employee relation and the Child relation has the semantics of nesting such as ‘For each Employee tuple, retrieve its name, dept#, and the c_name of the matching tuple in the Child relation.’ The outer relation is called a *source* relation and the inner relation is called a *destination* relation in our

²Objects also support a generalization hierarchy through is-a relationship, inheriting part of the attributes from parent objects. We regarded the inherited attributes as well as the local attributes uniformly as belonging to the objects.

work. The attribute mapping part of the view shows the aggregation hierarchy of object attributes and their mapping to relation attributes. The mapping is one-to-one as long as there is no derived attribute among the object attributes. We use the *key* attribute of one of the relations as the source of object identifier (id). In Figure 2, the key *ssn* of the *Employee* relation is retrieved to become the id of the *Employee* object. Object id's are not explicitly defined in the type definition but assumed to exist implicitly. The *dept* attribute of an *Employee* object has type *Department*. We call an attribute whose type is another object type as a *reference* attribute. In object-oriented paradigm, a reference is implemented with the id of the referenced object. In our framework, the value of a reference attribute is retrieved from the key of a database relation which is mapped to the id of the *referenced* object. Thus, in Figure 2, the *dept* attribute of an *Employee* object is retrieved from the *dept#* of the *Department* relation³. The *children* attribute defines a 'subobject' of the *Employee* object, and has its own attributes – name and birthDate. Like the *Employee* object, a *children* subobject is assumed to have its object id, but the object id is not actually retrieved from a database relation⁴.

3 Problem Formulation

3.1 The Two Operators

In the introduction, we mentioned the need of two operators: a left outer join and a non-null filter for instantiating objects from relational databases through views. A left outer join is different from an inner join in that it retrieves null tuples when there is no matching tuple in the destination relation for a given source relation. A non-null filter is a selection condition for eliminating any nulls of an attribute from a base relation⁵. Formal definitions of the left outer join and the non-null filter are as follows.

Definition 3.1 (Left Outer Join) Given two relations R_1 and R_2 , a left outer join from R_1 to R_2 , denoted by $R_1 \llcorner R_2$, is defined as follows.

$$R_1 \llcorner R_2 = (R_1 \bowtie R_2) \cup (R_1 - \Pi_{R_1}(R_1 \bowtie R_2) \times \Lambda) \quad (1)$$

where \bowtie denotes an inner join, $\pi_{R_1}(R_1 \bowtie R_2)$ denotes the projection of $R_1 \bowtie R_2$ on the attributes of R_1 , and Λ denotes a null tuple consisting of nulls for all attributes of R_2 . In other words, $R_1 \llcorner R_2$ produces the following set of tuples.

$$\{ \langle t_1, t_2 \rangle \mid t_1 \in R_1 \wedge ((t_2 \in R_2 \wedge t_1 \theta t_2) \vee t_2 = \Lambda) \} \quad (2)$$

where θ denotes the join condition.

For the rest of this paper, we use a small size join symbol (\bowtie) to denote a join which can be (has not yet been determined to be) either an inner join (\bowtie) or a left outer join (\llcorner).

³Let us assume there is a type *Department* whose object id is retrieved from the *dept#* of the *Department* relation.

⁴The id's of the *children* subobjects are needed for a different purpose, which will be discussed in Section 6.3.

⁵A base relation is the relation defined by the relation schema of a database, neither a view nor an intermediate relation.

Definition 3.2 (Non-null filter) A non-null filter is a conjunction of predicates applicable to a base relation R , defined as follows.

$$R.A_1 \neq \text{null} \wedge R.A_2 \neq \text{null} \wedge \cdots \wedge R.A_i \neq \text{null} \quad (3)$$

where A_1, A_2, \dots, A_i are the attributes of R that are not allowed to have nulls.

3.2 Motivation

3.2.1 Why do we need left outer joins and non-null filters?

Objects are identified by their identifiers (id's) only. In other words, an object exists even if all its attributes are nulls as long as it has an object id. Let us consider the objects of type Employee shown in Figure 2. An Employee object exists only if it has its id retrieved from the ssn of the Employee relation. Assuming that the Employee object allows null for its *children* attribute, what will happen if the join between Employee relation and Child relation is evaluated by an inner join? Any employee tuple that has no matching tuple in the Child relation will be discarded. In other words, any employee without children will not be retrieved. Therefore, it is certain we must evaluate the join by an *outer join* to prevent the loss of employees without children. What we need is not a bilateral outer join but a unilateral outer join because we are not interested in retrieving a Child tuple that has no matching tuple in the Employee relation, that is, a child without parent. Therefore, a left outer join is adequate assuming that the source, here the Employee, relation is the left hand side operand of the join. We assume the source relation is always on the left hand side of a join and thus use only left outer joins for the rest of this paper.

Now let us assume the Employee objects prohibit nulls for the dept attribute since a department affiliation is required of every employee. As mentioned in Section 2.3, the dept attribute is retrieved from the dept# of the Employee relation. The join between the Employee relation and Child relation is immaterial to the retrieval of dept# attribute. Rather, nulls of the dept# attribute stored in the tuples of the relation Employee should not be retrieved. Therefore, we must filter the Employee relation with a selection condition 'dept# \neq null'. We call this selection condition a *non-null filter*.

As explained with the above examples, we frequently need left outer joins [19] to prevent the loss of wanted objects, and non-null filters to prevent the retrieval of unwanted nulls.

3.2.2 Why do we want the system to do it?

Null-related semantics of object types are hard to understand and hence likely to induce errors. For example, the Employee type definition shown in Figure 2 does not distinguish between the semantics of 'employees and their zero or more children' and the semantics of 'employees with at least one child'. A left outer join is needed for the former while an inner join is needed for the latter. The distinction is entirely the programmer's responsibility. Even if the semantics is clear, it is an effort for the programmer to determine the left outer joins and non-null filters given a type and the corresponding view, especially if the view defines many joins. Therefore mechanization of the process will be useful.

3.2.3 Why do we want to reduce the number of left outer joins and non-null filters?

The query is processed more efficiently if we can eliminate a non-null filter ‘ $R.A \neq \text{null}$ ’ without affecting the query result, and thus avoid evaluating unnecessary selection conditions. Sometimes it is known at the semantic level that the column A of a relation R contains no null. An example is when A is the key of R and the entity integrity [20] is preserved.

The query also becomes more efficient if we reduce the number of left outer joins and still retrieve the same result. Sometimes left outer joins produce the same tuples as inner joins. For example in Figure 2, if every employee has one or more children, then the same tuples are produced by either join method. We know this fact at the semantic level, provided that the system enforces the referential integrity [20] from `Employee.ssn` to `Child.ssn`. As another example, let us consider the following directed join graph.

$$R_1 \longrightarrow R_2 \xrightarrow{\text{LO}} R_3 \longrightarrow R_4$$

where the join from R_2 to R_3 is a left outer join and the others are inner joins. If it is known there always exists a matching tuple of R_3 for every tuple of R_2 , then the result of $R_1 \bowtie R_2 \llcorner R_3$ is the same as $R_1 \bowtie R_2 \bowtie R_3$. Now, if we evaluate the join as an inner join, then the optimizer considers the three joins and will choose the most efficient order of joins. Let us assume the join order becomes $R_3 \rightarrow R_2 \rightarrow R_1$ in the optimal plan. On the other hand, if we evaluate the join as a left outer join, the query optimizer can not consider reversing the order of $R_2 \llcorner R_3$ and thus can not achieve the same optimal plan. In general, converting a left outer join to an inner join allows the query optimizer to deal with a larger number of joins. This increases the number of alternative plans but will certainly never generate less optimal plan than when left outer joins are evaluated as such and, therefore, cannot be reordered.

3.3 Problem Statement

Our problem is thus to develop a mechanism for the system to decide whether the joins of a query should be evaluated by inner joins or left outer joins when objects are instantiated from relational databases through views. In addition, the system decides which relations should be filtered through non-null filters. For efficiency reason, the number of left outer joins and non-null filters should be reduced whenever possible.

4 Our Approach

The heterogeneity of the object-oriented model and the relational model causes several difficulties in mapping between the two models [21]. Hence we cannot expect a simple solution without a well-defined system model. The system model should satisfy the following criteria.

- It provides the context in which we can develop a simple solution to the problem.
- It is based on a standard model and can be easily implemented in many existing systems.

Given the system model, we develop a mechanism for solving the problem. We use only one criterion that users should provide to the system. It is a *non-null option* on the object attribute

as will be explained in Section 5.1. Users do not even have to know what a left outer joins is. To prevent losing nonmatching tuples when nulls are allowed (by default), all joins of a query are initialized to left outer joins. The semantics of the non-null options are interpreted as *non-null constraints*⁶ on object attributes, and mapped to corresponding referential integrity constraints on the query result. Then we replace some joins by inner joins and add non-null filters to some relations accordingly. Finally, the number of left outer joins and non-null filters are reduced using the integrity constraints of the data model.

In the rest of the paper, we first develop the rigorous system model to facilitate the mapping between objects and relations in Section 5. The mechanism is developed in Section 6, and conclusion follows in Section 7.

5 System Model

The system model has three elements: an object type model, a view model, and a data model. The object type model defines the structure of objects. No object type model has gained universal acceptance [22, 23]. Therefore we define a model which is common to many existing object-oriented models [1, 6, 7, 4, 5]. The data model is the relational model proposed by Codd [14, 15]. The view model contains a relational query⁷ and defines a mapping between objects and relations. We restrict the query to an *acyclic select-project-join* query.

5.1 Object Type Model

Many existing object-oriented models [1, 6, 7, 4, 5] support *aggregation* through nested structure and references. For example, the Employee object of Figure 2 is an aggregation of name, dept, and children where dept is a reference to a Department object, and children is an aggregation of name and birthDate. The children attribute defines an embedded substructure of the Employee object. Thus our object type has a similar structure as the complex object [25, 26, 27].

We use value-oriented object id's [30, 31] and retrieve them from the keys of relations⁸. Those relations providing object id's are called *pivots* [11]. Sometimes an object is mapped semantically to an abstract relation rather than a base relation. Figure 3 illustrates these concepts. In Figure 3a, the Employee relation is the pivot for the Employee object and provides its key *ssn* as the object id. Figure 3b shows the abstract relation Project-manager of Figure 1, which becomes the pivot for the Project-manager object. It is defined by Employee \bowtie Project, and the key *ssn* of Employee in the join result is retrieved as the object id. ssn=managed-by


We do not consider derived attributes for our object type. Derived attributes have no direct mapping to relation attributes and, therefore, are computed separately from relation attributes.

An object type is defined formally as a tuple of attributes, $[A_1, A_2, \dots, X_1, X_2, \dots]$ where each


⁶These constraints require the existence of an object attribute given the id of an object. We would call this constraint as an *existence constraint* if this term were not already used in [15] to mean the same concept as the referential integrity.

⁷We do not assume the usage of any specific query language for our work.

⁸Tuple identifiers are usable as well. Otherwise we assume the system maintains a mapping between system-generated object id's and the keys of the corresponding relations.



(a) Pivot as a base relation



(b) Pivot as an abstract relation

Figure 3: The concept of a pivot

A_i is a simple attribute, and each X_i is a complex attribute⁹. An attribute is described in Backus-Naur Form as follows.

attribute ::= simple attribute complex attribute simple attribute ::= internal attribute external attribute complex attribute ::= [attribute, attribute, ...]

A *simple* attribute has an atomic value. It is either internal or external to the object. An *internal* attribute has a primitive data type such as string, integer, etc., while an *external* (or *reference*) attribute has another object type as its data type. The value of an external attribute is the id of the referenced object. A *complex* attribute defines a subobject by embedding its type definition within the object type. In the same way as an object id is mapped from the key of a pivot relation, a subobject also has an associated id which is mapped from the key of a base relation. However, the id of a subobject is *not* retrieved while the id of its (super)object is retrieved from the pivot key¹⁰.

We need a way of telling the system whether the value of an object attribute is allowed to be null or not. This is done by attaching a *non-null* option to an object attribute. This option deliberately declares that a null value is not allowed for the attribute. It is equivalent to specifying the constraint of ‘minimum cardinality > 0’ on the attribute¹¹. Attributes without non-null options are allowed to have null values by default.

An example is shown in Figure 4. The Project attribute defines its own attributes and becomes a subobject of the Programmer object. It has its object id mapped from a pivot key in the same way the Programmer object does. However, only the id’s of the Programmer objects are actually retrieved. This Programmer object example will be used throughout the rest of this paper.

Here we introduce two components derivable from the object type: *object set* (Oset) and *object chain* (Ochain). These will be used to facilitate mapping between objects and relations.

Definition 5.1 (Oset) Given an object O , $Oset(O)$ is defined as the set whose elements are the object O and all of its subobjects. The subobjects are recursively defined by nested complex attributes.

⁹Each attribute is either local to the object or inherited from its parent

¹⁰A subobject of an object is not a stand-alone object because it has no object id.

¹¹Many commercial tools for building object-oriented system applications, KEE[28, 29] for example, support this option.

```

Type Programmer
[ name: string non-null, dept: Department non-null, salary: integer,
  manager: Employee, task: string,
  Project: [ title: string non-null, sponsor: string, leader: string,
            depart: Department non-null ] ]

```

Figure 4: An example object type

For example, since the Programmer object contains one subobject Project, $Oset(Programmer) = \{ Programmer, Project \}$. Note each element of an Oset has its object id mapped to a pivot key.

Definition 5.2 (Ochain) Given an object O of type $[A_1, A_2, \dots, X_1, X_2, \dots]$, $Ochain(O, s_0)$ is defined as the chain of object-subobject relationships from O to an attribute s_0 , i.e., $O_0.O_1 \dots O_n.s_0$. Here $O_0 \equiv O$, O_i is a subobject of O_{i-1} for $i = 1, 2, \dots, n$, and s_0 an attribute of O_n .

For example, $Ochain(Programmer, title) = Programmer.Project.title$ and $Ochain(Programmer, Project) = Programmer.Project$.

5.2 Data Model

Integrity constraints are a part of the data model. Two kinds of integrity constraints are used in our work: referential integrity constraints and entity integrity constraints [20]. As mentioned in Section 3.2.1, these integrity constraints are useful to reduce the number of left outer joins and non-null filters.

The referential integrity constraint is defined as follows.

Definition 5.3 (Referential integrity) A referential integrity constraint from $R.A$ to $S.B$ requires that either $R.A$ be null or there exist a matching value of $S.B$ for every non-null $R.A$. That is:

$$\forall a \in R.A (a = \text{null} \vee \exists b \in S.B (a = b)) \quad (4)$$

Let us denote the referential integrity constraint by an arrow as in $R.A \rightarrow S.B$. Figure 5 shows the schema and referential integrity constraints of a sample database.

5.3 View Model

Figure 6 shows the components of the view model. A view consists of two parts: a query part and a mapping part. The mapping part in turn consists of an attribute mapping function (AMF) and a pivot description (PD). The AMF defines the mapping between object attributes (S_o) and relation attributes (S_r). The PD consists of a set of pivots (PS) and a pivot mapping function (PMF). The PMF defines the mapping between the pivots and the (sub)objects¹².

¹²Or equivalently, between the pivot keys and the id's of the (sub)objects.

```

/* Underlined attributes are keys. */
Division(name, manager, super-division, location)
Dept(name, budget, phone#)
Emp(ssn, name, salary, dept)
Engineer(ssn, degree, specialty)
Proj-Assign(emp, proj, task)
Project(proj#, dept, leader, sponsor)
Sponsor(name, phone#, address)
Proj-Title(proj#, title)

```

(a) Database schema

```

/* Arrows denote referential integrity constraints. */
Division.manager → Emp.name           Proj-Assign.emp → Engineer.ssn
Division.super-division → Division.name Proj-Assign.proj → Project.proj#
Dept.name → Division.name             Project.dept → Dept.name
Emp.dept → Dept.name                  Project.leader → Emp.ssn
Engineer.ssn → Emp.ssn                Project.sponsor → Sponsor.name
                                       Proj-title.proj# → Project.proj#

```

(b) Referential integrity constraints

The keys of all relations shown in the database schema are disallowed from having nulls. In addition, Emp.dept and Emp.name are prohibited from having nulls as well.

(c) Entity integrity constraints

Figure 5: A sample database

There can be designed a high level language for defining a view. The view should be preprocessed to generate the mapping components as well as the query.

5.3.1 Query Part

Figure 7 shows the query graph for the Programmer object. A query graph (QG) is a directed connected graph. Each vertex is represented by the node of a relation R labeled with a filter f and with the set of attributes π projected from R . Two occurrences of the same relation are distinguished by a tuple variable denoted as a subscript. Each edge represents a join specified in the query. A join is either an inner join or a left outer join. Since left outer joins are not symmetric, the edges are directed.

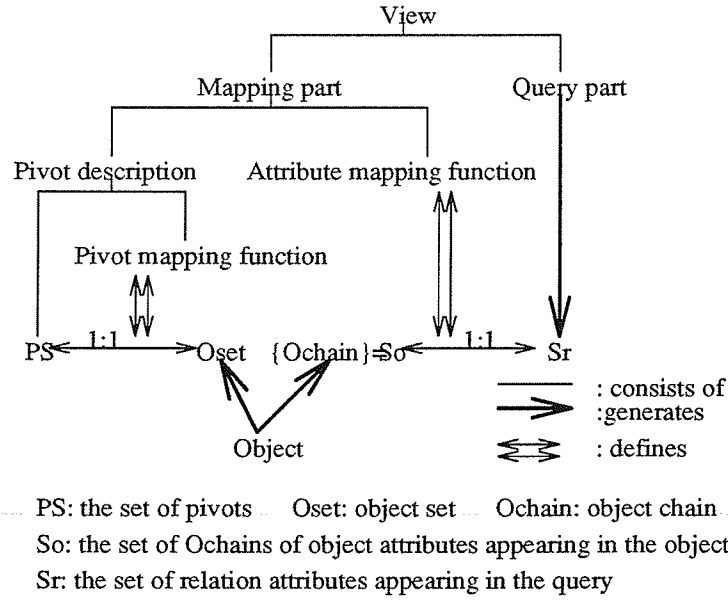


Figure 6: Mapping between objects and relations

5.3.2 Mapping Part

Now we give a more rigorous description of the mapping part. The set of object attributes S_o is represented as the set of Ochains as follows.

$$S_o = \{\text{Ochain}(O, s_o) \mid s_o \in \text{Attr}(O)\}$$

$\text{Ochain}(O, s_o)$ was defined in Definition 5.2. The set of relation attributes S_r is defined as follows.

$$S_r = \{R_i.A \mid A \subseteq \text{Attr}(R_i)\}$$

R_i denotes the i -th occurrence of the relation R .

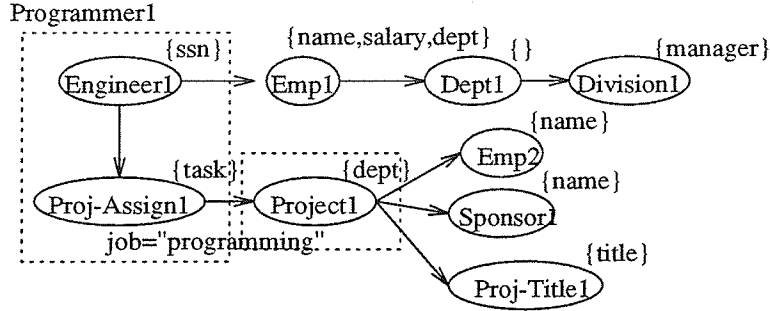
Since we assume no derived attribute, there exists a *one-to-one* mapping between S_o and S_r . This mapping information is contained in the attribute mapping function. The following example shows the mapping between the S_o and S_r of the Programmer object.

Example 5.1 (AMF)

```

Programmer.name ↔ Emp1.name,
Programmer.dept ↔ Emp1.dept,
Programmer.salary ↔ Emp1.salary,
Programmer.manager ↔ Division1.manager,
Programmer.task ↔ Proj-Assign1.task,
Programmer.Project.title ↔ Proj-Title1.title,
Programmer.Project.sponsor ↔ Sponsor1.name,
Programmer.Project.leader ↔ Emp2.name,
Programmer.Project.depart ↔ Project1.dept

```



(The keys of Engineer1 and Project1 are mapped to the id's of the Programmer object and the Project subobject respectively. Dotted lines denote pivots.)

Figure 7: The query graph for the Programmer object

As shown in Figure 3, a pivot is either a base relation or an abstract relation. If it is a base relation, its key is mapped to the object id. If it is an abstract relation, the key of one of its base relations is mapped to the object id. For example, the query for the Programmer object has two pivots, Programmer₁ and Project₁. Here Project₁ is a base relation and Programmer₁ is an abstract relation defined by $\langle \text{Engineer}_1, \{ \text{Engineer}_1 \bowtie_{\text{ssn}=\text{ssn}} \sigma_{\text{job} = \text{'programming'}} \text{Proj-Assgn}_1 \} \rangle$. A formal definition of an abstract relation is as follows.

Definition 5.4 An abstract relation of an object type O is an ordered pair $\langle R_b, E \rangle$ where R_b is a base relation whose key is mapped to the id of the object type O , and E is a *select-join*¹³ expression such that, for arbitrary instances of the relations in E :

- $\Pi_{\text{Key}(R_b)} E \subseteq \Pi_{\text{Key}(R_b)} R_b$
- $\neg \exists E' (E' \neq E \wedge \Pi_{\text{Key}(R_b)} E' = \Pi_{\text{Key}(R_b)} E)$

That is, the result of evaluating E produces a subset of the keys available from R_b and there is no other select-join expression E' which, when evaluated, produces the same set of keys.

For every object and its subobject, there always exists one and only one relation occurrence whose key is mapped to the id. In other words, there is a *one-to-one* mapping between the object set defined in Definition 5.1 and the set of pivots (PS). This mapping information is contained in the pivot mapping function. For example, the mapping between the Oset and PS of the Programmer object is as follows.

Example 5.2 (PMF) Programmer \leftrightarrow Programmer₁, Project \leftrightarrow Project₁

As mentioned in Section 5.1, we associate value-oriented object id's with an object and its subobjects. These id's are invisible in the type definition and their mappings to relation attributes are not explicitly specified in the attribute mapping function. These mappings are derived from the information stored in the pivot description using the following algorithm.

¹³Selection is not required while join is required.

Algorithm 5.1

```
For each pivot  $p \in PS$  begin
  If  $p$  is a base relation
    then append 'Ochain( $O$ , PMF( $p$ )).id  $\leftrightarrow$   $p$ .Key( $p$ )' to AMF.
  else /*  $p$  is an abstract relation */ begin
    Find the base pivot  $R_b$  of  $p$ .
    Append 'Ochain( $O$ , PMF( $p$ )).id  $\leftrightarrow$   $R_b$ .Key( $R_b$ )' to AMF.
  end.
end.
```

For example, given the set of pivots and the pivot mapping function of the Programmer view, Algorithm 5.1 derives the following mappings between the id's of the Programmer object and its Project subobject and their corresponding pivot keys. These are appended to the AMF.

{ Programmer.id \leftrightarrow Engineer₁.ssn, Programmer.Project.id \leftrightarrow Project₁.proj# }.

The attribute mapping function is essential for making it simple to map between objects and relations, as will be demonstrated in the following section.

6 Development of the Mechanism

Now we describe the mechanism for prescribing joins in a query as inner joins or left outer joins, and also for generating non-null filters for some relations in the query. We first present an overview of our mechanism, and then discuss each step in detail.

6.1 Overview

There are two source of nulls retrieved from databases. One is from the nulls stored in the tuples, the other is from any outer join failure. Inner joins create nulls from the first source only, while outer joins create nulls from both sources. Objects allow nulls by default, and need only one kind of outer join, left outer join, as explained in Section 3.2.1. Therefore our strategy is to initialize all joins of a query as left outer joins and then replace part of them by inner joins at each step of our mechanism.

The steps of our mechanism is as follows.

1. Compile the object type O and generate the object set (Oset) and the set of Ochain(O , s_0)'s for all the attributes defined in O .
2. Preprocess the view and generate the query and the mapping part: AMF, PMF, and PS.
3. Derive the mappings between object id's and pivot keys using Algorithm 5.1, and add the result to the attribute mapping function.
4. Initialize all joins of the query as left outer joins.
5. Replace all joins within abstract relations by inner joins. (See Section 6.2.)

6. Map non-null options on object attributes to non-null constraints on the query result. Replace some joins by inner joins and add non-null filters to some relations accordingly. (See Section 6.3 and Section 6.4.)
7. Find the left outer joins which produce the same tuples as inner joins due to referential integrity constraints, and replace those left outer joins by inner joins. Find also the relations whose non-null filtered attributes cannot have nulls due to entity integrity constraints, and remove the non-null filters from those relations. (See Section 6.5.)

6.2 Joins within an Abstract Relation

As mentioned in Section 2.2, an abstract relation is a conceptual relation derived from base relations via a select-join expression, and provides an abstraction of base relations so that the semantics of the abstract relation directly matches the semantics of the instantiated objects.

All joins specified within an abstract relation must be *inner* joins, as shown by the following theorem.

Theorem 6.1 Let us consider an object type O and an abstract relation $\langle R_1, E \rangle$ defined according to Definition 5.4. If $E = R_1 \bowtie R_2 \bowtie \dots \bowtie R_n$, then all the joins from R_1 through R_n are inner joins.

Proof: If we assume a join from R_i to R_{i+1} is a left outer join for an arbitrary $i \in [1, n]$ while the others are inner joins, then the following is true.

$$\Pi_{\text{Key}(R_1)}(R_1 \bowtie R_2 \bowtie \dots \bowtie R_i \text{ } \bowtie \text{ } R_{i+1} \bowtie \dots \bowtie R_n) = \Pi_{\text{Key}(R_1)}(R_1 \bowtie R_2 \bowtie \dots \bowtie R_i) \quad (5)$$

That is, there exists another select-join expression which, when evaluated, produces the same set of keys available from R_1 . This violates the second condition required of E in Definition 5.4. Therefore, all the joins in E must be inner joins. Q.E.D.

For example, given an abstract relation $\langle \text{Engineer}_1, \{ \text{Engineer}_1 \underset{\text{ssn}=\text{ssn}}{\bowtie} \sigma_{\text{job} = \text{'programming'}} \text{Proj-Assign}_1 \} \rangle$ defined to provide the semantics of the Programmer object, the join between Engineer_1 and Proj-Assign_1 must be an inner join. If the join is evaluated as a left outer join, it retrieves all tuples of Engineer_1 , not just those corresponding to programmers, who are defined as the engineers working on a programming job in the assigned projects.

Thus, given the set of pivots (PS):

Algorithm 6.1

1. For each abstract relation $\langle R_b, E \rangle$ in the set of pivots (PS),
replace all joins in E by inner joins.

6.3 Mapping Non-null options to Non-null Constraints on the Query Result

Let us consider an object O whose attribute s_0 has a non-null option. It requires there should exist a non-null s_0 given the id of the object. Let us denote this non-null constraint as $O.\text{id} \Rightarrow s_0$. If

s_0 is a simple attribute, it is non-null if its value is not null. On the other hand if s_0 is a complex attribute, it defines a subobject. An object is non-null only if its id is non-null. We thus interpret the semantics of non-null s_0 according to the following rule of non-null constraint.

Rule 6.1 (Non-null constraint) Let us denote $\text{Ochain}(O, O_n) \equiv O_0.O_1 \cdots O_n$ by $\Omega_{0,n}$ where O_n is the (sub)object containing s_0 as its attribute or subobject. If s_0 has a non-null option then, given $O_n.\text{id}$,

- If s_0 is a simple attribute, i.e., $O_n.\text{id} \Rightarrow s_0$, then s_0 cannot be null.
- If s_0 is a complex attribute, i.e., $O_n.\text{id} \Rightarrow s_0.\text{id}$, then $s_0.\text{id}$ cannot be null.

For example, given the Programmer object of Figure 4, the non-null options on name and dept attributes are interpreted as $\text{Programmer.id} \Rightarrow \text{name}$ and $\text{Programmer.id} \Rightarrow \text{dept}$, respectively, because name and dept are simple attributes. Besides, the non-null options on title and depart are interpreted as $\text{Project.id} \Rightarrow \text{title}$ and $\text{Project.id} \Rightarrow \text{depart}$, respectively. Beware they are *not* interpreted as $\text{Programmer.id} \Rightarrow \text{title}$ and $\text{Programmer.id} \Rightarrow \text{depart}$ because title and depart are the (direct) attributes of Project subobject instead of the Programmer object. On the other hand, if there were a non-null option on Project, it would be interpreted as $\text{Programmer.id} \Rightarrow \text{Project.id}$ because Project is a complex attribute.

Can we map the non-null constraint defined by Rule 6.1 to the corresponding non-null constraint on the query result? It is possible in our model because the id of each (sub)object always has a corresponding pivot key. The attribute mapping function in Example 5.1 showed this correspondence for the Programmer object. Using the correspondence, the non-null constraints on the name and dept attributes of the Programmer object are mapped to $\text{Engineer}_1.\text{ssn} \Rightarrow \text{Emp}_1.\text{name}$ and $\text{Engineer}_1.\text{ssn} \Rightarrow \text{Emp}_1.\text{dept}$, respectively. Likewise, if Project had the non-null option, its constraint would be mapped to $\text{Engineer}_1.\text{ssn} \Rightarrow \text{Project}_1.\text{proj\#}$. The non-null option on the title attribute is mapped not to $\text{Engineer}_1.\text{ssn} \Rightarrow \text{Proj-Title}_1.\text{title}$ but to $\text{Project}_1.\text{proj\#} \Rightarrow \text{Proj-Title}_1.\text{title}$ because title is defined not as an attribute of Programmer object but as an attribute of Project subobject. For the same reason, the non-null option on the depart attribute of Project is mapped to $\text{Project}_1.\text{proj\#} \Rightarrow \text{Project}_1.\text{dept}$.

More formally, a non-null option on the attribute s_0 of an object type O is translated into the non-null constraint on the query result as follows.

Algorithm 6.2

1. $\Omega_{0,n}.s_0 := \text{Ochain}(O, s_0) \equiv O_0.O_1 \cdots O_n.s_0$.
2. $R_p.A := \text{AMF}(\Omega_{0,n}.\text{id})$. /* A is always the key of R_p . */
3. If s_0 is a simple attribute
 then $R_s.B := \text{AMF}(\Omega_{0,n}.s_0)$
 else $R_s.B := \text{AMF}(\Omega_{0,n}.s_0.\text{id})$. /* If s_0 is a complex attribute, B is the key of R_s . */
4. Output the constraint ' $R_p.A \Rightarrow R_s.B$ '.

6.4 Prescribing Joins and Generating Non-null Filters

With the non-null constraints on the query result, we translate them into the corresponding inner joins and non-null filters of the query. Given the constraint ' $R_p.A \Rightarrow R_s.B$ ' obtained from Algorithm 6.2, it is done as follows.

Algorithm 6.3

1. Replace the filter f_s on R_s by $f_s \wedge (B \neq \text{null})$. /* Generate a non-null filter. */
2. /* Prescribe a join. */
 - (a) Find all directed join paths from R_p to R_s .
 - (b) For each path found in Step 2a,
replace all joins on the path by inner joins.

For example, given the non-null constraints established in Section 6.3, the following non-null filters are generated in the query of the Programmer object: $\text{Emp}_1.\text{name} \neq \text{null}$, $\text{Emp}_1.\text{dept} \neq \text{null}$, $\text{Project}_1.\text{dept} \neq \text{null}$, $\text{Proj-Title}_1.\text{title} \neq \text{null}$. Besides, the following left outer joins are replaced by inner joins: $\text{Engineer}_1 \bowtie \text{Emp}_1$, $\text{Project}_1 \bowtie \text{Proj-Title}_1$.

Now we prove the correctness of Algorithm 6.3 with the following theorem.

Theorem 6.2 Given a join path $R_1 \bowtie R_2 \bowtie \dots \bowtie R_n$ and a non-null constraint $R_1.A_1 \Rightarrow R_n.A_n$ on the join result, the materialized join result satisfies this non-null constraint if and only if all the joins are inner joins and R_n is filtered by $A_n \neq \text{null}$.

Proof:

If part: If all joins on the join path are inner joins, any nonmatching tuples are discarded. Then, the attribute A_n in the join result can have nulls only if A_n is *not* a join attribute and some tuples of R_n have null A_n . (If it is a join attribute, any tuple of R_n with null A_n is discarded by an inner join.) However, tuples with null A_n are removed from R_n by the given non-null filter. Therefore the constraint is satisfied.

Only if part: We prove this part by contradiction. Let us first assume $R_i \bowtie R_{i+1}$ is a left outer join for some i although the constraint is satisfied and let R_{i+1} have non-matching tuples. Then a null $R_n.A_n$ is retrieved from the null tuples appended to the tuples of R_i which have no matching tuples in R_{i+1} . This contradicts the assumed constraint. Therefore all the joins must be inner joins. Next, let us assume R_n is *not* filtered by $A_n \neq \text{null}$ although the constraint is satisfied and all joins are inner joins. Then null $R_n.A_n$ is retrieved from the nulls stored in $R_n.A_n$ if A_n is not a join attribute. This contradicts the assumed constraint. Q.E.D.

6.5 Reducing the Number of Left Outer Joins and Non-null Filters

We can further reduce the number of left outer joins and non-null filters by using integrity constraints.

Considering entity integrity constraints, some non-null filters are removed if they are defined on attributes which cannot have null. A typical case is when the attribute is a key or any other

non-null attribute designated in the schema definition. For example, we can remove $\text{Emp}_1.\text{name} \neq \text{null}$ and $\text{Emp}_1.\text{dept} \neq \text{null}$ among the four non-null constraints generated in Section 6.4 because, as it was shown in Figure 5c, those two attributes are key attributes and hence prohibited from having nulls.

We can also replace some left outer joins by inner joins if we consider referential integrity constraints. Since a referential integrity $R.A \rightarrow S.B$ allows $R.A$ to be null, we define a stronger condition by introducing a variable min as follows.

Definition 6.1 (min) Given a join $R_i \bowtie R_j$, let min_{ij} denote the minimum number of matching tuples in R_j for each tuple in R_i . Note min_{ij} is not necessarily the same as min_{ji} .

Using only the semantics of min without considering the instances of relations¹⁴, we define the following rules for deciding whether min is greater than zero or not.

Rule 6.2

- Given a single join predicate $A\theta B$ for the join between two relations R_i and R_j , $\text{min}_{ij} > 0$ if the join is an equijoin ($\theta = '='\text{'}$) and $R_i.A$ is a non-null attribute and $R_i.A \rightarrow R_j.B$ and the filter f_j on R_j is empty. Otherwise $\text{min}_{ij} = 0$.
- Given a *conjunctive* join predicate $A_1\theta_1B_1 \wedge A_2\theta_2B_2 \wedge \dots \wedge A_k\theta_kB_k$ for the join between R_i and R_j , $\text{min}_{ij} > 0$ for the conjunction of join predicates if $\text{min}_{ij} > 0$ for *every* single join predicate. Otherwise $\text{min}_{ij} = 0$.
- Given a *disjunctive* join predicate $A_1\theta_1B_1 \vee A_2\theta_2B_2 \vee \dots \vee A_k\theta_kB_k$ for the join between R_i and R_j , $\text{min}_{ij} > 0$ for the disjunction of join predicates if $\text{min}_{ij} > 0$ for *at least one* join predicate. Otherwise $\text{min}_{ij} = 0$.
- Given a join path between two relations, such as $R_i \bowtie R_{i+1} \bowtie \dots \bowtie R_j$, $\text{min}_{ij} > 0$ if $\text{min}_{k,k+1} > 0$ for every join on the path. Otherwise $\text{min}_{ij} = 0$.

Now assuming $\text{min}_{ij} > 0$ for a join path between R_i and R_j , can we replace all joins on the path by inner joins and still get the same query result? The answer is no. The case analysis in Figure 8 shows us why. Five exhaustive cases are shown for a join path between R_1 and R_3 . The joins from R_1 through R_3 have $\text{min}_{ij} > 0$ for all cases. All cases except Case 1 also show a join from R_3 to R_4 , whose min_{34} is either greater than 0 or equal to 0 and which is either an inner join or a left outer join. We see that, for all cases except Case 5, no tuple of R_3 is discarded from the join between R_3 and R_4 , and hence, the materialized join results are the same whether the joins from R_1 through R_3 are inner joins or left outer joins.

Now we describe an algorithm for reducing the number of left outer joins using min .

Algorithm 6.4

¹⁴In other words, we ignore the fact that min may be accidentally zero at the instance level although it is judged to be greater than zero at the semantic level.

1. Case 1: $R_1 \xrightarrow{\min>0} R_2 \xrightarrow{\min>0} R_3$
2. Case 2: $R_1 \xrightarrow{\min>0} R_2 \xrightarrow{\min>0} R_3 \xrightarrow{\min>0,LO} R_4$
3. Case 3: $R_1 \xrightarrow{\min>0} R_2 \xrightarrow{\min>0} R_3 \xrightarrow{\min>0,I} R_4$
4. Case 4: $R_1 \xrightarrow{\min>0} R_2 \xrightarrow{\min>0} R_3 \xrightarrow{\min=0,LO} R_4$
5. Case 5: $R_1 \xrightarrow{\min>0} R_2 \xrightarrow{\min>0} R_3 \xrightarrow{\min=0,I} R_4$

(Each edge represents a join and is labeled with either $\min > 0$ or $\min = 0$. The edge between R_3 and R_4 is additionally labeled with 'LO' for the left outer join or 'I' for the inner join.)

Figure 8: Case analysis of a join path

1. Find all join paths between pairs of nodes, such as R_i and R_j , that satisfy the following conditions.
 - $\min_{ij} > 0$.
 - There does not exist an inner join from R_j to another node R_k not on the same join path such that $\min_{jk} = 0$. (This is to exclude Case 5 of Figure 8.)
2. For each join path found in Step 1, replace all joins on the path with inner joins.

For example in the query of Programmer object, we find a join path from $Engineer_1$ to $Division_1$ for which all three joins have $\min > 0$ because, as shown in Figure 5, there are referential integrities $Engineer_1.ssn \rightarrow Emp_1.ssn$, $Emp_1.dept \rightarrow Dept_1.name$, $Dept_1.name \rightarrow Division_1.name$, and there are integrity constraints prohibiting nulls for $Engineer_1.ssn$, $Emp_1.dept$, and $Dept_1.name$, and none of the relations on the join path has a non-empty filter. We also find a join path from $Proj-Assign_1$ to $Project_1$ for which the $\min > 0$. All these joins are replaced by inner joins. Note $Project_1 \bowtie Emp_2$ and $Project_1 \bowtie Sponsor_1$ can not be replaced with inner joins because $Project.leader$ and $Project.sponsor$ are not non-null attributes.

6.6 Summary of the Mechanism

Given a query with initial left outer joins, the overall mechanism developed in Section 6 is as follows.

Algorithm 6.5

1. /* Replace all joins within abstract relations with inner joins. */
For each abstract relation $\langle R_b, E \rangle$ in the set of pivots (PS),
replace all joins in E by inner joins.

2. For each attribute s_0 of the object O that has a non-null option,
 - (a) /* Map the non-null option to a non-null constraint on the query result */
 - i. $\Omega_{0,n}.s_0 := \text{Ochain}(O, s_0) \equiv O_0.O_1 \cdots O_n.s_0$.
 - ii. $R_p.A := \text{AMF}(\Omega_{0,n}.id)$. /* A is always the key of R_p . */
 - iii. If s_0 is a simple attribute
 - then $R_s.B := \text{AMF}(\Omega_{0,n}.s_0)$
 - else $R_s.B := \text{AMF}(\Omega_{0,n}.s_0.id)$. /* If s_0 is a complex attribute, B is the key of R_s . */
 - iv. Output the non-null constraint ' $R_p.A \Rightarrow R_s.B$ '.
 - (b) /* Generate a non-null filter and prescribe a join. */
 - i. Replace the filter f_s on R_s by $f_s \wedge (B \neq \text{null})$. /* Generate a non-null filter. */
 - ii. /* Prescribe a join. */
 - A. Find all directed join paths from R_p to R_s .
 - B. For each path found in Step 2(b)iiA,
 - replace all joins on the path by inner joins.
3. /* Remove all non-null filters which can be shown to be redundant using the entity integrity constraint. */

Remove ' $R.A \neq \text{null}$ ' such that A is a non-null attribute.
4. /* Replace left outer joins if they prove to be equivalent to equijoins.*/
 - (a) Find all join paths between pairs of nodes, such as R_i and R_j , that satisfy the following conditions.
 - $\min_{ij} > 0$.
 - There does not exist an inner join from R_j to another node R_k not on the same join path such that $\min_{jk} = 0$.
 - (b) For each join path found in Step 1,
 - replace all joins on the path with inner joins.

The graph of the query for the Programmer object, labeled with joins and non-null filters, is shown in Figure 9. All the joins of the query except those between Project_1 and Emp_2 and between Project_1 and Sponsor_2 have been prescribed as inner joins. Two non-null filters have been attached as the selection conditions on the Project_1 and Proj-Title_1 nodes.

7 Conclusion

We developed a mechanism for automatically prescribing inner or left outer joins for the joins of a query used to instantiate objects from a relational database. It also generates non-null filters for some of the relations in the query. We developed a rigorous system model that facilitates the mapping between object and relations. The system model consists of an object type model, a view model, and a relational data model. These models are based on a standard model or well-known models. We added a few new components to the object type model and view model. These components are easily implementable in existing systems.

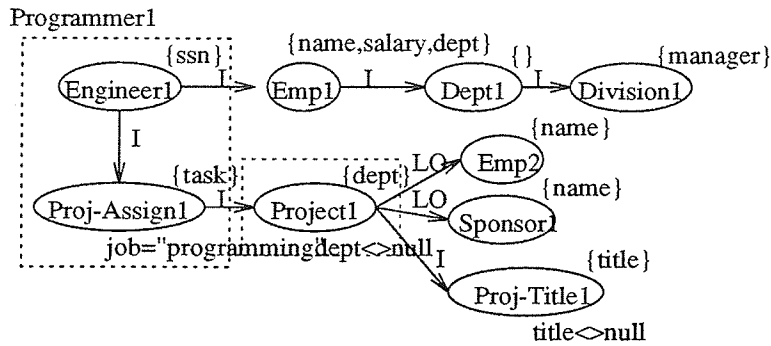


Figure 9: The query graph for the Programmer object with joins and non-null filters

Our result demonstrates how simple the mechanism becomes once the system model is established. The only criterion for the mechanism to use is the non-null option on object attributes, whose semantics is mapped to the non-null constraint on the query result. The number of left outer joins and non-null filters is reduced whenever possible using the integrity constraints so that the query is processed more efficiently.

Acknowledgement

We had fruitful discussions with Tore Risch, Peter Rathmann, Thierry Barsalou, and Dallan Quass in the early stage of this work. Linda DeMichiel, Peter Rathmann, Arthur Keller, and Witold Litwin gave valuable comments on the revised draft of this paper. This work was performed as part of the KBMS project, supported by DARPA Contract No. N039-84-C-02111. Ongoing research is partially sponsored by the Center for Integrated Facility Engineering at Stanford University.

References

- [1] Bancilhon, F., et al., "The Design and Implementation of O₂, an Object-Oriented Database System," in 'Advances in Object-Oriented Database Systems,' Springer-Verlag, September 1988.
- [2] Agrawal, R. and Gehani, N., "ODE (Object Database and Environment): The Language and the Data Model," Proc. the ACM SIGMOD International Conference on Management of Data, Portland, Oregon, May-June 1989.
- [3] Paepcke, A., "PCLOS: A Flexible Implementation of CLOS Persistence," Proc. the European Conference on Object-Oriented Programming, Oslo, Norway, August 1988.
- [4] Maier, D. and Stein, J., "Development of an Object-Oriented DBMS," Proc. International Conference on Object-Oriented Programming Systems, Languages, and Applications, September 1986, pp. 472 - 482.
- [5] Ford, S., et al., "Zeitgeist: Database Support for Object-Oriented Programming," Proc. International Workshop on Object-Oriented Database Systems, 1988, pp. 23 - 42.

- [6] Kim, W., Chou, N. and Garza, J., "Integrating an Object-Oriented Programming System with a Database System," Proc. OOPSLA International Conference on Object-Oriented Programming Systems, Languages, and Applications, September 1988, pp. 142 - 152.
- [7] Fishman, D., et al., "Iris: An Object-Oriented Database Management System," ACM Trans. on Office Information Systems, Vol. 5, No. 1, January 1987, pp. 48 - 69.
- [8] Stonebraker, M. and Rowe, L., "The Design of POSTGRES," Proc. ACM SIGMOD International Conference on Management of Data, 1986, pp. 340 - 354.
- [9] Learmont, T., and Cattell, R., "An Object-Oriented Interface to a Relational Database," Proc. the International Workshop on Object-Oriented Database Systems, 1987.
- [10] Wiederhold, G., "Views, Objects, and Databases", IEEE Computer, December 1986, pp. 37-44.
- [11] Barsalou, T. and Wiederhold, G., "Knowledge-based Mapping of Relations into Objects," The International Journal of Artificial Intelligence in Engineering, Computational Mechanics Publ., UK, 1989.
- [12] Wiederhold, G., Barsalou, T. and Chaudhuri, S., "Managing Objects in a Relational Framework," Stanford Technical report CS-89-1245, Stanford University, January 1989.
- [13] Rubenstein, W., Kubicar, M. and Cattell, R., "Benchmarking Simple Database Operation," Proc. ACM SIGMOD International Conference on Management of Data, May 1987.
- [14] Codd, E., "A relational model of data for large shared data banks," in 'Readings in Database Systems' edited by Michael Stonebraker, pp. 5 - 15, Morgan Kaufmann Publishers, Inc., 1988.
- [15] Ullman, J., "Principles of Database and Knowledge-Base Systems," Computer Science Press, 1988.
- [16] Morris, K., Ullman, J. and Van Gelder, A., "Design Overview of the NAIL! System," Proc. International Logic Programming Conference, 1986.
- [17] Tsur, S., and Zaniolo, C., "LDL: A Logic-based Data-Language," Proc. 12th VLDB conference, Kyoto, August 1986.
- [18] Chimenti, D., O'Hare, A., Krishnamurthy, R. and Zaniolo, C., "An Overview of the LDL System," IEEE Data Engineering, Vol. 10, No. 4, December 1987.
- [19] Date, C., "The Outer Join," Proc. Second International Conference on Databases, Cambridge, Britain, September 1983.
- [20] Date, C., "An Introduction to Database Systems," Vol. 1, Fourth edition, Addison-Wesley Publishing Company, Inc., 1986.
- [21] Bancilhon, F., "Object-Oriented Database Systems," Invited lecture, 7th ACM SIGART-SIGMOD-SIGACT Symposium on Principles of Database Systems., Austin, Texas, March 1988.
- [22] Maier, D., "Why Isn't There an Object-Oriented Data Model?," Proc. IFIP 11th World Computer Congress, San Francisco, California, September 1989.

- [23] Joseph, J., Thatte, S., Thompson, C. and Wells, D., "Report on the Object-Oriented Database Workshop," SIGMOD Record, Vol. 18, No. 3, September 1989.
- [24] Tsichritzis, D. and Lochovsky, F., "Data Models," Prentice-Hall, Inc. 1982, pp. 210 - 225.
- [25] Wilkes, W., Klahold, P., and Schlageter, G., "Complex and Composite Objects in CAD/CAM Databases" Proc. 5th IEEE International Conference on Data Engineering, Los Angeles, February 1989.
- [26] Buneman, P., Davidson, S. and Watters, A. "A Semantics for Complex Objects and Approximate Queries" ACM Symposium on Principles of Database Systems, 1988.
- [27] Kim, W., Banerjee, J., and Chou, H., "Composite Object Support in an Object-Oriented Database System," Proc. OOPSLA International Conference on Object-Oriented Programming Systems, Languages, and Applications, October 1987, pp. 118 - 125.
- [28] "IntelliCorp KEETM Software Development System User's Manual," Document No. 3.0-U-1, Intellicorp, July 1986.
- [29] Kempf, R. and Stelzner, M., "Teaching Object-Oriented Programming with the KEE System," Proc. International Conference on Object-Oriented Programming Systems, Languages, and Applications, October 1987, pp. 11 - 25.
- [30] Khoshafian, S. and Copeland, G., "Object Identity," Proc. International Conference on Object-Oriented Programming Systems, Languages, and Applications, 1986.
- [31] Abiteboul, S. and Kanellakis, P., "Object Identity as a Query Language Primitive," Proc. ACM SIGMOD International Conference on Management of Data, Portland, Oregon, May-June 1989.

