Optimizing Object Queries Using an Effective Calculus

LEONIDAS FEGARAS
The University of Texas at Arlington
and
DAVID MAIER
Oregon Graduate Institute of Science & Technology

Object-oriented databases (OODBs) provide powerful data abstractions and modeling facilities, but they generally lack a suitable framework for query processing and optimization. The development of an effective query optimizer is one of the key factors for OODB systems to successfully compete with relational systems, as well as to meet the performance requirements of many nontraditional applications. We propose an effective framework with a solid theoretical basis for optimizing OODB query languages. Our calculus, called the monoid comprehension calculus, captures most features of ODMG OQL, and is a good basis for expressing various optimization algorithms concisely. This article concentrates on query unnesting (also known as query decorrelation), an optimization that, even though it improves performance considerably, is not treated properly (if at all) by most OODB systems. Our framework generalizes many unnesting techniques proposed recently in the literature, and is capable of removing any form of query nesting using a very simple and efficient algorithm. The simplicity of our method is due to the use of the monoid comprehension calculus as an intermediate form for OODB queries. The monoid comprehension calculus treats operations over multiple collection types, aggregates, and quantifiers in a similar way, resulting in a uniform method of unnesting queries, regardless of their type of nesting.

Categories and Subject Descriptors: H.2.1 [**Database Management**]: Logical Design

General Terms: Design, Experimentation, Performance

Additional Key Words and Phrases: Object-oriented databases, nested relations, query decorrelation, query optimization

This work is supported in part by the National Science Foundation under grants IIS-9811525 and IRI-9619977, and by the Texas Higher Education Advanced Research Program under grant 003656-0043-1999.

Authors' addresses: L. Fegaras, Department of Computer Science and Engineering, The University of Texas at Arlington, 416 Yates Street, P.O. Box 19015, Arlington, TX 76019-19015; email: fegaras@cse.uta.edu; D. Maier, Department of Computer Science and Engineering, Oregon Graduate Institute of Science & Technology, 20000 N.W. Walker Road, Beaverton, OR 97006; email: maier@cse.ogi.edu.

Permission to make digital/hard copy of part or all of this work for personal or classroom use is granted without fee provided that the copies are not made or distributed for profit or commercial advantage, the copyright notice, the title of the publication, and its date appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists, requires prior specific permission and/or a fee.
© 2001 ACM 0362-5915/00/1200–0457 $5.00

1. INTRODUCTION

One of the reasons for the commercial success of relational database systems is that they perform well on many business applications, mostly due to sophisticated query processing and optimization techniques. However, many nontraditional database applications require more advanced data structures and more expressive operations than those provided by the relational model. Currently, there are two major approaches for addressing the needs of these applications [Carey and DeWitt 1996]. One is object-relational databases (ORDBs), which extend the relational model with object-oriented features, and the other is object-oriented databases (OODBs), which add database capabilities to object-oriented programming languages.

OODBs provide powerful data abstractions and modeling facilities, but they usually lack a suitable framework for query processing and optimization. In addition, many early OODB systems lacked a declarative language for associative access to data. Instead, they used simple pointer chasing to perform object traversals, which does not offer many opportunities for optimization. There is an increasing number of commercial OODB systems today, such as Poet, GemStone, ObjectStore, Objectivity, Jasmine, and Versant, that support a higher-lever query language interface (often along the lines of SQL). In addition, there is now an object data standard, called ODMG 3.0 [Cattell 2000], which supports a declarative higher-level object query language called OQL.

The development of an effective query optimizer is one of the key factors for OODB systems to successfully compete with relational systems, as well as to meet the performance requirements of many nontraditional applications. There are many aspects of the OODB query optimization problem that can benefit from the already proven relational query-optimization technology. However many key features of OODB languages present new and difficult problems not adequately addressed by this technology, including object identity, methods, encapsulation, subtype hierarchy, user-defined type constructors, large multimedia objects, multiple collection types, arbitrary nesting of collections, and nesting of query expressions.

We propose an effective framework with a solid theoretical basis for optimizing OODB query languages. We have chosen to concentrate on the OQL language of the ODMG 3.0 standard, which closely resembles the query language of the O2 OODBMS [Deux 1990]. We chose OQL because even though it is a small language, and hence easier to comprehend, it contains most of the language features that are showing up in other object query languages and proposed relational extensions such as SQL3 [Beech 1993], now called SQL:1999 [Eisenberg and Melton 1999]. If we are able to handle OQL in full generality, we believe our work will be widely applicable to other query languages.

OQL queries in our framework are translated into a calculus format that serves as an intermediate form, and then are translated into a version of the nested relational algebra. We use both a calculus and an algebra as

intermediate forms because the calculus closely resembles current OODB languages and is easy to put into canonical form, while the algebra is lower-level and can be directly translated into the execution algorithms supported by database systems.

Our calculus is called the *monoid comprehension calculus* [Fegaras and Maier 1995]. As we will demonstrate in this article, our calculus captures most features of OQL and is a good basis for expressing various optimization algorithms concisely. It is based on monoids, a general template for a data type, which can capture most collection and aggregate operators currently in use for relational and object-oriented databases. Monoid comprehensions—similar to set former notation, but applicable to types other than sets—give us a uniform way to express queries that simultaneously deal with more than one collection type and also naturally compose in a way that mirrors the allowable query nesting in OQL. Further, comprehension expressions permit easy integration of functional subexpressions. We demonstrate the expressive power of our calculus by showing how to map the major features of OQL into the monoid calculus.

We show in this article that monoid calculus is amenable to efficient evaluation. We give initial evidence by exhibiting a simple normalization system for putting expressions into a canonical form. The main contribution of this article is the description of a framework for optimizing monoid comprehension. We focus on a very important optimization problem, query unnesting (sometimes called query decorrelation), and propose a practical, effective, and general solution [Fegaras 1998b]. Our method generalizes many unnesting techniques proposed recently in the literature. Our framework is capable of removing *any* form of query nesting in our calculus using a very simple and efficient algorithm. The simplicity of our method is due to the monoid comprehension calculus as an intermediate form for OODB queries. Monoid comprehension calculus treats operations over multiple collection types, aggregates, and quantifiers in a similar way, resulting in a uniform method to unnest queries, regardless of their type of nesting. Our unnesting algorithm is compositional, that is, the translation of an embedded query does not depend on the context in which it is embedded, Instead, each subquery is translated independently, and all translations are composed to form the final unnested query. This property makes the soundness and completeness of the algorithm easy to prove. Our unnesting algorithm is complete, efficient, and easy to implement. It is the first decorrelation algorithm developed for a nontrivial database language *proven* to be complete. In addition, we believe, and we plan to show in future work, that it can be adapted to handle object-relational and relational queries.

This article is organized as follows. Section 2 describes the monoid comprehension calculus. Section 3 gives the translation of some OQL query forms into monoid comprehension. Section 4 describes a simple normalization algorithm that unnests most simple forms of nesting that do not require outer-joins, outer-unnests, or grouping. Section 5 proposes extensions to the basic framework for alleviating restrictions on idempotent operations and for capturing advanced list operations, vectors, and arrays.

Section 6 describes a version of the nested-relational algebra that supports aggregation, quantification, and the handling of null values (using outer-joins and outer-unnests). The semantics of these operations is given in terms of the monoid calculus. Section 7 presents transformation rules for unnesting OODB queries. Section 8 reports on the implementation of an ODMG-based system that is dependent on our theoretical framework, and Section 9 evaluates the performance of our implementation using sample queries. Finally, Section 10 compares our work with related efforts, and Section 11 describes our current and future plans related to this project.

2. THE MONOID COMPREHENSION CALCULUS

Several recent proposals for object-oriented database languages, including OQL, support multiple collection types, such as sets, bags, lists, and arrays. These approaches define a language syntax, but frequently fail to provide a concrete semantics. For example, is the join of a list with a set meaningful? If so, what is the resulting type of this join? More generally, what are the precise semantics of queries over multiple collection types? To answer these questions we need to form a theory that generalizes all collection types and their operations in a natural way. This theory must capture the differences between the collection types in such a way that no inconsistencies are introduced, and, more importantly, must abstract their similarities. By abstracting common semantic features, we derive a framework that treats collections uniformly in a simple and extensible language. The primary focus of such a framework is the bulk manipulation of collection types. Bulk operations are both the source of expressiveness and the basis for efficient execution.

Consider lists and sets. What are the semantic properties that make lists different from sets? Intuitively, one may exchange two elements of a set or insert a set element twice into a set without changing the set. These properties do not hold for lists. To formalize these observations, we indicate how sets and lists are constructed and then impose these properties on the set and list constructors. One way of constructing sets is to union together a number of singleton set elements, e.g., $\{1\} \cup \{2\} \cup \{3\}$ constructs the set $\{1, 2, 3\}$. Similarly, one way of constructing lists is to append singleton list elements, e.g., $[1] +\!\!+ [2] +\!\!+ [3]$ constructs the list $[1, 2, 3]$ (where $+\!\!+$ is the list append function). Both $\cup$ and $+\!\!+$ are associative operations, but only $\cup$ is commutative and idempotent (i.e., $\forall x : x \cup x = x$). It is the commutativity and idempotence properties of $\cup$ that make sets different from lists. As to their similarities, both the set and list constructors have an identity. The empty set $\{\}$ is the identity of $\cup$, while the empty list $[]$ is the identity of $+\!\!+$. Using terminology from abstract algebra, we say that both $(\cup, \{\})$ and $(+\!\!+, [])$ are *monoids*, and, more specifically, $(\cup, \{\})$ is a commutative and idempotent monoid.

Primitive types, such as integers and booleans, can be represented as monoids too, but possibly with a choice of monoid. For example, both

$(+, 0)$ and $(*, 1)$ are integer monoids and both $(\vee, \text{false})$ and $(\wedge, \text{true})$ are boolean monoids. As will become apparent when monoid operations are described, the choice of a monoid depends on a particular interpretation of a primitive type in a query. We call the monoids for primitive types *primitive monoids*.

In contrast to primitive types, collection types, such as sets, bags, and lists, are parametric types that capture homogeneous collections of objects. The types set(⟨name : string, age : int⟩) and set(set(int)) are examples of collection types. We call the monoids for collection types *collection monoids*. Each collection monoid must have a unit function that takes an element of some type as input and constructs a singleton value of the collection type. For example, the list unit function takes an element a and constructs the singleton list $[a]$. Any list can be constructed from the three monoid primitives for lists: the empty list [], the list unit function, and the list append $++$.

Since all types are represented as monoids, a query in our framework is a map from some monoids to a particular monoid. These maps are called *monoid homomorphisms*. For example, a monoid homomorphism from lists to sets in our framework is captured by an operation of the form

$$\text{hom}[++, \cup](f)A,$$

where A is a list and f is a function that takes an element x of A and returns a set $f(x)$. Basically, this monoid homomorphism performs the following computation:

```
result := {};
foreach x in A do
    result := result ∪ f(x);
return result;
```

In other words, $\text{hom}[++, \cup](f)A$ replaces [] in A by {}; $++$ by $\cup$; and the singleton list $[x]$ by $f(x)$. Therefore, if A is the list $[a_1, \ldots, a_n]$, which is generated by $[a_1] ++ \cdots ++ [a_n]$, then the result is the set $f(a_1) \cup \cdots \cup f(a_n)$.

A monoid homomorphism captures a divide-and-conquer computation, since any list $A = [a_1, \ldots, a_n]$ can be divided into two lists A_1 and A_2 such that $A = A_1 ++ A_2$. In that case, the operation $\text{hom}[++, \cup](f)A$ can be decomposed into $(\text{hom}[++, \cup](f)A_1) \cup (\text{hom}[++, \cup](f)A_2)$.

Unfortunately, not all monoid homomorphisms are well-formed. For example, due to nondeterminism, sets cannot be converted unambiguously into lists. This semantic restriction is purely syntactic in our framework. That is, it depends on the static properties of the monoids involved in a homomorphism, such as the commutativity and the idempotence properties. Thus, well-formed queries will not result in any ill-defined homomorphisms.

The monoid homomorphism is the only form of bulk manipulation of collection types supported in our framework. But as we will demonstrate, monoid homomorphisms are very expressive. In fact, a small subset of these functions, namely the monoid homomorphisms from sets to sets, captures the nested relational algebra (since these homomorphisms are equivalent to the extension operator $ext(f)$ for sets, shown to be at least as expressive as the nested relational algebra [Buneman et al. 1995; Breazu-Tannen et al. 1992b]). But monoid homomorphisms go beyond the algebra to capture operations over multiple collection types, such as the join of a list with a bag returning a set, as well as predicates and aggregates. For example, an existential predicate over a set is a monoid homomorphism from the set monoid to the boolean monoid ($\vee$, false), while an aggregation, such as summing all elements of a list, is a monoid homomorphism from the list monoid to the numeric monoid $(+, 0)$.

We also define a calculus for this algebra, called the *monoid comprehensions calculus*, which captures operations involving multiple collection types in declarative form. Monoid comprehension is defined in terms of monoid homomorphisms, but any monoid homomorphism can be expressed by some monoid comprehension. Programs expressed in our calculus are far easier to understand and manipulate than the equivalent algebraic forms. In a way, monoid comprehension resembles the tuple relational calculus, but here query variables may range over multiple collection types, while the output of the comprehension may be of yet another collection type.

For example, the following monoid comprehension:

$$\cup\{\!\{ (a, b) \;[\!]\; a \leftarrow [1, 2, 3], b \leftarrow \{\!\{4, 5\}\!\} \}\!\}$$

joins the list $[1, 2, 3]$ with the bag $\{\!\{4, 5\}\!\}$ and returns the following set (it is a set because the comprehension is tagged by $\cup$):

$$\{(1, 4), (1, 5), (2, 4), (2, 5), (3, 4), (3, 5)\}.$$

Another example is

$$+\{\!\{ a \;[\!]\; a \leftarrow [1, 2, 3], a \geq 2 \}\!\},$$

which uses the merge function $+$ of the monoid $(+, 0)$ to add integers. This expression returns 5, the sum of all list elements greater than or equal to 2.

The rest of this section gives a formal definition of monoids and monoid operations.

2.1 Monoids

Definition 1 (Monoid). A monoid of type T is a pair $(\oplus, \mathcal{Z}_\oplus)$, where $\oplus$ is an associative function of type $T \times T \rightarrow T$ and $\mathcal{Z}_\oplus$ is the left and right identity of $\oplus$.

The type $T \times T \rightarrow T$ in the above definition denotes the type of binary functions that take two values of type T as input and return a value of type T. Function $\oplus$ is called the *merge* function and value $\mathcal{Z}_\oplus$ is called the *zero* element of the monoid $(\oplus, \mathcal{Z}_\oplus)$. According to the definition, the zero

Table I. Examples of Monoids

type $T_\oplus$	$\oplus$	$\mathcal{Z}_\oplus$	C/I
int	$+$	0	C
int	$\times$	1	C
int	max	0	CI
bool	$\vee$	false	CI
bool	$\wedge$	true	CI

A. Primitive Monoids

type $T_\oplus$	$\oplus$	$\mathcal{Z}_\oplus$	$\mathcal{U}_\oplus(a)$	C/I
set (α)	$\cup$	$\{\}$	$\{a\}$	CI
bag (α)	$\uplus$	$\{\{\}\}$	$\{\{a\}\}$	C
list (α)	$++$	$[]$	$[a]$	

B. Collection Monoids

element satisfies $\mathcal{Z}_\oplus \oplus x = x \oplus \mathcal{Z}_\oplus = x$ for every x. A monoid $(\oplus, \mathcal{Z}_\oplus)$ may be a *commutative* monoid (i.e., when $\oplus$ is commutative) or an *idempotent* monoid (i.e., when $\forall x : x \oplus x = x$), or both. In our framework, if a monoid is not idempotent, it must be *anti-idempotent*: $\forall x \neq \mathcal{Z}_\oplus : x \oplus x \neq x$, and if it is not commutative, it must be anti-commutative: $\forall x \neq y, x \neq \mathcal{Z}_\oplus, y \neq \mathcal{Z}_\oplus : x \oplus y \neq y \oplus x$, for reasons discussed in Section 2.2. For example, $(+, 0)$ is a commutative and anti-idempotent monoid, while $(\cup, \{\})$ is a commutative and idempotent monoid.

Since the merge function uniquely identifies a monoid, we often use the merge function name as the monoid name. In addition, we use the notation $T_\oplus$ to represent the type of the monoid with merge function $\oplus$.

Table I presents some examples of monoids. The C/I column indicates whether the monoid is a commutative or idempotent monoid (if a monoid is not tagged by I, it is anti-idempotent, and if it is not tagged by C, it is anti-commutative). An example of an idempotent and anti-commutative monoid that captures vectors and arrays is presented in Section 5.3.

The monoids in Table I.A are called *primitive monoids* because they construct values of a primitive type. The monoids $\cup$, $\uplus$, and $++$ in Table I.B capture the well-known collection types for sets, bags, and linear lists (where $\uplus$ is the additive union for bags). They are called *collection monoids*. Each collection monoid $(\oplus, \mathcal{Z}_\oplus)$ needs the additional definition of a *unit function*, $\mathcal{U}_\oplus$, which, along with merge and zero, allows the construction of all possible values of the monoid type $T_\oplus$. For example, the unit function for the set monoid is $\lambda x.\{x\}$, where expression $\lambda x.e$ denotes the function f with $f(x) = e$. That is, the set unit function takes a value x as input and constructs the singleton set $\{x\}$ as output. Consequently, a collection monoid is associated with a triple $(\oplus, \mathcal{Z}_\oplus, \mathcal{U}_\oplus)$ and corresponds to a parametric type $T_\oplus(\alpha)$, where α is a type parameter that can be bound to any type with equality, including another collection type.[1]

[1]Using terminology from abstract algebra, $(\oplus, \mathcal{Z}_\oplus, \mathcal{U}_\oplus)$ is a free monoid and $\mathcal{U}_\oplus$ is a natural transformation from the identity to the $T_\oplus$ functor that satisfies the universal mapping property (also known as uniqueness property, or adjunction) that for any $f : \alpha \to T_\oplus(\alpha)$ there is a unique function hom$[\oplus, \otimes](f)$ (given in Definition 2) such that hom$[\oplus, \otimes](f) \circ \mathcal{U}_\oplus = f$. Function hom$[\oplus, \otimes](f)$ is called the homomorphic extension of f. The interested reader is referred to the book by Pierce [1991] for a more formal treatment of monoids and monoid homomorphisms.

We use the shorthand $\oplus\{e_1, \ldots, e_n\}$ to represent the construction $\mathcal{U}_\oplus(e_1) \oplus \cdots \oplus \mathcal{U}_\oplus(e_n)$ over the collection monoid $\oplus$. In particular, we use the shorthand $\{e_1, \ldots, e_n\}$ for sets, $\{\{e_1, \ldots, e_n\}\}$ for bags and $[e_1, \ldots, e_n]$ for lists.

We define the mapping ψ from monoids to the powerset of $\{C, I\}$ as: $C \in \psi(\oplus)$ if and only if $\oplus$ is commutative and $I \in \psi(\oplus)$ if and only if $\oplus$ is idempotent. The partial order $\preceq$ between monoids is defined as

$$\otimes \preceq \oplus \equiv \psi(\otimes) \subseteq \psi(\oplus).$$

For example, $+\!\!+ \preceq \uplus \preceq \cup$, since $\cup$ is commutative and idempotent, $\uplus$ is commutative but not idempotent, and $+\!\!+$ is neither commutative nor idempotent.

2.2 Monoid Homomorphisms

We now define our algebraic operator, which is parameterized by input and output monoids. The main purpose of this operator is to give formal semantics to monoid comprehension, discussed in Section 2.3. Our monoid calculus is defined in terms of monoid comprehension exclusively.

Definition 2 (Monoid Homomorphism). A homomorphism $\mathrm{hom}[\oplus, \otimes](f)A$ from the collection monoid $(\oplus, \mathcal{Z}_\oplus, \mathcal{U}_\oplus)$ to any monoid $(\otimes, \mathcal{Z}_\otimes)$, where $\oplus \preceq \otimes$ is defined by the following inductive equations:

$$\begin{aligned}
\mathrm{hom}[\oplus, \otimes](f)(\mathcal{Z}_\oplus) &= \mathcal{Z}_\oplus \\
\mathrm{hom}[\oplus, \otimes](f)(\mathcal{U}_\oplus(a)) &= f(a) \\
\mathrm{hom}[\oplus, \otimes](f)(x \oplus y) &= (\mathrm{hom}[\oplus, \otimes](f)(x)) \otimes (\mathrm{hom}[\oplus, \otimes](f)(y)).
\end{aligned}$$

That is, A is a collection of type $T_\oplus(T')$ for some type T' and f is a function from T' to $T_\otimes$. Basically, the expression $\mathrm{hom}[\oplus, \otimes](f)A$ replaces $\mathcal{Z}_\oplus$ in A by $\mathcal{Z}_\otimes$, $\oplus$ by $\otimes$, and $\mathcal{U}_\oplus$ by f. For example, for $H = \mathrm{hom}[+\!\!+, \cup](f)$, Definition 2 is equivalent to

$$\begin{aligned}
H([]) &= \{\} \\
H([a]) &= f(a) \\
H(x +\!\!+ y) &= H(x) \cup H(y).
\end{aligned}$$

If A is the list $[a_1, \ldots, a_n]$, then $H(A)$ computes the set $f(a_1) \cup \cdots \cup f(a_n)$. For instance, if $f(x) = \{x + 1\}$, then $H(A) = \{a_1 + 1, \ldots, a_n + 1\}$.

The condition $\oplus \preceq \otimes$ in Definition 2 is important. If the collection monoid $\oplus$ is a commutative or idempotent monoid, then $\otimes$ must be too. For example, the bag cardinality function can be expressed as $\mathrm{hom}[\uplus, +](\lambda x.1)A$, which is well-formed, while the similar function for sets $\mathrm{hom}[\cup, +](\lambda x.1)A$ is not (since $+$ is commutative but not idempotent).

Without this restriction we would have [Breazu-Tannen and Subrahmanyam 1991]:

$$\begin{aligned} 1 &= \text{hom}[\cup, +](\lambda x.1)(\{a\}) \\ &= \text{hom}[\cup, +](\lambda x.1)(\{a\} \cup \{a\}) \\ &= (\text{hom}[\cup, +](\lambda x.1)(\{a\})) + (\text{hom}[\cup, +](\lambda x.1)(\{a\})) \\ &= 1 + 1 = 2. \end{aligned}$$

This restriction also prohibits the conversion of sets into lists (since $\cup \not\preceq +\!\!+$). Furthermore, Definition 2 justifies the restriction that noncommutative monoids should be anti-commutative and nonidempotent monoids should be anti-idempotent. If we allow a nonidempotent monoid $\oplus$ to satisfy $x \oplus x = x$ for at least one $x = x_0$ (but not for all x, since $\oplus$ is not idempotent), then we have $\text{hom}[\oplus, \otimes](f)(x_0) = \text{hom}[\oplus, \otimes](f)(x_0 \oplus x_0) = (\text{hom}[\oplus, \otimes](f)(x_0)) \otimes (\text{hom}[\oplus, \otimes](f)(x_0))$, which is not always true for an arbitrary nonidempotent monoid $\otimes$. The observations above are generalized in the following theorem.

THEOREM 1. *A well-formed monoid homomorphism* $\text{hom}[\oplus, \otimes](f)$ *preserves the properties of the monoid* $\otimes$.

PROOF. Let $H = \text{hom}[\oplus, \otimes](f)$, then $\oplus \preceq \otimes$, since H is well-formed. We prove that, for any valid instance x of the monoid $\oplus$, the value $H(x)$ is a valid instance of $\otimes$. We consider idempotence only; commutativity can be handled in a similar way. If $\oplus$ is idempotent, then $H(x \oplus x) = H(x)$. Thus, $H(x) \otimes H(x) = H(x)$, which means that $\otimes$ is not anti-idempotent. This is true since $\oplus \preceq \otimes$, which implies that $\otimes$ is idempotent. If $\oplus$ is anti-idempotent, then $\otimes$ can be either idempotent or anti-idempotent. If $\otimes$ is anti-idempotent, then for $H(x) \neq \mathcal{Z}_{\otimes} : H(x) \otimes H(x) \neq H(x) \Rightarrow H(x \oplus x) \neq H(x)$, which implies that $x \oplus x \neq x$. If $\otimes$ is idempotent, then $H(x) \otimes H(x) = H(x) \Rightarrow H(x \oplus x) = H(x)$, which does not necessarily conflict with $x \oplus x \neq x$. □

The following are examples of well-formed monoid homomorphisms:

$$\begin{aligned} \text{map}(f)x &= \text{hom}[\cup, \cup](\lambda a.\{f(a)\})x \\ x \times y &= \text{hom}[\cup, \cup](\lambda a.\text{hom}[\cup, \cup](\lambda b.\{(a, b)\})y)x \\ e \in x &= \text{hom}[\cup, \vee](\lambda a.(a = e))x \\ \text{filter}(p)x &= \text{hom}[\cup, \cup](\lambda a.\textbf{if}\ p(a)\ \textbf{then}\ \{a\}\ \textbf{else}\ \{\})x \\ \text{length}(x) &= \text{hom}[+\!\!+, +](\lambda a.1)x, \end{aligned}$$

where $\text{map}(f)x$ maps the function f over all elements of the set x, $x \times y$ computes the Cartesian product of the sets x and y, and $\text{filter}(p)x$ selects all elements a of the set x that satisfy the predicate $p(a)$.

2.3 Monoid Comprehension

Queries in our calculus are expressed in terms of monoid comprehension. Informally, a monoid comprehension over the monoid $\oplus$ takes the form $\oplus\{\!\{ e \,[]\, \bar{q} \}\!\}$. The merge function $\oplus$ is called the *accumulator* of the comprehension and the expression e is called the *head* of the comprehension. Each term q_i in the term sequence $\bar{q} = q_1, \ldots, q_n$, $n \geq 0$ is called a *qualifier*, and is either

- a *generator* of the form $v \leftarrow e'$, where v is a *range variable* and e' is an expression (the generator domain) that constructs a collection, or
- a *filter* pred, where pred is a predicate.

Formally, monoid comprehension is defined in terms of monoid homomorphisms.

Definition 3 (Monoid Comprehension). Monoid comprehension over a primitive or collection monoid $\oplus$ is defined by the following inductive equations:

$$\oplus\{\!\{ e \,[]\, \}\!\} = \begin{cases} \mathcal{U}_\oplus(e) & \text{for a collection monoid } \oplus \\ e & \text{for a primitive monoid } \oplus \end{cases} \quad \text{(D1)}$$

$$\oplus\{\!\{ e \,[]\, x \leftarrow u, \bar{q} \}\!\} = \text{hom}[\otimes, \oplus](\lambda x.\, \oplus\{\!\{ e \,[]\, \bar{q} \}\!\})u \quad \text{(D2)}$$

$$\oplus\{\!\{ e \,[]\, \text{pred}, \bar{q} \}\!\} = \textbf{if } \text{pred } \textbf{then } \oplus\{\!\{ e \,[]\, \bar{q} \}\!\} \textbf{ else } \mathcal{Z}_\oplus \quad \text{(D3)}$$

where $\otimes$ is the collection monoid associated with the expression u, and is possibly different from $\oplus$.

While the monoid $\oplus$ of the output is specified explicitly, the collection monoid $\otimes$ associated with the expression u in $x \leftarrow u$ can be inferred (by the typing rules given in Section 2.5). The main purpose of the definition above is to give semantics to monoid comprehensions, not to suggest in any way how they are implemented. Monoid comprehensions can be translated effectively into efficient evaluation algorithms, as we show in Section 7.

Note that $+\!\!+\{\!\{ x \,[]\, x \leftarrow \{1, 2\} \}\!\}$ is not a well-formed comprehension, since it is translated into a homomorphism from $\cup$ to $+\!\!+$ (it maps sets into lists), while $+\{\!\{ x \,[]\, x \leftarrow \{\!\{1, 2\}\!\} \}\!\}$ is well-formed and computes $1 + 2 = 3$. That is, if one of the generators in a monoid comprehension is over a commutative or idempotent monoid, then the comprehension accumulator must be a commutative or idempotent monoid, respectively. This condition can be checked statically, since the commutativity and idempotence properties of a monoid are specified explicitly when this monoid is defined (see Section 2.5).

Relational joins can be represented directly as comprehensions. The join of two sets x and y is

$$\cup\{\!\{ \text{f(a, b)} \,[]\, \text{a} \leftarrow \text{x}, \text{b} \leftarrow \text{y}, \text{p(a, b)} \}\!\},$$

where p is the join predicate and function f constructs an output set element given two elements from x and y. For example, if p(a,b) = (a.C=b.C) $\wedge$ (a.D > 10) and f(a,b) = $\langle$ C=a.C, D=b.D $\rangle$ (i.e., a record construction), then this comprehension becomes

$$\cup\{\langle C = a.C,\ D = b.D\rangle \,[\!]\, a \leftarrow x,\ b \leftarrow y,\ a.C = b.C,\ a.D > 10\}.$$

If we use the rules in Definition 3, this comprehension is translated into algebraic form as

$$\begin{aligned}\text{hom}[\cup, \cup](\lambda a.\ \text{hom}][\cup, \cup](\lambda b.\ &\textbf{if } (a.C{=}b.C) \wedge (a.D > 10)\\ &\textbf{then}\{\langle\ C{=}a.C,\ D{=}b.D\ \rangle\}\\ &\textbf{else } \{\})\ y)\ x.\end{aligned}$$

This evaluation resembles a nested loop, but we show in Section 7 that comprehensions similar to the one above can be effectively translated into joins.

Furthermore, comprehensions can be used to join different collection types. For example,

$$\cup\{\,(x, y) \,[\!]\, x \leftarrow [1, 2],\ y \leftarrow \{\!\{3, 4, 3\}\!\}\,\} = \{(1, 3), (1, 4), (2, 3), (2, 4)\}.$$

Another example is nest$(k)x$ equal to

$$\cup\{\,\langle \text{KEY} = k(e),\ \text{INFO} = \cup\{\,a \,[\!]\, a \leftarrow x,\ k(e) = k(a)\,\}\rangle \,[\!]\, e \leftarrow x\,\},$$

which is similar to the nesting operator for nested relations, but it is a bit more general and gives output in a different form. Similarly, the inverse function, unnest, is

$$\text{unnest}(x) = \cup\{\,e \,[\!]\, s \leftarrow x,\ e \leftarrow s.\text{INFO}\,\}.$$

The last comprehension is an example of a *dependent join* in which the value of the collection s.INFO in the second generator depends on the value of s, which is an element of the first relation x. Dependent joins are a convenient way of traversing nested collections. Other examples of comprehensions are the following:

$$\begin{aligned}
\text{filter}(p)x &= \cup\{\,e \,[\!]\, e \leftarrow x,\ p(e)\,\} & \text{count}(x, a) &= +\{\,1 \,[\!]\, e \leftarrow x,\ e = a\,\}\\
\text{flatten}(x) &= \cup\{\,e \,[\!]\, s \leftarrow x,\ e \leftarrow s\,\} & \text{maximum}(x) &= max\{\,e \,[\!]\, e \leftarrow x\,\}\\
x \cap y &= \cup\{\,e \,[\!]\, e \leftarrow x,\ e \in y\,\} & \exists a \in x : e &= \vee\{\,e \,[\!]\, a \leftarrow x\,\}\\
\text{length}(x) &= +\{\,1 \,[\!]\, e \leftarrow x\,\} & \forall a \in x : e &= \wedge\{\,e \,[\!]\, a \leftarrow x\,\}\\
\text{sum}(x) &= +\{\,e \,[\!]\, e \leftarrow x\,\} & a \in x &= \vee\{\,a = e \,[\!]\, e \leftarrow x\,\}.
\end{aligned}$$

The expression sum(x) adds the elements of any non-idempotent monoid x, e.g., sum$([1, 2, 3]) = 6$. The expression count(x, a) counts the number of

occurrences of a in the bag x, e.g., count($\{\{1, 2, 1\}\}$, 1) = 2. Recall that the $(max, 0)$ monoid is both commutative and idempotent, and thus x in maximum(x) can be of any collection type.

THEOREM 2. *Monoid comprehensions and monoid homomorphisms have the same expressive power.*

PROOF. Definition 3 expresses comprehensions in terms of monoid homomorphisms. Thus, we need to prove that any monoid homomorphism can be expressed in terms of monoid comprehension. In particular, we prove the equations $\text{hom}[\oplus, \otimes](f)(A) = \otimes\{ y \;[]\; x \leftarrow A, y \leftarrow f(x) \}$ for a collection monoid $\otimes$ and $\text{hom}[\oplus, \otimes](f)(A) = \otimes\{ f(x) \;[]\; x \leftarrow A \}$ for a primitive monoid $\otimes$. The universal property for homomorphisms indicates that, for a collection monoid $\otimes$, $\forall x : \text{hom}[\otimes, \otimes](\mathcal{U}_\otimes)x = x$. According to Definition 3, for a collection monoid $\otimes$, $\otimes\{ y \;[]\; x \leftarrow A, y \leftarrow f(x) \} = \text{hom}[\oplus, \otimes](\lambda x.\text{hom}[\otimes, \otimes](\lambda y.\mathcal{U}_\otimes(y))(f(x)))(A) = \text{hom}[\oplus, \otimes](f)(A)$, which proves the first equation. For a primitive monoid $\otimes$, $\otimes\{ y [] \; x \leftarrow A, y \leftarrow f(x) \} = \text{hom}[\oplus, \otimes]$ $(\lambda x.\text{hom}[\otimes, \otimes](\lambda y.y)(f(x)))(A)$, which in turn is equal to $\text{hom}[\oplus, \otimes](f)(A)$, which proves the second equation. □

2.4 Monoid Types

In our treatment of queries we consider the following types to be valid:

Definition 4 (Monoid Type). A *monoid type* takes one of the following forms:

class_name	(a reference to a user-defined class)
T	(a primitive type, such as int or bool)
$T(t)$	(a collection type, such as list(int))
$\langle A_1 : t_1, \ldots, A_n : t_n \rangle$	(a record type)

where $t, t_1, \ldots, t_n$ are monoid types. Thus, monoid types can be nested freely.

Our type system is rich enough to capture most ODMG ODL data types. For our queries, we use the university schema shown in Figure 1 as a running example. Note that the class Instructor is a subclass of the class Person. ODL relationships can be 1:1, 1:N, or N:M. When used in OQL queries, a binary relationship between the classes A and B offers two ways of relating A and B objects: given an instance of A to retrieve the related instance(s) of B and given an instance of B to retrieve the related instance(s) of A. A class extent, such as Instructors for the class Instructor, is a set of all persistent instances of this class. Extents are the only entry points to the database. We assume persistence by reachability, which indicates that not only the members of an extent but also all the objects that can be reached from these members, are persistent. In our treatment of types, relationships are handled as attributes and extents as sets of class

```
class Person ( extent Persons key ssn )
{   attribute long ssn;
    attribute string name;
    attribute struct( street: string, zipcode: string ) address;
};
class Instructor extends Person ( extent Instructors )
{   attribute long salary;
    attribute string rank;
    attribute set⟨string⟩ degrees;
    relationship Department dept inverse Department::instructors;
    relationship set⟨Course⟩ teaches inverse Course::taught_by;
};
class Department ( extent Departments keys dno, name )
{   attribute long dno;
    attribute string name;
    attribute Instructor head;
    relationship set⟨Instructor⟩ instructors inverse Instructor::dept;
    relationship set⟨Course⟩ courses_offered inverse Course::offered_by;
};
class Course ( extent Courses keys code, name )
{   attribute string code;
    attribute string name;
    relationship Department offered_by inverse Department::courses_offered;
    relationship Instructor taught_by inverse Instructor::teaches;
    relationship set⟨Course⟩ is_prerequisite_for inverse Course::has_prerequisites;
    relationship set⟨Course⟩ has_prerequisites inverse Course::is_prerequisite_for;
};
```

Fig. 1. The ODL schema of the university database.

references such as Courses: set ⟨Course⟩. A class itself is treated as an aggregation (a record) of its attributes.

2.5 The Monoid Comprehension Calculus

Our monoid comprehension calculus consists of a number of syntactic forms that operate over monoid types. These syntactic forms can be composed to construct yet more complicated terms. Not every combination of syntactic forms is well-formed though. Ill-formed expressions can be detected and terms can be assigned unique types with the help of a typing system, which is usually specified by a set of typing rules. We first present the monoid calculus and then the typing rules of the calculus.

Definition 5 (Monoid Comprehension Calculus). The monoid calculus consists of the syntactic forms in Figure 2, where $e, e_1, \ldots, e_n$ are terms in the monoid calculus, v is a variable, t is a monoid type, and $q_1, \ldots, q_n$ are qualifiers of the form $v \leftarrow e$ or e.

For example, the following expression:

$$\cup\{ \text{e.name} \;[]\; \text{el} \leftarrow \uplus\{ \text{d.instructors} \;[]\; \text{d} \leftarrow \text{Departments}, \text{d.name}=\text{“CSE”} \},\\ \text{e} \leftarrow \text{el}, \vee\{ \text{c.name}=\text{“cse5331”} \;[]\; \text{c} \leftarrow \text{e.teaches} \} \},$$

which is based on our example OODB schema, is a valid expression in the calculus. It returns all CSE instructors who teach cse5331. For each

NULL	null value
c	constant
v	variable
$e.A$	record projection
$\langle A_1 = e_1, \ldots, A_n = e_n \rangle$	record construction
if e_1 **then** e_2 **else** e_3	if-then-else statement
$e_1 \,\text{op}\, e_2$	where op is a primitive binary function, such as $+, =, <, >$
$\lambda v\!:\!t.\, e$	function abstraction
$e_1(e_2)$	function application
$\mathcal{Z}_\oplus$	zero element
$\mathcal{U}_\oplus(e)$	singleton construction
$e_1 \oplus e_2$	merging
$\oplus\{\!\{ e \,[\!]\, q_1, \ldots, q_n \}\!\}$	comprehension

Fig. 2. The monoid comprehension calculus.

department with name CSE, the inner bag comprehension returns all instructors in the department. Thus, the result of the bag comprehension is of type bag(bag(Instructor)). Since variable el ranges over the result of the bag comprehension, the type of el is bag(Instructor). Variable e ranges over el, which indicates that the type of e is Instructor. The $\vee$-comprehension captures an existential quantification; it returns true if there is at least one course c in e.teaches that has c.name= "cse5331".

We use the shorthand $x \equiv u$ to represent binding the variable x to the value u. The meaning of this construct is given by the following reduction rule:

$$\oplus\{\!\{ e \,[\!]\, \bar{r}, x \equiv u, \bar{s} \}\!\} \longrightarrow \oplus\{\!\{ e[u/x] \,[\!]\, \bar{r}, \bar{s}[u/x] \}\!\}, \tag{N1}$$

where $e[u/x]$ is the expression e with u substituted for all the free occurrences of x. Another shorthand is the pair (x, y), which is equivalent to the record $\langle \text{fst} = x, \text{snd} = y \rangle$. In addition, as syntactic sugar, we allow irrefutable patterns in place of lambda variables, range variables, and variables in bindings. In functional programming languages [Peyton Jones 1987], a pattern is irrefutable if it is either a variable or a tuple consisting of irrefutable patterns. Patterns such as these can be compiled away using standard pattern decomposition techniques [Peyton Jones 1987]. For example, $\cup\{\!\{ x + y \,[\!]\, (x, (y, z)) \leftarrow A, z = 3 \}\!\}$ is equivalent to $\cup\{\!\{ a.\text{fst} + a.\text{snd.fst} \,[\!]\, a \leftarrow A, a.\text{snd.snd} = 3 \}\!\}$.

Figure 3 gives the typing rules of the terms in the monoid calculus. The name v in Figure 3 indicates a variable name, names starting with A are attribute names, names starting with e represent terms in our calculus, α is a type variable, and names starting with t represent types. The notation $\sigma \vdash e : t$ indicates that the term e is assigned the type t under the substitution σ. If a type equation is a fraction, the numerator is the premise, while the denominator is the consequence. For example, Eq. (T4) indicates that if e is a tuple of type $\langle A_1 : t_1, \ldots, A_n : t_n \rangle$, then the type of $e.A_i$ is the ith component of the tuple type t_i. The substitution list σ binds

$$\sigma \vdash c : \text{type of the constant } c \quad \text{(T1)}$$

$$\sigma \vdash \text{NULL} : \alpha \quad \text{(T2)}$$

$$\sigma \vdash v : \sigma(v) \quad \text{(T3)}$$

$$\frac{\sigma \vdash e : \langle A_1 : t_1, \ldots, A_n : t_n \rangle}{\sigma \vdash e.A_i : t_i} \quad \text{(T4)}$$

$$\frac{\sigma \vdash e_1 : t_1, \ldots, \sigma \vdash e_n : t_n}{\sigma \vdash \langle a_1 = e_1, \ldots, a_n = e_n \rangle : \langle a_1 : t_1, \ldots, a_n : t_n \rangle} \quad \text{(T5)}$$

$$\frac{\sigma \vdash e_1 : \text{bool},\ \sigma \vdash e_2 : t,\ \sigma \vdash e_3 : t}{\sigma \vdash \textbf{if } e_1 \textbf{ then } e_2 \textbf{ else } e_3 : t} \quad \text{(T6)}$$

$$\frac{\sigma[t_1/v] \vdash e : t_2}{\sigma \vdash \lambda v{:}t_1.\, e : t_1 \to t_2} \quad \text{(T7)}$$

$$\frac{\sigma \vdash e_1 : t_1 \to t_2,\ \sigma \vdash e_2 : t_1}{\sigma \vdash e_1(e_2) : t_2} \quad \text{(T8)}$$

$$\frac{\oplus \text{ is collection monoid}}{\sigma \vdash \mathcal{Z}_\oplus : T_\oplus(\alpha)} \quad \text{(T9)}$$

$$\frac{\sigma \vdash e : t}{\sigma \vdash \mathcal{U}_\oplus(e) : T_\oplus(t)} \quad \text{(T10)}$$

$$\frac{\oplus \text{ is primitive monoid} \quad \sigma \vdash e_1 : T_\oplus,\ \sigma \vdash e_2 : T_\oplus}{\sigma \vdash e_1 \oplus e_2 : T_\oplus} \quad \text{(T11)}$$

$$\frac{\oplus \text{ is collection monoid} \quad \sigma \vdash e_1 : T_\oplus(t),\ \sigma \vdash e_2 : T_\oplus(t)}{\sigma \vdash e_1 \oplus e_2 : T_\oplus(t)} \quad \text{(T12)}$$

$$\frac{\sigma \vdash e : T_\oplus,\ \oplus \text{ is primitive monoid}}{\sigma \vdash \oplus\{\, e \,[\!]\, \} : T_\oplus} \quad \text{(T13)}$$

$$\frac{\sigma \vdash e : t,\ \oplus \text{ is collection monoid}}{\sigma \vdash \oplus\{\, e \,[\!]\, \} : T_\oplus(t)} \quad \text{(T14)}$$

$$\frac{\sigma \vdash e_2 : T_\otimes(t_2),\ \otimes \preceq \oplus, \quad \sigma[t_2/v] \vdash \oplus\{\, e_1 \,[\!]\, \bar{r} \,\} : t_1}{\sigma \vdash \oplus\{\, e_1 \,[\!]\, v \leftarrow e_2, \bar{r} \,\} : t_1} \quad \text{(T15)}$$

$$\frac{\sigma \vdash e_2 : \text{bool},\ \sigma \vdash \oplus\{\, e_1 \,[\!]\, \bar{r} \,\} : t}{\sigma \vdash \oplus\{\, e_1 \,[\!]\, e_2, \bar{r} \,\} : t} \quad \text{(T16)}$$

Fig. 3. Typing rules of the monoid calculus.

variable names to types ($\sigma(v)$) and returns the binding of v in σ,, and $\sigma[t/v]$ extends σ with the binding from v to t. The λ-variable v in Eq. (T7) is annotated with its type t_1. This annotation is omitted in our examples because in most cases it can be inferred from context. This equation indicates that if the type of e is t_2 under the assumption that the type of v is t_1, then the type of $\lambda v : t_1.e$ is $t_1 \to t_2$ (a function from t_1 to t_2). Equation (T15) checks the validity of a comprehension by testing whether $\otimes \leq \oplus$. Note that according to these typing rules the term $1 + 2 \times 3$ is type-correct, since $T_+ = T_\times = \text{int}$ (from Table I.A), while $x +\!\!+\, y \cup z$ is not type-correct.

3. TRANSLATING OQL TO THE MONOID CALCULUS

Nearly all OQL expressions have a direct translation into the monoid calculus, with the exception of indexed OQL collections. (Section 5 presents a monoid for indexed collections as well as a comprehension syntax to capture complex vector and array operations.) Here we illustrate how to translate the main constructs of OQL into the monoid calculus.

A select-from-where OQL statement of the form

select e
from x_1 **in** $e_1, \ldots, x_n$ **in** e_n
where *pred*

is translated into

$$\uplus\{\, e \,[\!]\, x_1 \leftarrow e_1, \ldots x_n \leftarrow e_n, pred \,\}.$$

Note that e, e_i, and *pred* could all contain nested comprehension expressions, which supports the capability in OQL to have nested queries in the *select*, *where*, and *from* clauses. For example, the following nested OQL query:

```
select e.name, e.address
from el in (select d.instructors
            from d in Departments
            where d.name= "CSE"),
     e in el
where exists c in e.teaches: c.name= "cse5331"
```

is expressed in the comprehension syntax, as follows:

$$\uplus\{ \langle \text{Name: e.name, Address: e.address} \rangle \; [] \; \text{el} \leftarrow \uplus\{ \text{d.instructors} \; [] \; \text{d} \leftarrow \text{Departments, d.name=“CSE”} \}, \; \text{e} \leftarrow \text{el}, \vee\{ \text{c.name=“cse5331”} \; [] \; \text{c} \leftarrow \text{e.teaches} \} \}.$$

The **select-distinct** OQL statement is translated into a set comprehension in a similar way. For example, the following query against the university schema finds all departments that have at least one instructor who teaches more than 8 courses per year:

```
select distinct e.dept.name
from e in Instructors
where count(e.teaches) > 8.
```

This query is translated into the following comprehension:

$$\cup\{ \text{e.dept.name} \; [] \; \text{e} \leftarrow \text{Instructors}, \; +\{ 1 \; [] \; \text{c} \leftarrow \text{e.teaches} \} > 8 \}.$$

OQL universal and existential quantifications are mapped into comprehensions too, as shown in the following query that finds the instructors who teach courses that all have cse2320 as a prerequisite:

```
select e.name
from e in Instructors
where for all c in e.teaches:
exists d in c.has_prerequisites: d.name = "cse2320".
```

This query is translated into the following comprehension:

$$\uplus\{ \text{e.name} \; [] \; \text{e} \leftarrow \text{Instructors}, \wedge\{ \vee\{ \text{d.name=“cse2320”} \; [] \; \text{d} \leftarrow \text{c.has_prerequisites} \} \; [] \; \text{c} \leftarrow \text{e.teaches} \} \}.$$

The OQL group-by syntax is essentially syntactic sugar, since the same results can be accomplished with nested queries. For example, the query

```
select dept: dname, largest: max(e.salary)
from e in Instructors
group by dname: e.dept.name,
```

prints for each department the largest salary of all instructors working in the department is equivalent to

```
select dept: e.dept.name,
       largest: max(select s.salary
                    from s in Instructors
                    where e.dept.name = s.dept.name)
from e in Instructors,
```

which, in turn, can be expressed in the calculus as follows:

$\uplus\{ \langle$ dept: e.dept.name,
largest: $max\{$ s.salary [] s $\leftarrow$ Instructors, e.dept.name=s.dept.name $\}$ $\rangle$
[] e $\leftarrow$ Instructors $\}$.

We see in Section 7 that nested query forms such as these can be translated into efficient execution plans that make use of group-by algorithms. In general, a group-by OQL query takes the following form:

select e $(A_1, \ldots, A_n,$ partition$)$
from x_1 **in** u_1, x_2 **in** $u_2(x_1), \ldots, x_m$ **in** $u_m(x_1, \ldots, x_{m-1})$
where $p(x_1, \ldots, x_m)$
group by $A_1 : e_1(x_1, \ldots, x_m), \ldots, A_n : e_n(x_1, \ldots, x_m)$
having h $(A_1, \ldots, A_n,$ partition$)$,

where the variable *partition* is bound to the set of all groups, where each group is a record of all x_i values that have the same $A_1, \ldots, A_n$ values (it is used for complex aggregations). This OQL form is equivalent to

$$\begin{array}{l} \uplus\{ e\,(A_1, \ldots, A_n, \text{ partition}) \\ \;[]\; x_1 \leftarrow u_1, x_2 \leftarrow u_2(x_1), \ldots, x_m \leftarrow u_m(x_1, \ldots x_{m-1}), \\ \quad A_1 \equiv e_1(x_1, \ldots, x_m), \ldots, A_n \equiv e_n(x_1, \ldots, x_m), \\ \quad \text{partition} \equiv \uplus\{ \langle x_1 : y_1, \ldots, x_m : y_m \rangle, \\ \qquad\qquad [] \; y_1 \leftarrow u_1, y_2 \leftarrow u_2(y_1), \ldots, y_m \leftarrow u_m(y_1, \ldots, y_{m-1}), \\ \qquad\qquad\; e_1(y_1, \ldots, y_m) = A_1, \ldots, e_n(y_1, \ldots, y_m) = A_n \}, \\ \quad p(x_1, \ldots x_m), h(A_1, \ldots A_n, \text{partition}) \}. \end{array}$$

For example, the query

```
select salary: max(e.salary), dept: dname
from e in Instructors
group by dname: e.dept.name,
         teaches: count(select *
                        from c in e.teaches
                        where c.offered_by.name = "CSE")
```

is equivalent to

```
⊎{ ⟨salary: max{ p.e.salary [] p ← partition }, dept: dname ⟩
  [] e ← Instructors,
    dname ≡ e.dept.name,
    teaches ≡ +{ 1 [] c ← e.teaches, c.offered_by.name="CSE" }
    partition ≡ ⊎{ ⟨ e: s ⟩ [] s ← Instructors,
                            s.dept.name=d.name
                            +{ 1 [] c ← s.teaches, c.offered_by.name="CSE" }
                                   = teaches } }.
```

The following table gives the translation of other OQL expressions into monoid calculus:

e_1 **intersect** e_2	$\rightarrow$	$\cup\{ x \;[]\; x \leftarrow e_1, x \in e_2 \}$
for all x **in** $e : pred$	$\rightarrow$	$\wedge\{ pred \;[]\; x \leftarrow e \}$
exists x **in** $e : pred$	$\rightarrow$	$\vee\{ pred \;[]\; x \leftarrow e \}$
e_1 **in** e_2	$\rightarrow$	$\vee\{ x = e_1 \;[]\; x \leftarrow e_2 \}$
count(e)	$\rightarrow$	$+\{ 1 \;[]\; x \leftarrow e \}$
sum(e)	$\rightarrow$	$+\{ x \;[]\; x \leftarrow e \}$
flatten(e)	$\rightarrow$	$\cup\{ x \;[]\; s \leftarrow e, x \leftarrow s \}$.

$$(\lambda v.e_1)\, e_2 \longrightarrow e_1[e_2/v] \qquad \text{beta reduction} \qquad \text{(N2)}$$

$$\langle A_1 = e_1, \ldots, A_n = e_n\rangle.A_i \longrightarrow e_i \qquad \text{(N3)}$$

$$\oplus\{\, e \,\|\, \overline{q},\, v \leftarrow (\textbf{if}\, e_1\, \textbf{then}\, e_2\, \textbf{else}\, e_3),\, \overline{s}\,\} \longrightarrow (\oplus\{\, e \,\|\, \overline{q},\, e_1,\, v \leftarrow e_2,\, \overline{s}\,\}) \oplus (\oplus\{\, e \,\|\, \overline{q},\, \neg e_1,\, v \leftarrow e_3,\, \overline{s}\,\}) \qquad \text{(N4)}$$

for commutative $\oplus$ or empty $\overline{q}$

$$\oplus\{\, e \,\|\, \overline{q},\, v \leftarrow \mathcal{Z}_\otimes,\, \overline{s}\,\} \longrightarrow \mathcal{Z}_\oplus \qquad \text{(N5)}$$

$$\oplus\{\, e \,\|\, \overline{q},\, v \leftarrow \mathcal{U}_\otimes(e'),\, \overline{s}\,\} \longrightarrow \oplus\{\, e \,\|\, \overline{q},\, v \equiv e',\, \overline{s}\,\} \qquad \text{(N6)}$$

$$\oplus\{\, e \,\|\, \overline{q},\, v \leftarrow (e_1 \otimes e_2),\, \overline{s}\,\} \longrightarrow (\oplus\{\, e \,\|\, \overline{q},\, v \leftarrow e_1,\, \overline{s}\,\}) \oplus (\oplus\{\, e \,\|\, \overline{q},\, v \leftarrow e_2,\, \overline{s}\,\}) \qquad \text{(N7)}$$

for commutative $\oplus$ or empty $\overline{q}$

$$\oplus\{\, e \,\|\, \overline{q},\, v \leftarrow \otimes\{\, e' \,\|\, \overline{r}\,\},\, \overline{s}\,\} \longrightarrow \oplus\{\, e \,\|\, \overline{q},\, \overline{r},\, v \equiv e',\, \overline{s}\,\} \qquad \text{(N8)}$$

$$\oplus\{\, e \,\|\, \overline{q},\, \vee\{\, pred \,\|\, \overline{r}\,\},\, \overline{s}\,\} \longrightarrow \oplus\{\, e \,\|\, \overline{q},\, \overline{r},\, pred,\, \overline{s}\,\} \qquad \text{for idempotent } \oplus \qquad \text{(N9)}$$

$$\oplus\{\, \oplus\{\, e \,\|\, \overline{r}\,\} \,\|\, \overline{s}\,\} \longrightarrow \oplus\{\, e \,\|\, \overline{s},\, \overline{r}\,\} \qquad \text{for a primitive monoid } \oplus \qquad \text{(N10)}$$

Fig. 4. The normalization algorithm.

The restriction on monoid comprehensions relating idempotence and commutativity of the monoids involved turns out not to be a limitation in translation. OQL *select* statements always return sets or bags, and the only explicit conversion function on collections is **listtoset**, which the calculus allows. Nevertheless, there are OQL language constructs that cannot be simulated in the monoid calculus. They include bag difference and set or bag element (which retrieves an element from a singleton set or bag, respectively).

4. PROGRAM NORMALIZATION

The monoid calculus can be put into a canonical form by an efficient rewrite algorithm, called the *normalization algorithm*. The evaluation of these canonical forms generally produces fewer intermediate data structures than the initial unnormalized programs. Moreover, in many cases, the normalization algorithm improves program performance. It generalizes many optimization techniques already used in relational algebra, such as pushing a selection before a join and fusing two selections into one selection. While normalization does not perform cost-based optimization, it does provide a good starting point for it.

Figure 4 gives the normalization rules. For example, Rule (N5) applies when a comprehension has a generator whose domain is a zero value (i.e., $\mathcal{Z}_\otimes$) that appears between two possibly empty sequences of generators $\overline{q}$ and $\overline{s}$. In that case, the comprehension is normalized into the value $\mathcal{Z}_\oplus$, since no values are generated for the variable v. Rule (N8) is the most important: it flattens a nested comprehension (i.e., a comprehension that contains a generator whose domain is another comprehension). Rule (N9) unnests an existential quantification. There are other cases of query unnesting in Section 7.

One advantage of the normalization algorithm, or any algorithm on calculus expressions expressed via pattern-based rewrite, is that it can be shown to correctly preserve meaning by proving each rewrite transformation correct.

THEOREM 3. *The normalization rules in Figure 4 are meaning-preserving.*

A proof sketch of this theorem is given in Appendix A. In addition, this set of rewrite rules is terminating and confluent (i.e., it satisfies the Church-Rosser property). That is, it does not matter in which sequence rules are applied to a term, as all sequences result in the same canonical form at the end. The rewrite rules are terminating because they decrease the number of nesting levels. They are confluent, since it can be seen from Rule (N8) that

$$\oplus\{\, e \;[]\; \overline{q},\, v_1 \leftarrow \otimes\{\, e_1 \;[]\; \overline{r_1}\,\},\, \overline{r},\, v_2 \leftarrow \otimes\{\, e_2 \;[]\; \overline{r_2}\,\},\, \overline{s}\,\}$$

$$= \oplus\{\, e \;[]\; \overline{q},\, \overline{r_1},\, v_1 \equiv e_1,\, \overline{r},\, v_2 \leftarrow \otimes\{\, e_2\, \overline{r_2}\,\},\, \overline{s}\,\}$$

$$= \oplus\{\, e \;[]\; \overline{q},\, \overline{r_1},\, v_1 \equiv e_1,\, \overline{r},\, \overline{r_2},\, v_2 \equiv e_2,\, \overline{s}\,\},$$

which is also the result of $\oplus\{\, e \;[]\; \overline{q},\, v_1 \leftarrow \otimes\{\, e_1 \;[]\; \overline{r_1}\,\},\, \overline{r},\, \overline{r_2},\, v_2 \equiv e_2,\, \overline{s}\,\}$. This identity shows that both chains of rewrites (left inner comprehension first or right inner comprehension first) lead to the same reduced form.

Rules (N8) and (N9) may require some variable renaming to avoid name conflicts. If there is a generator $v' \leftarrow e_1$ in $\overline{q}$ and a generator $v' \leftarrow e_2$ in $\overline{r}$, then variable v' in $\overline{r}$ should be renamed. For example, filter(p)(filter(q) x) is

```
∪{ a [] a ← ∪{ a [] a ← x, q(a) }, p(a) }
= ∪{ a [] a ← ∪{ b [] b ← x, q(b) }, p(a) }
```

(by renaming the inner variable a to b) and is normalized into

```
∪{ a [] b ← x, q(b), a ≡ b, p(a) }
→ ∪{ b [] b ← x, q(b), p(b) }
```

(by Rules (N8) and (N1)), which is a filter whose predicate is the conjunction of p and q. As another example of normalization, consider the following nested OQL query:

```
select distinct r
   from r in R
   where r.B in (select distinct s.D
                    from s in S
                    where r.C=s.C),
```

which is expressed in the monoid calculus, as follows:

```
∪{ r [] r ← R, ∨{ x=r.B [] x ← ∪{ s.D [] s ← S, r.C=s.C } } }
```

and is normalized into

```
∪{ r [] r ← R, x ← ∪ { s.D [] s ← S, r.C=s.C }, x=r.B }
→ ∪{ r [] r ← R, s ← S, r.C=s.C, x ≡ s.D, x=r.B }
→ ∪{ r [] r ← R, s ← S, r.C=s.C, s.D=r.B }
```

(by Rules (N9), (N8), and (N1)), which is equivalent to the following OQL query:

```
select distinct r
  from r in R, s in S
  where r.C=s.C and s.D = r.B.
```

Note that although we were able to rewrite this unnested form back into OQL, not every comprehension has a direct translation back to OQL.

Now consider the query, presented in Section 3, which finds all CSE instructors who teach cse5331:

$$\uplus\{\!\{ \langle \text{Name: e.name, Address: e.address} \rangle \;[]\; \text{el} \leftarrow \uplus\{\!\{ \text{d.instructors} \;[]\; \text{d} \leftarrow \text{Departments, d.name=“CSE”} \}\!\}, \text{e} \leftarrow \text{el}, \vee\{\!\{ \text{c.name=“cse5331”} \;[]\; \text{c} \leftarrow \text{e.teaches} \}\!\} \}\!\}.$$

This query is normalized into

$$\begin{aligned}
&\uplus\{\!\{ \langle \text{Name: h.name, Address: h.address} \rangle \\
&\quad [] \; \text{d} \leftarrow \text{Departments, d.name=“CSE”, el} \equiv \text{d.instructors}, \\
&\quad\; \text{e} \leftarrow \text{el}, \vee\{\!\{ \text{c.name=“cse5331”} \;[]\; \text{c} \leftarrow \text{e.teaches} \}\!\} \}\!\} \\
&\rightarrow \uplus\{\!\{ \langle \text{Name: h.name, Address: h.address} \rangle \\
&\qquad [] \; \text{d} \leftarrow \text{Departments, d.name=“CSE”}, \\
&\qquad\; \text{e} \leftarrow \text{d.instructors}, \vee\{\!\{ \text{c.name=“cse5331”} \;[]\; \text{c} \leftarrow \text{e.teaches} \}\!\} \}\!\} \\
&\rightarrow \uplus\{\!\{ \langle \text{Name: h.name, Address: h.address} \rangle \\
&\qquad [] \; \text{d} \leftarrow \text{Departments, d.name=“CSE”}, \\
&\qquad\; \text{e} \leftarrow \text{d.instructors, c} \leftarrow \text{e.teaches, c.name=“cse5331”} \}\!\}
\end{aligned}$$

(by Rules (N9), (N8), and (N1)).

Canonical forms: A path *path* is a bound variable, the name of a class extent, or an expression *path′.name*, where *name* is an attribute name of a record and *path′* is a path. If a comprehension has a commutative accumulator, then the generator domains in this comprehension can be normalized into paths. If, in addition, all predicates of a comprehension are moved at the end of the comprehension that forms a conjunction *pred* of predicates, then monoid comprehensions can be put into the following canonical form:

$$\oplus\{\!\{ e \;[]\; v_1 \leftarrow path_1, \ldots v_n \leftarrow path_n, pred \}\!\},$$

where each $path_i$ is a path. The proof of this statement is easy. First, it is easy to see that the previous form is a canonical form (i.e., it cannot be normalized further), since there is no normalization rule that reduces generators over paths. Second, if the domain e of a generator $v \leftarrow e$ in a monoid comprehension is a form other than a path, then this domain is reduced to a simpler form by the normalization algorithm. This situation becomes apparent if we consider all the different forms that e can take. For each such form there is a normalization rule, thus leaving the only form that cannot be reduced, namely a path.

Performance: Our normalization algorithm unnests all type N and J nested queries [Kim 1982] (using Rules (N8) and (N9), respectively). The important question, though, is whether normalization always improves performance. Unfortunately not. Consider the term $\cup\{\!\{ (v, v) \;[]\; v \leftarrow \cup\{\!\{ E \;[]\; w \leftarrow X \}\!\} \}\!\}$, where E is an expensive query. This term is normalized into

$\cup\{\!\{ (E, E) \,[]\, w \leftarrow X \}\!\}$, which performs the computation of E twice. In this case, the normalized form may perform worse than the original term. Cases such as these occur frequently in lazy functional languages [Peyton Jones 1987]. In those languages, function application is evaluated using beta reduction (Rule (N2)), which, if it is implemented naively as term substitution, may repeat computations (if v appears more than once in e_1). To avoid situations like these, the evaluators of these languages use graph-reduction techniques [Peyton Jones 1987] in which all occurrences of v in e_1 share the same term by pointing to the same memory address, thus forming an acyclic graph. When this term is reduced to a value, the term is replaced by this value in the graph, thus avoiding the need to compute this value twice. If we apply this technique to our normalization algorithm, the normalized form $\cup\{\!\{ (E, E) \,[]\, w \leftarrow X \}\!\}$ will not repeat the evaluation of E; instead it will use two pointers to the same copy of the term E.

If we use graph reduction during normalization, then, under the naive cost function described below, normalization always improves cost. Of course, this result may not necessarily be true for realistic cost functions and physical algorithms. The purpose of this cost analysis is to give a simple explanation why normalization is an effective optimization technique. It should be noted though that, like any relational algebraic optimization, normalization is a heuristic that improves performance in most cases. The true effectiveness of an optimization algorithm can only be asserted by extensive testing applied to real-world applications.

We assume that we know the average sizes of all sets reached by paths in the database schema: For example, in $\cup\{\!\{ e \,[]\, x \leftarrow X, y \leftarrow x.A \}\!\}$, the size of $x.A$ is the average of all sets $x.A$ for every $x \in X$. The following cost function assumes a naive nested-loop evaluation that does not depend on the order of the joins:

$$
\begin{aligned}
\text{size}[\![\oplus\{\!\{ e \,[]\, \bar{r} \}\!\}]\!] &= \text{size}[\![\bar{r}]\!] \\
\text{size}[\![\bar{r}, v \leftarrow e]\!] &= \text{size}[\![e]\!] \times \text{size}[\![\bar{r}]\!] \\
\text{size}[\![\bar{r}, \mathit{pred}]\!] &= \text{size}[\![\bar{r}]\!] \times \text{selectivity}[\![\mathit{pred}]\!] \\
\text{size}[\![\bar{r}, v \equiv e]\!] &= \text{size}[\![\bar{r}]\!] \\
\text{size}[\![]\!] &= 1
\end{aligned}
$$

$$
\begin{aligned}
\text{cost}[\![\oplus\{\!\{ e \,[]\, \bar{r} \}\!\}]\!] &= \text{cost}[\![e]\!] + \text{cost}[\![\bar{r}]\!] + \text{size}[\![r]\!] \\
\text{cost}[\![\bar{r}, v \leftarrow e]\!] &= \text{cost}[\![e]\!] + \text{cost}[\![\bar{r}]\!] \\
\text{cost}[\![\bar{r}, \mathit{pred}]\!] &= \text{cost}[\![\mathit{pred}]\!] + \text{cost}[\![\bar{r}]\!] \\
\text{cost}[\![\bar{r}, v \equiv e]\!] &= \text{cost}[\![e]\!] + \text{cost}[\![\bar{r}]\!] \\
\text{cost}[\![]\!] &= 0
\end{aligned}
$$

where size$[\![e]\!]$ of a collection e estimates the size of e, while selectivity $[\![\mathit{pred}]\!]$ is the selectivity of the predicate *pred*. The variable-binding cost

equation (for $v \equiv e$) assumes the use of graph reduction, i.e., the computation of e is done only once, even when v is used more than once.

For example, the join $X \bowtie_{x.A=y.B} Y$ has cost cost$[\![\cup\{ (x, y) \,[\!]\, x \leftarrow X, y \leftarrow Y, x.A = y.B \}]\!]$ equal to size$[\![X]\!]$ × size$[\![Y]\!]$ × selectivity$[\![x.A = y.B]\!]$. The size function satisfies size$[\![\bar{r}, \bar{s}]\!]$ = size$[\![\bar{r}]\!]$ × size$[\![\bar{s}]\!]$ and the cost function satisfies cost$[\![\bar{r}, \bar{s}]\!]$ = cost$[\![\bar{r}]\!]$ + cost$[\![\bar{s}]\!]$. Under this cost function, we can prove that each normalization rule improves cost. For example, the left-hand side of Rule (N8) has cost

$$\begin{aligned}
&\text{cost}[\![\oplus\{ e \,[\!]\, \bar{q}, v \leftarrow \cup\{e' \,[\!]\, \bar{r} \}, \bar{s} \}]\!] \\
&= \text{cost}[\![e]\!] + \text{cost}[\![\bar{q}, v \leftarrow \cup\{ e' \,[\!]\, \bar{r} \}, s]\!] + \text{size}[\![\bar{q}, v \leftarrow \cup\{ e' \,[\!]\, \bar{r} \}, s]\!] \\
&= \text{cost}[\![e]\!] + \text{cost}[\![\bar{q}]\!] + \text{cost}[\![\cup\{ e' \,[\!]\, \bar{r} \}]\!] + \text{cost}[\![\bar{s}]\!] + \text{size}[\![\bar{q}]\!] \times \text{size}[\![\cup\{ e' \,[\!]\, \bar{r} \}]\!] \times \text{size}[\![\bar{s}]\!] \\
&= \text{cost}[\![e]\!] + \text{cost}[\![\bar{q}]\!] + \text{cost}[\![e']\!] + \text{cost}[\![\bar{r}]\!] + \text{size}[\![\bar{r}]\!] + \text{cost}[\![\bar{s}]\!] + \text{size}[\![\bar{q}]\!] \times \text{size}[\![\bar{r}]\!] \times \text{size}[\![\bar{s}]\!]
\end{aligned}$$

which has an excess of size$[\![\bar{r}]\!]$ over the cost of the right-hand side of the rule

$$\begin{aligned}
&\text{cost}[\![\oplus\{ e \,[\!]\, \bar{q}, \bar{r}, v \equiv e', \bar{s} \}]\!] \\
&= \text{cost}[\![e]\!] + \text{cost}[\![\bar{q}, \bar{r}, v \equiv e', \bar{s}]\!] + \text{size}[\![\bar{q}, \bar{r}, v \equiv e', \bar{s}]\!] \\
&= \text{cost}[\![e]\!] + \text{cost}[\![\bar{q}]\!] + \text{cost}[\![\bar{r}]\!] + \text{cost}[\![e']\!] + \text{cost}[\![\bar{s}]\!] + \text{size}[\![\bar{q}]\!] \times \text{size}[\![\bar{r}]\!] \times \text{size}[\![\bar{s}]\!].
\end{aligned}$$

Limitations: Even though the normalization algorithm unnests many forms of nested queries, there are still some forms of queries that cannot be unnested by it. The following query contains three examples of such forms:

$\cup\{ \langle$ E=e, M= $\cup\{$ c [] c ← e.children, $\wedge\{$ c.age > d.age [] d ← e.manager.children } } $\rangle$
[] e ← Employees, e.salary>$max\{$ m.salary [] m ← Managers, e.age>m.age } }.

The subquery $\cup\{$ c [] c ← e.children, $\wedge\{$ c.age > d.age [] d ← e.manager.children } } cannot be unnested by the normalization algorithm because the computed set must be embedded in the result of every iteration of the outer-set comprehension. Similarly, the universal quantification (the $\wedge$-comprehension) and the aggregation (the max-comprehension) cannot be unnested by the normalization algorithm. These cases (which are types A and JA nested queries [Kim 1982]) require the use of outer-joins and grouping, and they will be treated properly in Section 7.

5. MODEL EXTENSIONS

5.1 Handling Inconsistencies Due to Idempotence

In Section 2.2 we showed that if we allow homomorphisms from idempotent to nonidempotent monoids, such as set cardinality, hom$[\cup, +](\lambda x.1)(A)$, semantic inconsistencies may result. The reason is that, by definition, monoid homomorphisms are divide-and-conquer computations. For set cardinality, this property implies that

$$\text{hom}[\cup, +](\lambda x.1)(A \cup B) = (\text{hom}[\cup, +](\lambda x.1)(A) + \text{hom}[\cup, +](\lambda x.1)(B)).$$

But this equation is not valid when $A \cap B \neq \{\}$, because it counts the common elements between A and B twice. The correct equation is

$$\text{hom}[\cup, +](\lambda x.1)(A \cup B) = (\text{hom}[\cup, +](\lambda x.1)(A) + \text{hom}[\cup, +](\lambda x.1)(B - A)).$$

This observation can be generalized to any monoid. The third equation of Definition 2 should be

$$\text{hom}[\oplus, \otimes](f)(x \oplus y) = (\text{hom}[\oplus, \otimes](f)(x) \otimes \text{hom}[\oplus, \otimes](\lambda v.F(v, x))(y)),$$

where function F is defined as follows:

$$F(v, x) = \begin{cases} \textbf{if } v \in x \textbf{ then } \mathcal{Z}_\oplus \textbf{ else } f(v) & \text{for idempotent } \otimes \\ f(v) & \text{otherwise} \end{cases}$$

Given this definition, if $\otimes$ is an idempotent collection monoid, then according to Theorem 2:

$$\begin{aligned} \text{hom}[\oplus, \otimes](\lambda v.F(v, x))(x) &= \otimes\{\!\{ w \;[]\; v \leftarrow x, w \leftarrow F(v, x) \}\!\} \\ &= \otimes\{\!\{ w \;[]\; v \leftarrow x, w \leftarrow (\textbf{if } v \in x \textbf{ then } \mathcal{Z}_\oplus \textbf{ else } f(v)) \}\!\} \\ &= \otimes\{\!\{ w \;[]\; v \leftarrow x, w \leftarrow \mathcal{Z}_\oplus \}\!\} = \mathcal{Z}_\otimes. \end{aligned}$$

Consequently,

$$\begin{aligned} \text{hom}[\oplus, \otimes](f)(x \oplus x) &= (\text{hom}[\oplus, \otimes](f)(x)) \otimes (\text{hom}[\oplus, \otimes](F)(x)) \\ &= (\text{hom}[\oplus, \otimes](f)(x)) \otimes \mathcal{Z}_\otimes \\ &= \text{hom}[\oplus, \otimes](f)(x), \end{aligned}$$

which holds even when $\oplus$ is anti-idempotent. A similar proof holds for an idempotent primitive monoid $\otimes$.

In addition to the definition above, Rule (N7) of the normalization algorithm (Figure 4) must be modified accordingly:

$$\oplus\{\!\{ e \;[]\; \bar{q}, v \leftarrow (e_1 \otimes e_2), \bar{s} \}\!\}$$

$$\rightarrow \begin{cases} (\oplus\{\!\{ e \;[]\; \bar{q}, v \leftarrow e_1, \bar{s} \}\!\}) \oplus (\oplus\{\!\{ e \;[]\; \bar{q}, v \leftarrow e_2, v \notin e_1 \bar{s} \}\!\}) & \text{if } \oplus \not\preceq \otimes \\ (\oplus\{\!\{ e\; \bar{q}, v \leftarrow e_1, \bar{s} \}\!\}) \oplus (\oplus\{\!\{ e \;[]\; \bar{q}, v \leftarrow e_2, \bar{s} \}\!\}) & \text{otherwise} \end{cases}$$

for commutative $\oplus$ and $\otimes$. Note that the first alternative is when $\oplus \not\preceq \otimes$, that is, when $\oplus$ is commutative and $\otimes$ is both commutative and idempotent. For example, $+\{\!\{ 1 \;[]\; x \leftarrow (\{1\} \cup \{1\}) \}\!\}$ is normalized into $(+\{\!\{ 1 \;[]\; x \leftarrow \{1\} \}\!\}) + (+\{\!\{ 1 \;[]\; x \leftarrow \{1\}, x \notin \{1\} \}\!\})$, which is equal to $1 + 0 = 1$.

5.2 Advanced List Operations

A primitive monoid useful for lists is the *function composition monoid*, $(\circ, \lambda x.x)$, where the function composition, $\circ$, is defined by the equation $(f \circ g)x = f(g(x))$ and is associative, but neither commutative nor idempotent. Even though the type of this monoid, $T_\circ(\alpha) = \alpha \leftarrow \alpha$, is parametric, it is still a primitive monoid. For a list $X = [a_1, a_2, \ldots, a_n]$, $\circ\{\!\{ \lambda x.f(a, x) \;[]\; a \leftarrow X \}\!\}$ is equal to $(\lambda x.f(a_1, x)) \circ (\lambda x.f(a_2, x)) \circ \cdots \circ (\lambda x.f(a_n, x))$. Consequently, $\circ\{\!\{ \lambda x.f(a, x) \;[]\; a \leftarrow X \}\!\}(v)$ computes the value $f(a_1, f(a_2, \ldots, f(a_n, v)))$. For example, the reverse of the list X is $\circ\{\!\{ \lambda x.x +\!\!+ [a] \;[]\; a \leftarrow X \}\!\}([\,])$, and the first element of X is $\circ\{\!\{ \lambda x.a \;[]\; a \leftarrow X \}\!\}(\text{NULL})$.

The function composition monoid offers enormous expressive power to list operations because it can be used to compose functions that propagate a state during a list iteration. For example, the n^{th} element of a list X, $\text{nth}(X, n)$, is

$$(\circ\{\!\{ \lambda(x, k).(\textbf{if } k = n \textbf{ then } a \textbf{ else } x, k - 1) \;[]\; x \leftarrow X \}\!\}(\text{NULL}, \text{length}(X))).\text{fst},$$

which uses a counter k as the state to propagate through the iteration over X. A more complex example is sorting the elements of a list X using a total order $\leq$:

$$\text{sort}(X) = \circ\{\!\{ \lambda z.\text{insert}(a, z) \;[]\; a \leftarrow X \}\!\}([\,])$$
$$\text{insert}(a, Y) = (\circ\{\!\{ \lambda(z, p).\textbf{if } p \vee a \leq b \textbf{ then } ([b] +\!\!+ z, p) \textbf{ else } ([b, a] +\!\!+ z, \text{true}) \;[]\; b \leftarrow Y \}\!\}$$
$$([\,], \text{false})).\text{fst}.$$

List comprehensions of the form $\circ\{\!\{ \lambda x.f(a, x) \;[]\; a \leftarrow X \}\!\}(v)$ are equivalent to list folds [Fegaras 1993], also known as a catamorphism [Meijer et al. 1991]. Like list folds, list comprehensions can capture the list primitive recursive function, $\text{P}(f, v)(X)$, defined by the following recursive equations:

$$\begin{aligned} \text{P}(f, v)([\,]) &= v \\ \text{P}(f, v)([a] +\!\!+ x) &= f(a, x, \text{P}(f, v)(x)) \end{aligned}$$

as follows:

$$\text{P}(f, v)(X) = (\circ\{\!\{ \lambda(x, r).(f(a, r, x), [a] +\!\!+ r) \;[]\; a \leftarrow X \}\!\}(v, [\,])).\text{fst}.$$

Note that $\text{P}(f, v)(X)$ represents a paramorphism [Meijer et al. 1991], which can be simulated easily by a catamorphism (i.e., a fold) by returning a pair that carries the rest of the list along with the result. When some simple syntactic restrictions are applied to the list primitive recursive functions, these functions precisely characterize the polynomial time space [Bellantoni and Cook 1992; Leivant 1993].

Unfortunately, the expressive power gained by the function composition monoid comes with a high cost. Comprehensions of the form $\oplus\{\!\{ e \;[]\; \bar{r}, v \leftarrow \circ\{\!\{ f \;[]\; \bar{s} \}\!\}(u), \bar{q} \}\!\}$ are very difficult to normalize and unnest.

5.3 Vectors and Arrays

Vectors and arrays are important collection types for scientific and other applications [Maier and Vance 1993; Libkin et al. 1996]. In contrast to other collection types, there is no obvious monoid that captures vectors effectively. Vector operations should provide random access through indexing as well as bulk manipulation. We first propose an effective form of vector comprehensions and then describe a monoid that captures these comprehensions. An example of a vector manipulation is vector reverse, which computes $x[n - i - 1]$ for $i = 0, 1, \ldots, n - 1$ from a vector x of size n. This function can be specified by

$$\Box\{\!\{ a[n - i - 1] \,[\!]\, a[i] \leftarrow x \}\!\},$$

where $\Box$ is the anticipated monoid for vectors. Note that we want to access both the value a and the associated index i from the vector x, but we do not want to impose any order on the way $a[i]$ is accessed. The generator $a[i] \leftarrow x$ accesses the pairs (a, i) in some unspecified order, much as elements of a set x are retrieved in a comprehension by the generator $a \leftarrow x$. But the constructed element of the vector comprehension above is $a[n - i - 1]$, which means that we should be able to store an element at any position in the resulting vector. That is, even though vector elements are accessed in bulk fashion (in pairs of value-index values), vectors are constructed in random fashion by specifying which value is stored at what place. But there is a problem here: What if we have two different values stored in the same place in a vector? We need to perform a merge operation on vectors. One solution is to merge the vector elements individually [Fegaras and Maier 1995]. That is, when two elements are stored at the same place in a vector, the resulting vector value is computed by merging these two elements. Another solution, which we adopt here, is to choose the latter of the two values if this value is not null, otherwise to choose the first. That way, the last non-null value overwrites the previous values.

We now formalize these observations. We introduce a new collection monoid $(\Box, \mathcal{Z}_\Box, \mathcal{U}_\Box)$ to represent vectors of any type. A vector of type vector(α) is represented by the type set(int $\times$ α), which specifies a partial mapping from positive integers (i.e., vector indexes) to values (of type α). This monoid has the following primitives: (The unit function here is binary, but in functional languages any n-ary function is equivalent to a unary function applied to a tuple.)

$$\begin{aligned}
\mathcal{Z}_\Box &= \{\} \\
\mathcal{U}_\Box(i, a) &= \{(i, a)\} \\
x \Box y &= y \cup (\cup\{\!\{ (i, a) \,[\!]\, (i, a) \leftarrow x, \wedge\{\!\{ i \neq j \,[\!]\, (j, b) \leftarrow y \}\!\} \}\!\}).
\end{aligned}$$

That is, when the vector elements of x are merged with the vector elements of y, the result is y union, the elements of x that are not indexed in y. This

monoid is idempotent but not commutative. In database terminology, the integer index (the first component of the pair int $\times\ \alpha$) is the key of set(int $\times\ \alpha$). The vector merge function, $\square$, reflects the common knowledge that in a vector assignment $A[i] := v$, the new value v overwrites the old value $A[i]$. Like any other monoid, the only purpose of the vector monoid is to provide meaning to vector operations, not to suggest in any way how vectors should be implemented.

For example, the set $\{(4, 10), (6, 13), (2, 11), (3, 12)\}$ corresponds to the vector <NULL, 11, 12, 10, NULL, 13>. If this set is merged with the set $\{(1, 5), (5, 7), (3, 6)\}$, which corresponds to the vector <5, NULL, 6, NULL, 7>, it constructs the set $\{(1, 5), (5, 7), (3, 6), (2, 11), (4, 10), (6, 13)\}$, which corresponds to <5, 11, 6, 10, 7, 13>.

The following are examples of vector manipulations. To make the comprehensions more readable, we represent a pair of the form (a, i) as $a[i]$.

$$
\begin{array}{lcl}
\text{sum_all}(x) & = & +\{\!\{\, a \;[\!]\; a[i] \leftarrow x \,\}\!\} \\
\text{subseq}(x, n, l) & = & \square\{\!\{\, a[i - n] \;[\!]\; a[i] \leftarrow x, i \geq n, i < l + n \,\}\!\} \\
\text{permute}(x, p) & = & \square\{\!\{\, a[b]\; a[i] \leftarrow x, b[j] \leftarrow p, i = j \,\}\!\} \\
\text{concat}(x, y, n) & = & x \;\square\; (\square\{\!\{\, a[n + i] \;[\!]\; a[i] \leftarrow y \,\}\!\}) \\
\text{inner}(x, y) & = & +\{\!\{\, a \times b \;[\!]\; a[i] \leftarrow x, b[j] \leftarrow y, i = j \,\}\!\}.
\end{array}
$$

Here subseq(x, n, l) returns the vector $x[i], i = n, \ldots, l + n - 1$, permute$(x, p)$ returns the vector $x[p[i]], i = 1, \ldots, n$, and inner$(x, y)$ returns the inner product of the two vectors x and y. A matrix can be represented as a vector of vectors, but the syntax becomes a little bit awkward:

$$
map(f)x = \square\{\!\{\, (\square\{\!\{\, f(b) \;[\!]\; b[j] \leftarrow a \,\}\!\})[i] \;[\!]\; a[i] \leftarrow x \,\}\!\}.
$$

Function map(f) maps function f to all the elements of a matrix. In the comprehension above, x is a matrix and thus a is a vector. Therefore, the head of the comprehension should be $z[i]$ where z is a vector. Thus, z must be another vector comprehension. Fortunately, there is a better approach to matrices. We can consider the monoid $\square$ to be of type set(index $\times\ \alpha$), where index is a type with equality. For simple vectors, index=int and for matrices, index= int $\times$ int. The actual type of the index can be inferred at compile time. The following are examples of matrix operations (we use the shorthand $a[i, j]$ for $(a, (i, j))$):

$$
\begin{array}{lcl}
\text{map}(f)x & = & \square\{\!\{\, (f(a))[i, j] \;[\!]\; a[i, j] \leftarrow x \,\}\!\} \\
\text{transpose}(x) & = & \square\{\!\{\, a[j, i] \;[\!]\; a[i, j] \leftarrow x \,\}\!\} \\
\text{row}(x, k) & = & \square\{\!\{\, a[j] \;[\!]\; a[i, j] \leftarrow x, i = k \,\}\!\} \\
\text{column}(x, k) & = & \square\{\!\{\, a[i] \;[\!]\; a[i, j] \leftarrow x, j = k \,\}\!\} \\
\text{multiply}(x, y) & = & \square\{\!\{\, v[i, j] \;[\!]\; a[i, j] \leftarrow x, v \equiv \text{inner}(\text{row}(x, i), \text{column}(y, j)) \,\}\!\}.
\end{array}
$$

Here row(x, k) returns the k^{th} row of matrix x while column(x, k) returns the k^{th} column. The only purpose of the generator $a[i, j] \leftarrow x$ in matrix multiplication multiply(x, y) is to generate the indexes i and j. The latter function is an example where the Haskell approach to vectors [Thompson 1998] (also adopted by Libkin et al. [1996]), in which vector indexes are first generated from integer domains and then used to access vectors, is more intuitive than ours.

Another example is matrix tiling, where a matrix X is partitioned into 4×4 tiles and each tile is replaced by the sum of the values in the tile:

$$\text{tiling}(X) = \Box\{\, (+\{\, b \,[\!]\, b[k, l] \leftarrow X, i - 4 \leq k \leq i + 4, j - 4 \leq l \leq j + 4 \,\})\left[\left\lceil \frac{i}{4} \right\rceil, \left\lceil \frac{j}{4} \right\rceil\right] \;[\!]\; a[i, j] \leftarrow x \,\}.$$

Normalization can be used to optimize vector and array operations. For example, transpose(transpose(x)) can be normalized into a comprehension equivalent to x.

OQL treats arrays like lists with indexing; it does not provide a convenient random bulk manipulation, as our vector comprehensions do. Although OQL can be extended with syntactic constructs to support our vectors. For example, map(f)x can be expressed in an OQL-like syntax as

select vector $(f(a))[i, j]$
from $a[i, j]$ **in** x.

6. INTERMEDIATE OBJECT ALGEBRA

In Section 2.3 we expressed comprehensions in terms of homomorphisms (Definition 3). If homomorphisms were implemented as iterations over collections, then comprehensions would be evaluated in a nested-loops fashion (since comprehensions with more than one generator are translated into nested homomorphisms). But in most cases we can do better than that. Research in relational query optimization has already addressed the related problem of efficiently evaluating join queries by considering different join orders, different access paths to data, and different join algorithms [Selinger et al. 1979]. To effectively adapt this technology to handle comprehensions, comprehensions must be expressed in terms of algebraic operators, which will eventually be mapped into physical execution algorithms like those found in database systems (e.g., merge join, hash join, etc). Section 7 describes a framework for translating terms in our calculus into efficient algorithms. This translation is done in stages. Queries in our framework are first translated into monoid comprehensions (as described in Section 3), which serve as an intermediate form, and are then translated into a version of the nested relational algebra that supports aggregation, quantification, outer-joins, and outer-unnests. At the end, the algebraic terms are translated into execution plans. This algebra is called the

$$X \bowtie_p Y = \{\!\{\, (v,w) \;[\!]\; v \leftarrow X,\, w \leftarrow Y,\, p(v,w) \,\}\!\} \tag{O1}$$

$$\sigma_p(X) = \{\!\{\, v \;[\!]\; v \leftarrow X,\, p(v) \,\}\!\} \tag{O2}$$

$$\mu_p^{path}(X) = \{\!\{\, (v,w) \;[\!]\; v \leftarrow X,\, w \leftarrow path(v),\, p(v,w) \,\}\!\} \tag{O3}$$

$$\Delta_p^{\oplus/e}(X) = \oplus\{\!\{\, e(v) \;[\!]\; v \leftarrow X,\, p(v) \,\}\!\} \tag{O4}$$

Fig. 5. The semantics of the basic intermediate algebra.

intermediate object algebra. We use both a calculus and an algebra as intermediate forms because the calculus closely resembles current OODB languages and is easy to normalize, while the algebra is lower-level and can be directly translated into the execution algorithms supported by database systems.

We classify the intermediate algebraic operators into two categories: basic and extended. We will prove that the basic algebra is equivalent to the monoid comprehension calculus, that is, all basic algebraic operators can be expressed in terms of the monoid calculus and vice versa. The purpose of the extended algebraic operators will become apparent when we present our query unnesting method (Section 7).

6.1 The Basic Intermediate Algebra

As a convenient notation for the semantics of the algebraic operators, we define $\{\!\{ e \,[\!]\, \bar{r} \}\!\}$ as follows:

$$\{\!\{\, e \;[\!]\; \bar{r} \,\}\!\} = \begin{cases} \cup\{\!\{\, e \;[\!]\; \bar{r} \}\!\} & \bar{r} \text{ has a commutative and idempotent generator} \\ \uplus\{\!\{\, e \;[\!]\; \bar{r} \,\}\!\} & \bar{r} \text{ has a commutative generator} \\ +\!\!+\{\!\{\, e \;[\!]\; \bar{r} \,\}\!\} & \text{otherwise.} \end{cases} \tag{D4}$$

Figure 5 defines the basic algebraic operators in terms of the monoid calculus. These operators generalize the operators of the nested relational algebra, so they have a straightforward explanation:

- **join**, $X \bowtie_p Y$, joins the collections X and Y using the join predicate p. This join is not necessarily between two sets; if, for example, X is a list and Y is a bag, then, according to Eq. (D4), the output is a bag.
- **selection**, $\sigma_p(X)$, selects all elements of X that satisfy the predicate p.
- **unnest**, $\mu_p^{path}(X)$, returns the collection of all pairs (x, y) for each $x \in X$ and for each $y \in x.path$ that satisfy the predicate $p(x, y)$.
- **reduce**, $\Delta_p^{\oplus/e}(X)$, collects the values $e(x)$ for all $x \in X$ that satisfy $p(x)$ using the accumulator $\oplus$. The reduce operator can be thought of as a generalized version of the relational projection operator.

$$\frac{\sigma \vdash X : T_{\oplus}(t_1),\ \sigma \vdash Y : T_{\otimes}(t_2),\ \sigma \vdash p : t_1 \times t_2 \rightarrow \text{bool},\ \odot = max(\oplus, \otimes)}{\sigma \vdash X \bowtie_p Y : T_{\odot}(t_1 \times t_2)} \quad \text{(T17)}$$

$$\frac{\sigma \vdash X : T_{\oplus}(t),\ \sigma \vdash p : t \rightarrow \text{bool}}{\sigma \vdash \sigma_p(X) : T_{\oplus}(t)} \quad \text{(T18)}$$

$$\frac{\sigma \vdash X : T_{\oplus}(t_1),\ \sigma \vdash path : t_1 \rightarrow T_{\otimes}(t_2),\ \sigma \vdash p : t_1 \times t_2 \rightarrow \text{bool},\ \odot = max(\oplus, \otimes)}{\sigma \vdash \mu_p^{path}(X) : T_{\odot}(t_2 \times t_1)} \quad \text{(T19)}$$

$$\frac{\sigma \vdash X : T_{\otimes}(t_1),\ \sigma \vdash e : t_1 \rightarrow T_{\oplus},\ \sigma \vdash p : t_1 \rightarrow \text{bool}}{\sigma \vdash \Delta_p^{\oplus/e}(X) : T_{\oplus}} \quad \text{(T20)}$$

$$\frac{\sigma \vdash X : T_{\otimes}(t_1),\ \sigma \vdash e : t_1 \rightarrow T_{\oplus},\ \sigma \vdash f : t_1 \rightarrow t_2 \quad \sigma \vdash p : t_1 \rightarrow \text{bool},\ \sigma \vdash g : t_1 \rightarrow t_3}{\sigma \vdash \Gamma_{p/g}^{\oplus/e/f}(X) : T_{\oplus}(t_1 \times T_{\oplus})} \quad \text{(T21)}$$

Fig. 6. The typing rules of the (basic and extended) algebraic operators.

The first four rules in Figure 6 are the typing rules for the basic algebraic operators.

THEOREM 4. *The basic intermediate algebra is equivalent to the monoid calculus.*

PROOF. The rules in Figure 5 translate the algebra into the calculus. The calculus-to-algebra translation algorithm is accomplished by the following rewrite rules:

$$[\![\oplus\{\, e \mid\!\mid \,\}]\!]_w E = \Delta_{\lambda w.\text{true}}^{\oplus/\lambda w.e}(E) \quad \text{(C1)}$$

$$[\![\oplus\{\, e_1 \mid\!\mid v \leftarrow e_2, \bar{r}\,\}]\!]_w E = [\![\oplus\{\, e_1 \mid\!\mid \bar{r}\,\}]\!]_{(w,\,v)}(\mu_{\lambda(w,\,v).\text{true}}^{\lambda w.e_2}(E)) \quad \text{(C2)}$$

$$[\![\oplus\{\, e \mid\!\mid p, \bar{r}\,\}]\!]_w E = [\![\{\, e \mid\!\mid \bar{r}\,\}]\!]_w(\sigma_{\lambda w.p}(E)) \quad \text{(C3)}$$

A monoid comprehension $\oplus\{\, e_1 \mid\!\mid v \leftarrow e_2, \bar{r}\,\}$ that may appear at any point in a calculus form (including the ones inside other comprehensions) is translated into an algebraic form by $[\![\oplus\{\, e_1 \mid\!\mid \bar{r}\,\}]\!]_v e_2$. The comprehension $\oplus\{\, e_1 \mid\!\mid \bar{r}\,\}$ is translated by compiling the qualifiers in $\bar{r}$ from left to right using the term E as a seed that grows at each step. That is, the term E is the algebraic tree derived at a particular point of compilation. The subscript variables w are all the range variables encountered so far during the translation. The validity of the above rules can be easily proved using the semantics in Figure 5. □

Note that Rule (C2) compiles every comprehension generator into an unnest. This step may not be necessary if e_2 is a class extent, in which case it may be better compiled into a join. In Section 7 we present an effective algorithm that translates the calculus into the algebra and at the same

time performs heuristic optimizations and query unnesting. As an example of how comprehensions are translated into intermediate algebraic forms, consider the following:

$$\cup\{\langle E = e.name, C = c.name\rangle \,[]\, e \leftarrow Employees, c \leftarrow e.children\},$$

which is translated into the algebraic form

$$\Delta^{\cup/\lambda(e,\,c).\langle E=e.name,\,C=c.name\rangle}_{\lambda(e,\,c).true}(\mu^{\lambda e.e.children}_{\lambda e.true}(Employees)).$$

Algebraic forms can be displayed as operator trees in which the tree leaves are collections of objects (e.g., class extents) and the output of the tree root is the output of the algebraic form. For example, the previous algebraic form has the following operator tree:

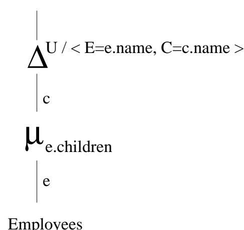

The functionality of an algebraic form can be better understood if we use a stream-based interpretation in which a stream of tuples flows from the leaves to the root of the tree. Under this interpretation, the algebraic form above generates a stream of tuples of type set(⟨ e: Employee ⟩) from the extent Employees. Variable e ranges over employees (i.e., it is of type Employee). The unnest operator, $\mu_{e.children}$, accepts the stream of tuples of type set(⟨ e: Employee ⟩) and constructs a stream of tuples of type set(⟨ e: Employee, c: Person ⟩), connecting each employee with one of his or her children. The reduce operator, $\Delta^{\cup/<E=e.name,\, C=c.name>}$, at the top of this algebraic form, is a generalization of the relational projection operator: It evaluates the expression ⟨E=e.name, C=c.name ⟩ for every input element and constructs a set from these tuples using ∪.

6.2 Extended Intermediate Algebra

This section presents an extension to the basic algebra well-suited to be the target of the query unnesting algorithm described in Section 7. This algorithm uses outer-unnests and outer-joins to relate the data between the inner and outer queries of a nested query and then reconstructs the result of the nested query using a group-by operator. The semantics of these operators cannot be given in terms of the monoid calculus if the outer query constructs a nonidempotent collection, since information about the exact number of copies of the data in the outer query may be lost after the outer join or unnest.

$$X \,\text{=}\!\bowtie_p Y \;=\; \{\!\{\, (v', w) \;[\!]\; v \leftarrow X,\; w \leftarrow \textbf{if} \wedge\{\!\{\, \neg p(v, w') \;[\!]\; v \neq \text{NULL},\; w' \leftarrow Y \,\}\!\} \;\textbf{then}\; [\text{NULL}] \;\textbf{else}\; \{\!\{\, w' \;[\!]\; w' \leftarrow Y,\; p(v, w') \,\}\!\} \,\}\!\} \quad \text{(O5)}$$

$$=\!\mu_p^{path}(X) \;=\; \{\!\{\, (v', w) \;[\!]\; v \leftarrow X,\; w \leftarrow \textbf{if} \wedge\{\!\{\, \neg p(v, w') \;[\!]\; v \neq \text{NULL},\; w' \leftarrow path(v) \,\}\!\} \;\textbf{then}\; [\text{NULL}] \;\textbf{else}\; \{\!\{\, w' \;[\!]\; w' \leftarrow path(v),\; p(v, w') \,\}\!\} \,\}\!\} \quad \text{(O6)}$$

$$\Gamma_{p/g}^{\oplus/e/f}(X) \;=\; \{\!\{\, (v, \oplus\{\!\{\, e(w) \;[\!]\; w \leftarrow X,\; g(w) \neq \text{NULL},\; v \doteq f(w),\; p(w) \,\}\!\}) \;[\!]\; v \leftarrow \textstyle\prod_f(X) \,\}\!\} \quad \text{(O7)}$$

Fig. 7. The semantics of the extended algebra.

We use the following new constructs to define the semantics of the extended algebra:

- Dot-equality $x \doteq y$ is defined if both x and y belong to the same collection C. If the collection C is idempotent, such as a set, then dot-equality is simply value-equality. Otherwise, every element x in C is dot-different from any other element y in C. Note that dot-equality is stronger than OID-equality. For example, the same object x may appear multiple times in a list, such as $[x, x]$, but the first element of this list is dot-different from the second. Even though dot-equality for C can easily be implemented by labeling the copies of an element in C with unique identifiers (such as unique TIDs), it cannot be expressed in terms of C using our basic algebra.
- $\prod_f(X)$ is equal to $\{\!\{ f(a) \;[\!]\; a \leftarrow X \}\!\}$, but with all duplicate elements under dot-equality removed.

These constructs satisfy the following properties for a function h:

$$\{\!\{ h(\oplus\{\!\{ e \;[\!]\; v' \leftarrow X, v' \doteq v, \bar{r} \}\!\}) \;[\!]\; v \leftarrow X \}\!\} = \{\!\{ h(\oplus\{\!\{ e \;[\!]\; v' \equiv v, \bar{r} \}\!\}) \;[\!]\; v \leftarrow X \}\!\} \quad \text{(P1)}$$

$$\textstyle\prod_{\lambda(v,\, w).v}(\{\!\{ (v', w') \;[\!]\; v' \leftarrow X, \bar{r} \}\!\}) = X \quad \text{(P2)}$$

Property (P1) indicates that the generator $v' \leftarrow X$ can be safely removed from the inner comprehension because the elements v' are dot-equal to the elements v of the outer comprehension. Property (P2) indicates that if we project over the elements of X and remove all duplicates, we get X.

The extended algebra operators are given in Figure 7 and are described below:

- **outer-join**, $X \,\text{=}\!\bowtie_p Y$, is a left outer-join between X and Y using the join predicate p. The domain of the second generator (the generator of w) in Eq. (O5) is always nonempty. If Y is empty or there are no elements that

can be joined with v (this condition is tested by universal quantification), then the domain is the singleton value [NULL], i.e., w becomes null. Otherwise each qualified element w of Y is joined with v.

- **outer-unnest**, $=\!\!\mu_p^{path}(X)$, is similar to $\mu_p^{path}(X)$, but if $x.path$ is empty for $x \in X$ or $p(x, y)$ is false for all $y \in x.path$, then the pair (x, NULL) appears in the output.

- **nest**, $\Gamma_{p/g}^{\oplus/e/f}(X)$, resembles $\Delta_p^{\oplus/e}(X)$; it combines the functionality of the nested relational operator, nest, with the functionality of reduce. In Eq. (O7), the nest operator uses the group-by function f: If two values v and w from a set X are dot-equal under f (i.e., $f(v) \doteq f(w)$), their images under e (i.e., $e(v)$ and $e(w)$) are grouped together in the same group. After a group is formed, it is reduced by the accumulator $\oplus$, and a pair of the group-by value, along with the result of the reduction of this group, is returned. Function g indicates which nulls to convert into zeros (i.e., into $Z_\oplus$). For example, the nest operation

$$\Gamma^{\cup/\lambda(d,e).e/\lambda(d,e).d}_{\lambda(d,e).\text{true}/\lambda(d,e).e}(X) = \cup\{\!\{ (d', \cup\{\!\{ e \,[\!]\, (d,e) \leftarrow X, e \neq \text{NULL}, d' \doteq d \}\!\}) \;[\!]\; d' \leftarrow \textstyle\prod_{\lambda(d,e).d}(X) \}\!\}$$

groups the input (which consists of pairs (d, e) of a department d and an employee e) by d and converts the null e's into empty sets. $\prod_{\lambda(d,e).d}(X)$ is a generalized projection that retrieves the d component of each pair (d, e) in X and removes the duplicate d's.

7. QUERY UNNESTING

There are many recent proposals for OODB query optimization that focus on unnesting nested queries (also known as query decorrelation) [Cluet and Moerkotte 1995a; 1995b; Claussen et al. 1997; Steenhagen et al. 1994]. Nested queries appear more often in OODB queries than in relational queries because OODB query languages allow complex expressions at any point in a query. In addition, OODB types are allowed to have attributes with collection values (i.e., nested collections), which lead naturally to nested queries. Current OODB systems typically evaluate nested queries in a nested-loop fashion, which does not leave many opportunities for optimization. Most unnesting techniques for OODB queries are actually based on similar techniques for relational queries [Kim 1982; Ganski and Wong 1987; Muralikrishna 1992]. For all but the trivial nested queries, these techniques require the use of outer-joins to prevent loss of data and grouping to accumulate the data and to remove the null values introduced by the outer-joins.

If considered in isolation, query unnesting itself does not result in performance improvement. Instead, it makes possible other optimizations that would not be possible otherwise. More specifically, without unnesting,

the only choice in evaluating nested queries is a naive nested-loop method: For each step of the outer query, all the steps of the inner query need to be executed. Query unnesting promotes all the operators of the inner query into the operators of the outer query. This operator mixing allows other optimization techniques to work, such as rearranging operators to minimize cost and the free movement of selection predicates between inner and outer operators, which enables operators to be more selective.

All early unnesting techniques were actually source-to-source transformations over SQL code, mostly due to the lack of a group-by operator to express grouping in the relational algebra. The absence of a formal theory in a form of an algebra to express these transformations resulted in a number of bugs (such as the infamous count bug [Ganski and Wong 1987]) which were eventually detected and corrected. Since OODB queries are far more complex than relational ones, it is even more crucial to express the unnesting transformations in a formal algebra that will allow us to prove the soundness and completeness of these transformations. The first work with that goal in mind was by Cluet and Moerkotte [1995a; 1995b], which covered many cases of nesting, including nested aggregate queries, and validated all the transformations. Their work was extended by Claussen et al. [1997] to include universal quantification. Cherniack and Zdonik [1998a; 1998b] have used formal methods to validate query unnesting transformations, but their work is a moderate generalization of Kim's query unnesting techniques for relational queries.

This section describes a query decorrelation algorithm that unnests all possible forms of query nesting in our language. Our unnesting algorithm extends previous work in two ways. First, it is not only sound, but it is also complete. That is, our framework is capable of removing any form of nesting. Second, our unnesting algorithm is more concise, more uniform, and more general than earlier work, mostly due to the use of the monoid comprehension calculus as an intermediate form for OODB queries. The monoid comprehension calculus treats operations over multiple collection types, aggregates, and quantifiers in a similar way, resulting in a uniform way of unnesting queries, regardless of the type of nesting. In fact, many forms of nested queries can be unnested by the normalization algorithm for monoid comprehensions. The remaining forms require the introduction of outer-joins and grouping. Our algorithm requires only two rewrite rules to unnest the queries that cannot be handled by the normalization algorithm.

7.1 Examples of Unnesting

As an example of how OODB queries are unnested in our model, consider the following nested comprehension calculus expression:

QUERY A:

$\cup\{\ \langle\ D{=}d,\ E{=}\cup\{\ e\ []\ e \leftarrow \text{Employees},\ e.\text{dno} = d.\text{dno}\ \}\ \rangle\ []\ d \leftarrow \text{Departments}\ \}$,

which, for each department, returns the employees of the department. The nesting inherent in this query can be avoided by using an outer-join combined with grouping [Muralikrishna 1992], as shown in Figure 8.A. The

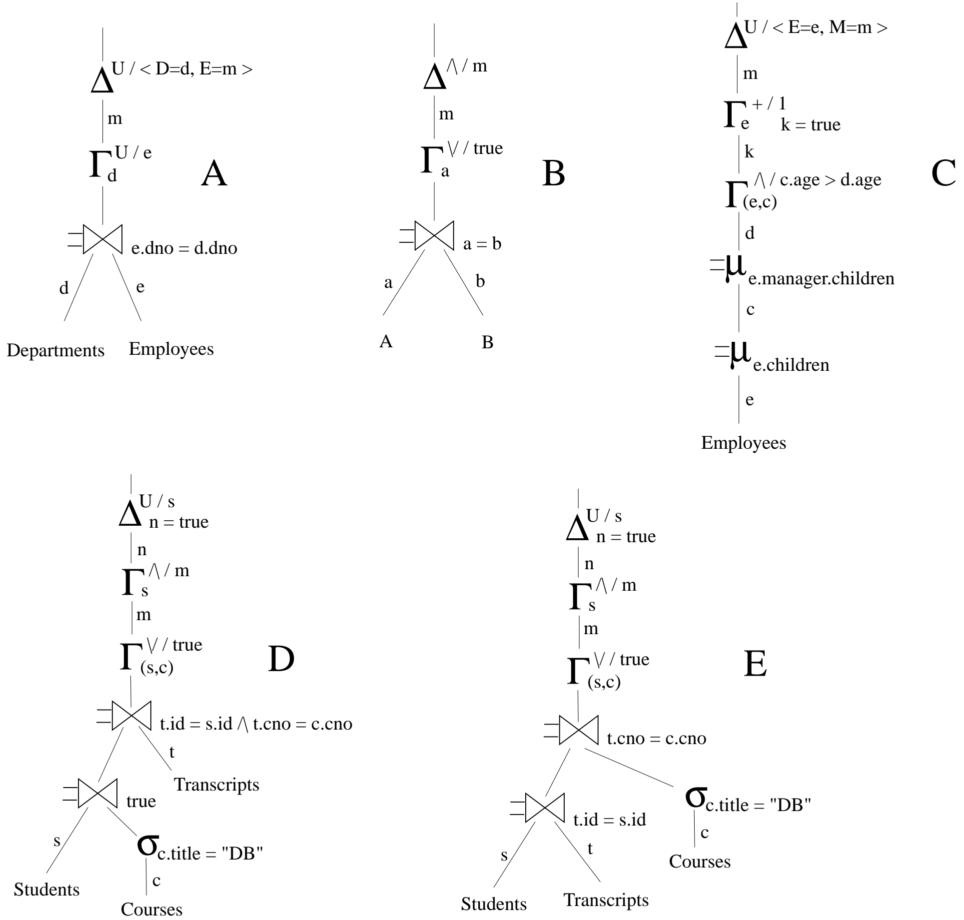

Fig. 8. The algebraic form of some OODB queries.

nest operator, $\Gamma_{\mathsf{d}}^{\cup/\mathsf{e}}$, groups the input by the range variable d, constructing a set of all e's that are associated with the value of d. This set becomes the value of the range variable m. That is, this nest operator reads a stream of tuples of type set(⟨ d: Department, e: Employee ⟩) and generates a stream of tuples of type set(⟨ d: Department, m: set(⟨ e: Employee ⟩) ⟩). The join before the nesting in Figure 8.A is a left outer-join: If there are no employees or the predicate e.dno=d.dno is false for all employees in a department d, then the result of the join is d associated with a null value, i.e., the value of the range variable e becomes null. The nest operator converts this null value into the empty set. That is, the outer-join introduces nulls and the nest operator converts nulls into zeros.

Another interesting example is the expression $A \subseteq B$, which is equivalent to $\forall a \in A : \exists b \in B : a = b$. It can be expressed in our calculus as follows:

QUERY B: $\wedge\{\ \vee\{$ true [] b ← B, a=b } [] a ← A }.

The inner comprehension, which captures the existential quantification, checks if there is at least one $b \in B$ with $a = b$ (it can also be written as ∨{ a=b [] b ← B }). The algebraic form of QUERY B after unnesting is shown in Figure 8.B. During the outer-join, if there is no a with b=a, then a will be matched with a null value. In that case, the nest operator will convert the null value into a false value, which is the zero element of ∨ . The reduction at the root tests whether all m's are true. If there is at least one false m, then there is at least one a in A that has no equal in B.

A more challenging example is the following double-nested comprehension:

```
QUERY C:
  ∪{ ⟨ E=e, M=+{ 1 [] c ← e.children,
                   ∧{ c.age > d.age [] d ← e.manager.children } } ⟩
     [] e ← Employees },
```

which, for each employee, computes the number of the employee's children who are older than the oldest child of the employee's manager. The algebraic form of QUERY C is shown in Figure 8.C. Here, instead of an outer-join, we use an outer-unnesting: the $=\!\mu_{\text{e.children}}$ operator, which introduces a new range variable, c, whose value is the unnesting of the set e.children. If this set is empty, then c becomes null, i.e., the employee e is padded with a null value. Since we have a double-nested query, we need to use two nested unnest-nest pairs. The top nest operator groups the input stream by each employee e, generating the number of children of e. Note that every operator in our algebra can be assigned a predicate (such as the predicate k=true in the top nest operation) to restrict the input data. The second nest operator groups the input stream by both e and c (i.e., by the tuple (e, c)) and, for each group, evaluates the predicate c.age > d.age and extends the output stream with an attribute k bound to the conjunction of all the predicates (indicated by the ∧ accumulator).

The following comprehension finds the students who have taken all database courses. It uses a predicate that has the same pattern as $A \subseteq B$:

```
QUERY D:
  ∪{ s [] s ← Students, ∧{ ∨{ true [] t ← Transcript, t.id=s.id, t.cno=c.cno }
                           [] c ← Courses, c.title = "DB" } }.
```

The algebraic form of this query is shown in Figure 8.D. This query can be evaluated more efficiently if we switch Courses with Transcripts, as shown in Figure 8.E. In that case, the resulting outer-joins are both assigned equality predicates, thus making them more efficient. Optimizations such as this justify query unnesting.

7.2 Our Query Unnesting Method

Figure 9 shows how our unnesting algorithm unnests QUERY D. Every comprehension is first normalized and then translated into an algebraic form consisting of regular joins, selections, unnests, and reductions—the last being the root of the algebraic form. This translation is straightforward. The outermost comprehension has one output (the result of the

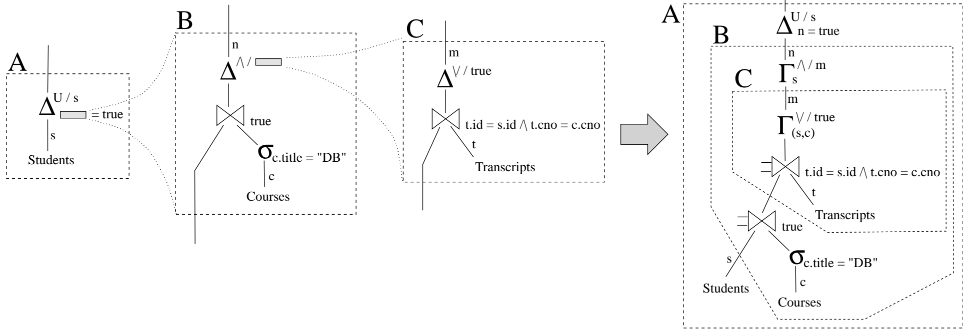

Fig. 9. Unnesting QUERY D.

reduction) and no input streams. For example, the dashed box A in Figure 9 represents the outer comprehension in QUERY D. The shaded box on the reduction represents a nested query. The algebraic form of an inner comprehension (i.e., a comprehension inside another comprehension) has one input stream and one output value: The input stream is the same stream of tuples as that of the operation in which this form is embedded and the output is the result of this form. For example, the dashed box B in Figure 9 represents the universally quantified comprehension (with accumulator $\wedge$): the input of this box is the same input stream as that of the reduction in box A (namely, the stream of employees), since box B is embedded in the predicate of the reduction. The output value of box B, n, is used in the predicate of the reduction in box A. Similarly, the existential quantification (with accumulator $\vee$) is translated into box C.

Our unnesting algorithm is simple: For each box that corresponds to a nested query (i.e., Boxes B and C in the current example), it converts reductions into nests (group-bys), and the joins and unnests that lay on the input-output path of the box into outer-joins and outer-unnests, respectively. At the same time, it embeds the resulting boxes at the points immediately before they are used. For example, box C will be embedded before the reduction in box B and the output value of box C will be used as the result of the innermost form. Similarly, box B will be embedded immediately before its output value is used in box A.

There is a very simple explanation why this algorithm is correct (a formal proof is given in Appendix A): A nested query is represented as a box, say box C, that consumes the same input stream as that of the embedding operation and computes a value m that is used in the embedding query (box B). If we want to splice this box onto the stream of the embedding query, we need to guarantee two things. First, box C should not block the input stream by removing tuples from the stream. This condition is achieved by converting the blocking joins into outer-joins and the blocking unnests into outer-unnests. Second, we need to extend the stream with the new value m of box C before it is used in the reduction of box B. This manipulation can be done by converting the reduction of box C into a nest, since the main difference between nest and reduce is that, while the reduce returns a

value (a reduction of a stream of values), nest embeds this value in the input stream. At the same time, the nest operator will convert null values to zeros so that the stream that comes from the output of the spliced box C will be exactly the same as it was before the splice. There are some important elements that we omitted here, but we will present them later when we describe the unnesting algorithm in detail. The most important factor is which nulls to convert to zeros each time: If we convert the null c's (i.e., the courses) to false in the second nest operation $\Gamma^{\vee/true}_{(s,\,c)}$ in the final unnested form in Figure 9, it will be too soon. This nest should convert null t's to false and the first nest should convert null c's to true. This difference is indicated by an extra parameter to the nest operator, which is not shown here.

7.3 The Query Unnesting Algorithm

The query unnesting algorithm is expressed in Figure 10 in terms of a set of rewrite rules. We assume that all comprehensions in a query have been put into the canonical form $\oplus\{\, e \;[]\; v_1 \leftarrow path_1, \ldots, v_n \leftarrow path_n, pred \,\}$ before this algorithm is applied. That is, all generator domains have been reduced to paths and all predicates have been collected to the right of the comprehension into *pred* by anding them together (*pred* is set to true if no predicate exists). The translation of a monoid comprehension $\oplus\{\, e \;[]\; \bar{r} \,\}$ is accomplished by using $[\![\oplus\{\, e \;[]\; \bar{r} \,\}]\!]^{u}_{w} E$. The comprehension $\oplus\{\, e \;[]\; \bar{r} \,\}$ is translated by compiling the qualifiers in $\bar{r}$ from left to right using the term E as a seed that grows at each step. That is, the term E is the algebraic tree derived at this point of compilation. The variables in w are all the variables encountered so far during the translation, and u are the variables that need to be converted to zeros during nesting if they are nulls. The situation $u = ()$ (i.e., when we have no variables in u) indicates that we are compiling an outermost comprehension (not a nested one). Rules (C4) through (C7) compile outermost comprehensions, while Rules (C8) through (C10) compile inner comprehensions. Rules (C11) and (C12) do the actual unnesting (here u can be of any value, including (), and $\otimes$ is not necessarily the same monoid as $\oplus$).

Rule (C4) is the first step of the unnesting algorithm: The comprehension must be the outermost comprehension; thus, the first generator must be over an extent X. In that case, the seed becomes a selection over X. The notation $p[v]$ specifies the part of the predicate p that refers to v exclusively, and does not contain any embedded comprehensions. The rest of the predicate is denoted by $p[\bar{v}]$ and satisfies $p[v] \wedge p[\bar{v}] = p$. This decomposition is used for pushing predicates to the appropriate operators. Rule (C5) is the last rule to be performed, after all generators have been compiled. Rule (C6) converts a generator over an extent (a variable) into a join. Here we split the predicate p into three parts: $p[v]$, which refers to v exclusively; $p[(w, v)]$, which refers to both w and v; and $p[\overline{(w, v)}]$ for the rest of the predicate. Note that all predicates with embedded comprehensions should

$$[\![\oplus\{\, e \,\|\, v \leftarrow X, \bar{r}, p \,\}]\!]^{()}_{()}\ \{()\} \;=\; [\![\oplus\{\, e \,\|\, \bar{r}, p[\bar{v}] \,\}]\!]^{()}_{v}\ (\sigma_{\lambda v.p[v]}(X)) \tag{C4}$$

$$[\![\oplus\{\, e \,\|\, p \,\}]\!]^{()}_{w}\ E \;=\; \Delta^{\oplus/\lambda w.e}_{\lambda w.p}(E) \tag{C5}$$

$$[\![\oplus\{\, e \,\|\, v \leftarrow X, \bar{r}, p \,\}]\!]^{()}_{w}\ E \;=\; [\![\oplus\{\, e \,\|\, \bar{r}, p[\overline{(w,v)}] \,\}]\!]^{()}_{(w,v)}\ (E \bowtie_{\lambda(w,v).p[(w,v)]} (\sigma_{\lambda v.p[v]}(X))) \tag{C6}$$

$$[\![\oplus\{\, e \,\|\, v \leftarrow path, \bar{r}, p \,\}]\!]^{()}_{w}\ E \;=\; [\![\oplus\{\, e \,\|\, \bar{r}, p[\bar{v}] \,\}]\!]^{()}_{(w,v)}\ (\mu^{\lambda w.path}_{\lambda(w,v).p[v]}(E)) \tag{C7}$$

$$[\![\oplus\{\, e \,\|\, p \,\}]\!]^{u}_{w}\ E \;=\; \Gamma^{\oplus/\lambda w.e/\lambda w.u}_{\lambda w.p/\lambda w.w\setminus u}(E) \tag{C8}$$

$$[\![\oplus\{\, e \,\|\, v \leftarrow X, \bar{r}, p \,\}]\!]^{u}_{w}\ E \;=\; [\![\oplus\{\, e \,\|\, \bar{r}, p[\overline{(w,v)}] \,\}]\!]^{u}_{(w,v)}\ (E \mathbin{=\!\!\bowtie}_{\lambda(w,v).p[(w,v)]} (\sigma_{\lambda v.p[v]}(X))) \tag{C9}$$

$$[\![\oplus\{\, e \,\|\, v \leftarrow path, \bar{r}, p \,\}]\!]^{u}_{w}\ E \;=\; [\![\oplus\{\, e \,\|\, \bar{r}, p[\bar{v}] \,\}]\!]^{u}_{(w,v)}\ ({=\!\!\mu}^{\lambda w.path}_{\lambda(w,v).p[v]}(E)) \tag{C10}$$

$$[\![\oplus\{\, e_1 \,\|\, \bar{s}, p(\otimes\{\, e_2 \,\|\, \bar{r} \,\}) \,\}]\!]^{u}_{w}\ E \;=\; [\![\oplus\{\, e_1 \,\|\, \bar{s}, p(v) \,\}]\!]^{u}_{(w,v)}\ ([\![\otimes\{\, e_2 \,\|\, \bar{r} \,\}]\!]^{w}_{w}\ E) \tag{C11}$$

$\otimes\{\, e_2 \,\|\, \bar{r} \,\}$ does not depend on the $\bar{s}$ generators

$$[\![\oplus\{\, f(\otimes\{\, e \,\|\, \bar{r} \,\}) \,\|\, p \,\}]\!]^{u}_{w}\ E \;=\; [\![\oplus\{\, f(v) \,\|\, p \,\}]\!]^{u}_{(w,v)}\ ([\![\otimes\{\, e \,\|\, \bar{r} \,\}]\!]^{w}_{w}\ E) \tag{C12}$$

Fig. 10. Rules for translating and unnesting comprehensions ($u \neq ()$ in Rules C8–C10).

be assigned to $p[\overline{(w,v)}]$, since otherwise the embedded comprehensions will not be unnested. Details of how predicates are split into parts will be given later in this section when we present the unnesting algorithm in a more operational form. Rule (C7) compiles generators with path domains into unnests.

Rules (C8) through (C10) apply to inner comprehensions and are similar to Rules (C5) through (C7), with the only difference that reductions become nests, joins become left outer-joins, and unnests become outer-unnests. The notation $w \setminus u$ indicates all the variables in w that do not appear in u. Those variables are the attributes to group-by (u are the attributes to convert into zeros when they are nulls). Rules (C11) and (C12) perform the actual unnesting. They do exactly what we have done in Figure 9 when we composed boxes: here the boxes are actually the results of the translation of the outer and inner comprehensions. Rule (C11) unnests a nested comprehension in the predicate p. It is applied as early as possible, that is, immediately, once the generators $\bar{s}$ do not affect the inner comprehension (i.e., when the free variables of the inner comprehensions do not depend on the generator variables in $\bar{s}$). Rule (C12) unnests a nested comprehension in the head of a comprehension. This unnesting is performed when all the generators of the outer comprehension have been reduced.

For example, QUERY C, Q_C, is compiled as follows:

$$
\begin{aligned}
Q_C &= [\![\cup\{\, \langle E = e, M = +\{\, 1 \,[\!]\, c \leftarrow e.\text{children}, \\
&\qquad \wedge\{\, c.\text{age} > d.\text{age} \,[\!]\, d \leftarrow e.\text{manager.children.true} \,\}\,\} \,[\!]\, e \leftarrow \text{Employees, true} \,\}]\!]^{()}_{()}\{()\} \\
&= [\![\cup\{\, \langle E = e, M = +\{\, 1 \,[\!]\, c \leftarrow e.\text{children}, \\
&\qquad\qquad \wedge\{\, c.\text{age} > d.\text{age} \,[\!]\, d \leftarrow e.\text{manager.children, true} \,\}\,\} \,[\!]\, \text{true} \,\}]\!]^{()}_{e}\text{Employees}
\end{aligned}
$$

from Rule (C4), if we ignore the selection over Employees (because it has a true predicate). From Rule (C12) we have,

$$
\begin{aligned}
&Q_C = [\![\cup\{\, \langle E = e, M = m\rangle \,[\!]\, \text{true} \,\}]\!]^{()}_{(e,\, m)} \\
&([\![+\{\, 1 \,[\!]\, c \leftarrow e.\text{children}, \wedge\{\, c.\text{age} > d.\text{age} [\!] d \leftarrow e.\text{manager.children, true} \,\}\,\}]\!]^{e}_{e}\text{Employees})
\end{aligned}
$$

to handle the inner + comprehension. From Rule (C10) we have,

$$
\begin{aligned}
Q_C &= [\![\cup\{\, \langle E = e, M = m\rangle \,[\!]\, \text{true} \,\}]\!]^{()}_{(e,\, m)} \\
&\qquad ([\![+\{\, 1 \,[\!]\, \wedge\{\, c.\text{age} > d.\text{age} \,[\!]\, d \leftarrow e.\text{manager.children, true} \,\}\,\}]\!]^{e}_{(e,\, c)} \\
&\qquad\qquad (=\!\mu^{\lambda e.e.\text{children}}_{\lambda(e,\, c).\text{true}}(\text{Employees})))
\end{aligned}
$$

to translate the e.children into an unnest:

$$
\begin{aligned}
Q_C &= [\![\cup\{\, \langle E = e, M = m\rangle \,[\!]\, \text{true} \,\}]\!]^{()}_{(e,\, m)} \\
&\qquad ([\![+\{\, 1 \,[\!]\, k \,\}]\!]^{e}_{((e,\, c),\, k)}([\![\wedge\{\, c.\text{age} > d.\text{age} \,[\!]\, d \leftarrow e.\text{manager.children, true} \,\}]\!]^{(e,\, c)}_{(e,\, c)} \\
&\qquad\qquad (=\!\mu^{\lambda e.e.\text{children}}_{\lambda(e,\, c).\text{true}}(\text{Employees}))))
\end{aligned}
$$

$$
\begin{aligned}
&= [\![\cup\{\, \langle E = e, M = m\rangle \,[\!]\, \text{true} \,\}]\!]^{()}_{(e,\, m)} \\
&\qquad ([\![+\{\, 1 \,[\!]\, k \,\}]\!]^{e}_{((e,\, c),\, k)}([\![\wedge\{\, c.\text{age} > d.\text{age} \,[\!]\, \text{true} \,\}]\!]^{(e,\, c)}_{((e,\, c),\, d)}(=\!\mu^{\lambda(e,\, c).e.\text{manager.children}}_{\lambda((e,\, c),\, d).\text{true}} \\
&\qquad\qquad (=\!\mu^{\lambda e.e.\text{children}}_{\lambda(e,\, c).\text{true}}(\text{Employees})))))
\end{aligned}
$$

$$
\begin{aligned}
&= [\![\cup\{\, \langle E = e, M = m\rangle \,[\!]\, \text{true} \,\}]\!]^{()}_{(e,\, m)} \\
&\qquad ([\![+\{\, 1 \,[\!]\, k \,\}]\!]^{e}_{((e,\, c),\, k)}(\Gamma^{\wedge/\lambda((e,\, c),\, d).c.\text{age}>d.\text{age}/\lambda((e,\, c),\, d).(e,\, c)}_{\lambda((e,\, c),\, d).\text{true}/\lambda((e,\, c),\, d).d}(=\!\mu^{\lambda(e,\, c).e.\text{manager.children}}_{\lambda((e,\, c),\, d).\text{true}} \\
&\qquad\qquad (=\!\mu^{\lambda e.e.\text{children}}_{\lambda(e,\, c).\text{true}}(\text{Employees})))))
\end{aligned}
$$

$$
\begin{aligned}
&= [\![\cup\{\, \langle E = e, M = m\rangle \,[\!]\, \text{true} \,\}]\!]^{()}_{(e,\, m)} \\
&\qquad (\Gamma^{+/\lambda((e,\, c),\, k).1/\lambda((e,\, c),\, k).e}_{\lambda((e,\, c),\, k).k/\lambda((e,\, c),\, k).(c,\, k)}(\Gamma^{\wedge/\lambda((e,\, c),\, d).c.\text{age}>d.\text{age}/\lambda((e,\, c),\, d).(e,\, c)}_{\lambda((e,\, c),\, d).\text{true}/\lambda((e,\, c),\, d).d}(=\!\mu^{\lambda(e,\, c).e.\text{manager.children}}_{\lambda((e,\, c),\, d).\text{true}} \\
&\qquad\qquad (=\!\mu^{\lambda e.e.\text{children}}_{\lambda(e,\, c).\text{true}}(\text{Employees})))))
\end{aligned}
$$

$$
\begin{aligned}
&= \Delta^{\cup/\lambda(e,\, m).(E=e,\, M=m)}_{\lambda(e,\, m).\text{true}} \\
&\qquad (\Gamma^{+/\lambda((e,\, c),\, k).1/\lambda((e,\, c),\, k).e}_{\lambda((e,\, c),\, k).k/\lambda((e,\, c),\, k).(c,\, k)}(\Gamma^{\wedge/\lambda((e,\, c),\, d).c.\text{age}>d.\text{age}/\lambda((e,\, c),\, d).(e,\, c)}_{\lambda((e,\, c),\, d).\text{true}/\lambda((e,\, c),\, d).d} \\
&\qquad\qquad (=\mu^{\lambda(e,\, c).e.\text{manager.children}}_{\lambda((e,\, c),\, d).\text{true}}(=\!\mu^{\lambda e.e.\text{children}}_{\lambda(e,\, c).\text{true}}(\text{Employees})))))
\end{aligned}
$$

(using Rules (C11), (C10), (C8), and (C5)), which is the algebraic form shown in Figure 8.C.

Figure 11 uses C-like pseudocode to present the unnesting rules in Figure 10 in a more operational way. The function call $T(e, u, w, E)$ corresponds to $[\![e]\!]^{u}_{w}E$. Thus, $T(e, (), (), \{()\})$ translates the term e into an

(1) Term T (Term e, Pattern u, Pattern w, Term E) {
(2) **if** $e = \oplus\{\, e_1 \,[\!]\, \overline{s},\, p \,\}$ and the outermost subterm $e' = \otimes\{\, e_2 \,[\!]\, \overline{r} \,\}$ of p
(if one exists) does not depend on $\overline{s}$;
(3) { Pattern v := new_variable();
(4) **return** $T(\, \oplus\{\, e_1 \,[\!]\, \overline{s},\, p[e'/v] \,\},\, u,\, (w,v),\, T(e', w, w, E)\,)$;
}
(5) **else if** $e = \oplus\{\, e_1 \,[\!]\, p \,\}$
(6) { Pattern v := new_variable();
(7) find the outermost term $e' = \otimes\{\, e_2 \,[\!]\, \overline{r} \,\}$ in term e_1;
(8) **if** found **return** $T(\, \oplus\{\, e_1[e'/v] \,[\!]\, p \,\},\, u, (w,v),\, T(e', w, w, E)\,)$;
(9) **else if** u=() **return** $\Delta^{\oplus/\lambda w.e_1}_{\lambda w.p}(E)$;
(10) **else return** $\Gamma^{\oplus/\lambda w.e_1/\lambda w.u}_{\lambda w.p/\lambda u \setminus v}(E)$;
}
(11) **else if** $e = \oplus\{\, e_1 \,[\!]\, v \leftarrow X,\, \overline{r},\, p \,\}$
(12) { $(p_1, p_2, p_3) \leftarrow$ split_predicate(p, w, v);
(13) **if** $u = ()$
(14) **if** $w = ()$ **return** $T(\, \oplus\{\, e_1 \,[\!]\, \overline{r},\, p_3 \,\},\, u,\, v,\, \sigma_{\lambda v.np}(E)\,)$;
else if X is a variable
(15) **return** $T(\, \oplus\{\, e_1 \,[\!]\, \overline{r},\, p_3 \,\},\, u,\, (w,v),\, E \bowtie_{\lambda(w,v).p_2}(\sigma_{\lambda v.p_1}(E))\,)$;
(16) **else return** $T(\, \oplus\{\, e_1 \,[\!]\, \overline{r},\, p_3 \,\},\, u,\, (w,v),\, \mu^{\lambda w.X}_{\lambda(w,v).p_1 \wedge p_2}(E)\,)$;
else if X is a variable
(17) **return** $T(\, \oplus\{\, e_1 \,[\!]\, \overline{r},\, p_3 \,\},\, u,\, (w,v),\, E \mathbin{=\!\!\bowtie}_{\lambda(w,v).p_2}(\sigma_{\lambda v.p_1}(E))\,)$;
(18) **else return** $T(\, \oplus\{\, e_1 \,[\!]\, \overline{r},\, p_3 \,\},\, u,\, (w,v),\, {=\!\!\mu}^{\lambda w.X}_{\lambda(w,v).p_1 \wedge p_2}(E)\,)$;
} }

Fig. 11. The query unnesting algorithm.

algebraic form. Terms are tree-like data structures that represent both calculus and algebraic terms. A pattern u is a binary tree structure that represents an empty pattern (), a variable v, or a pair of patterns (u_1, u_2). Central to this program is the function split_predicate(Term p, Pattern left, Pattern right) that returns a triple (p_1, p_2, p_3) of terms. Terms p, p_1, p_2, and p_3 are conjunctions of simple predicates with $p = p_1 \wedge p_2 \wedge p_3$. In particular, p_1 holds all simple predicates in p that refer to the pattern variables in right exclusively; p_2 holds all predicates that refer to the pattern variables in both left and right; and p_3 holds the rest of the predicates. Note that there are no predicates in p that refer to the pattern variables in left exclusively, since these predicates were assigned to the operators during previous steps. In addition, if a simple predicate is a comprehension (i.e., an embedded query in the predicate), then it is always assigned to p_3 in order to be unnested later by T. Line (2) in Figure 11 checks whether Rule () can be applied. This case holds when the outermost comprehension $e' = \otimes\{\, e_2 \,[\!]\, \overline{r} \,\}$ of the predicate p does not depend on the range variables of the generators $\overline{r}$. This case always applies when $\overline{r}$ is empty and there is an embedded comprehension in p. Lines (3) and (4) implement the right side of Rule (C11). The notation $p[e'/v]$ indicates that a copy of the term p is created, but with the subtree e' in p substituted for the variable v. Lines (5) through (7) check whether Rule (C12) can be applied and Line (8) applies this rule. Otherwise, either Rule (C5) or Rule

(C8) is applied, depending on whether $u = ()$ or not. Line (14) applies when both $u = ()$ and $w = ()$ (Rule (C4)). Otherwise we have four cases: whether $u = ()$ or not, whether the first generator X is over a path, or over a variable. Lines (15), (16), (17), and (18) capture Rules (C6), (C7), (C9), and (C10), respectively.

We can easily prove that the unnesting algorithm in Figure 10 is complete:

THEOREM 5. *The rules in Figure 10 unnest all nested comprehensions.*

PROOF. After normalization, the only places where we can find nested queries are the comprehension predicate and the head of the comprehension. These cases are handled by Rules (C11) and (C12), respectively. Even though Rule (C11) has a precondition, it will eventually be applied to unnest any nested query in a predicate. (In the worst case, it will be applied when all generators of the outer comprehension have been compiled by the other rules.) □

The soundness of our unnesting algorithm is a consequence of the following theorem:

THEOREM 6. *The rules in Figure 10 are meaning-preserving. That is,*

$$[\![\oplus\{\!\{ e \,[\!]\, \bar{r} \}\!\}]\!]^{()}_{()}\{()\} = \oplus\{\!\{ e \,[\!]\, \bar{r} \}\!\}.$$

The proof of this theorem is given in Appendix A (as Corollary 1 based on Theorem 8). Here we illustrate the validity of this theorem using an example. When we apply the rules in Figure 10, QUERY A becomes

$$[\![\cup\{\!\{ \langle D = d, M = \cup\{\!\{ e \,[\!]\, e \leftarrow \text{Employees}, e.\text{dno} = d.\text{dno} \}\!\}\rangle \,[\!]\, d \leftarrow \text{Department} \}\!\}]\!]^{()}_{()}\{()\}$$

$$= \Delta^{\cup/\lambda(d,m).(D=d,\,M=m)}_{\lambda(d,m).\text{true}}(\Gamma^{\cup/\lambda(d,e).e/\lambda(d,e).d}_{\lambda(d,e).\text{true}/\lambda(d,e).e}(\text{Department} \mathbin{=\!\bowtie}_{\lambda(d,e).e.\text{dno}=d.\text{dno}} \text{Employees}))$$

We prove that the algebraic form above is equivalent to the original comprehension. The tools that we use for this proof are the operator definitions in Figures 5 and 7 and the normalization algorithm. According to Rule (O4) in Figure 7, the above algebraic form is equal to

$$\cup\{\!\{ \langle D = d, M = m\rangle \,[\!]\, (d, m) \leftarrow (\Gamma^{\cup/\lambda(d,e).e/\lambda(d,e).d}_{\lambda(d,e).\text{true}/\lambda(d,e).e}(\text{Department} \mathbin{=\!\bowtie}_{\lambda(d,e).e.\text{dno}=d.\text{dno}} \text{Employees})) \}\!\}$$

$$\begin{aligned} = \cup\{\!\{ \langle D = d, M = m\rangle \,[\!]\, (d, m) \leftarrow \cup\{\!\{ (d, \cup\{\!\{ e' \,[\!]\, (d', e') \\ \leftarrow (\text{Department} \mathbin{=\!\bowtie}_{\lambda(d,e).e.\text{dno}=d.\text{dno}} \text{Employees}), e' \neq \text{NULL}, d \doteq d' \}\!\}) \\ [\!]\, d \leftarrow \textstyle\prod_{\lambda(d,e).d}(\text{Department} \mathbin{=\!\bowtie}_{\lambda(d,e).e.\text{dno}=d.\text{dno}} \text{Employees}) \}\!\} \}\!\} \end{aligned}$$

$$\begin{aligned} = \cup\{\!\{ \langle D = d, M = \cup\{\!\{ e' \,[\!]\, (d', e') \leftarrow (\text{Department} \mathbin{=\!\bowtie}_{\lambda(d,e).e.\text{dno}=d.\text{dno}} \text{Employees}), \\ e' \neq \text{NULL}, d \doteq d' \}\!\}\rangle \,[\!]\, d \leftarrow \textstyle\prod_{\lambda(d,e).d}(\text{Department} \mathbin{=\!\bowtie}_{\lambda(d,e).e.\text{dno}=d.\text{dno}} \text{Employees}) \}\!\} \end{aligned}$$

$$= \cup\{\!\{ \langle D = d, M = \cup\{\!\{ e' \,[\!]\, (d', e') \leftarrow \cup\{\!\{ (d, e) \,[\!]\, d \leftarrow \text{Department}, e \leftarrow F(d) \}\!\},$$

$$e' \neq \text{NULL}, d \doteq d' \}\rangle \,\|\, d \leftarrow \prod_{\lambda(d,e).d}(\uplus\{ (d, e) \,\|\, d \leftarrow \text{Department}, e \leftarrow F(d) \}) \},$$

where $F(d)$ is the right branch of the outer join in Rule (O5):

$$F(d) = \textbf{if } \wedge\{ e.\text{dno} \neq d.\text{dno} \,\|\, d \neq \text{NULL}, e \leftarrow \text{Employees} \} \textbf{ then } [\text{NULL}]$$
$$\textbf{else } \uplus\{ e \,\|\, e \leftarrow \text{Employees}, e.\text{dno} = d.\text{dno} \}.$$

But, according to Property (P2), $\prod_{\lambda(d,e).d}(\uplus\{ (d, e) \,\|\, d \leftarrow \text{Department}, e \leftarrow F(d) \})$ is equal to Department. Therefore, after normalization, we get

$$= \uplus\{ \langle D = d, M = \uplus\{ e' \,\|\, d' \leftarrow \text{Department}, e' \leftarrow F(d'), e' \neq \text{NULL}, d \doteq d' \}\rangle$$
$$\|\, d \leftarrow \text{Department} \}$$

$$= \uplus\{ \langle D = d, M = \uplus\{ e' \,\|\, d' \leftarrow \text{Department},$$
$$e' \leftarrow \uplus\{ e \,\|\, e \leftarrow \text{Employees}, e.\text{dno} = d.\text{dno} \}, d \doteq d' \}\rangle \,\|\, d \leftarrow \text{Department} \}$$

$$= \uplus\{ \langle D = d, M = \uplus\{ e \,\|\, d' \leftarrow \text{Department}, e \leftarrow \text{Employees}, e.\text{dno} = d.\text{dno}, d \doteq d' \}\rangle$$
$$\|\, d \leftarrow \text{Department} \}.$$

Finally, according to Property (P1), we can safely remove the generator $d' \leftarrow$ Department from the inner comprehension. The final form is

$$\uplus\{ \langle D = d, M = \uplus\{ e \,\|\, e \leftarrow \text{Employees}, e.\text{dno} = d.\text{dno} \}\rangle \,\|\, d \leftarrow \text{Department} \},$$

which is the original comprehension.

7.4 Simplifications

There is a large class of nested queries that can be further improved after unnesting. Consider for example the following OQL query that, for each department, finds the total salary of all the employees older than 30 in the department:

```
select distinct e.dno, sum(e.salary)
from Employees e
where e.age>30
group by e.dno.
```

Even though at a first glance this query does not seem to be nested, its translation to the monoid calculus in fact is nested:

$$\uplus\{ \langle E = e.\text{dno}, S = +\{ u.\text{salary} \,\|\, u \leftarrow \text{Employees}, u.\text{age} > 30, e.\text{dno} = u.\text{dno} \}\rangle$$
$$\|\, e \leftarrow \text{Employees}, e.\text{age} > 30 \}.$$

Our unnesting algorithm generates the algebraic form in Figure 12.A, but we prefer the form in Figure 12.B, which is more efficient.

This simplification can be accomplished easily with the help of the following transformation rule:

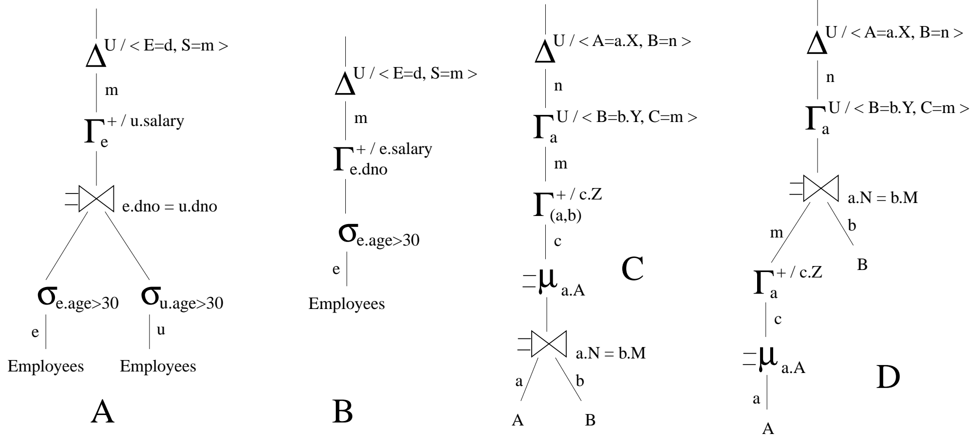

Fig. 12. Various simplifications.

$$\Gamma_{a}^{f(b)}(g(a) \bowtie_{a.M=b.M} g(b)) \rightarrow \Gamma_{a.M}^{f(a)}(g(a)),$$

where a and b are range variables in $g(a)$ and $g(b)$, respectively. Hence it does not make a difference whether this query is expressed as a group-by query or a nested query; both forms are optimized to the same efficient algebraic form.

As another simplification, we can use materialized views to handle some forms of uncorrelated nested queries [Cluet and Moerkotte 1995b]. For example, the following comprehension

$$\cup\{\, \langle C - c, D = \cup\{\, a.X + b.Y \,[]\, a \leftarrow A, b \leftarrow B, a.N = b.M \,\}\rangle \,[]\, c \leftarrow C \,\}$$

can be translated as follows:

$$\begin{aligned} &view := \cup\{\, a.X + b.Y \,[]\, a \leftarrow A, b \leftarrow B, a.N = b.M \,\} \\ &\cup\{\, \langle C = c, D = view \rangle \,[]\, c \leftarrow C \,\}, \end{aligned}$$

where the operator := creates a materialized view. All other cases of uncorrelated nested queries that cannot be handled with global views do not require any special treatment by our framework. For example, in the query

$$\begin{aligned} &\cup\{\, \langle A = a.X, B = \cup\{\, \langle A = a.X, B = +\{\, c.Z \,[]\, c \leftarrow a.A \,\}\rangle \,[]\, b \leftarrow B, a.N = b, M \,\}\rangle \\ &\qquad [] \, a \leftarrow A \,\}, \end{aligned}$$

the + comprehension has the same value for every a in A. Our unnesting algorithm will generate the algebraic form shown in Figure 12.C for this query, but it is an easy task for an optimizer to transform this algebraic form into the form shown in Figure 12.D (by pulling the outer-join up in the operator tree).

8. IMPLEMENTATION

We have already built a prototype ODMG database management system, called λ-DB, based on our framework [Fegaras et al. 2000]. Our system can handle most ODL declarations and process most OQL query forms. The λ-DB prototype is not ODMG-compliant. Instead it supports its own C++ binding, which provides a seamless integration between OQL and C++ with low impedance mismatch. It allows C++ variables to be used in queries and the results of queries to be passed back to C++ programs. Programs expressed in our C++ binding are compiled by a preprocessor that performs query optimization at compile time rather than runtime, as proposed by ODMG. In addition to compiled queries, λ-DB provides an interpreter that evaluates ad-hoc OQL queries at runtime.

The λ-DB evaluation engine is written in SDL (the SHORE Data Language) of the SHORE object management system [Carey et al. 1994], developed at the University of Wisconsin. ODL schemas are translated into SDL schemas in a straightforward way and are stored in the system catalog. The λ-DB OQL compiler is a C++ preprocessor that accepts a language called λ-OQL, which is C++ code with embedded DML commands, to perform transactions, queries, updates, etc. The preprocessor translates λ-OQL programs into C++ code that contains calls to the λ-DB evaluation engine. We also provide a visual query formulation interface, called VOODOO [Fegaras 1999b], and a translator from visual queries to OQL text, which can be sent to the λ-DB OQL interpreter for evaluation.

Our OQL optimizer is expressed in a very powerful optimizer specification language, called OPTL, and is implemented in a flexible optimization framework, called OPTGEN, which extends our earlier work on optimizer generators [Fegaras et al. 1993]. OPTL is a language for specifying query optimizers that captures a large portion of the optimizer specification information in a declarative manner. It extends C++ with a number of term-manipulation constructs and with a rule language for specifying query transformations. OPTGEN is a C++ preprocessor that maps OPTL specification into C++ code.

The λ-DB OQL optimizer proceeds in eight phases: (1) parsing of OQL queries and translation from OQL into calculus; (2) type checking; (3) normalization; (4) translation into algebra and query unnesting; (5) join permutation; (6) physical plan generation; (7) generation of intermediate evaluation code; and (8) translation from intermediate code into C++ code, or interpretation of the intermediate code.

Our unnesting algorithm, which is based on the framework described in earlier sections, turned out to be very easy to implement (it is only 160 lines of C++ code). In addition to query unnesting, λ-DB converts path expressions into pointer joins between class extents (when possible) and uses information about 1:N class relationships to convert unnests into pointer joins between extents (by using the inverse relationship).

Our query optimizer translates monoid comprehensions into algebraic operators, which, in turn, are mapped into physical execution algorithms like those found in relational database systems. This translation is done in stages: queries in our framework are first translated into monoid comprehensions, which serve as an intermediate form, and then are translated into a version of the nested relational algebra that supports aggregation, quantification, outer-joins, and outer-unnests. At the end, the algebraic terms are translated into execution plans.

A very important task of any query optimizer is finding a good order to evaluate the query operations. In the context of relational databases, this task is known as join-ordering. The λ-DB system uses a polynomial-time heuristic algorithm, called GOO [Fegaras 1998a], that generates a "good quality" order of the monoid algebra operators in query algebraic form. GOO is a bottom-up greedy algorithm that always performs the most profitable operations first. The measure of profit is the size of the intermediate result, but it can be easily modified to use real cost functions. GOO is based on query graphs and takes into account both the output sizes of the results constructed in earlier steps and the predicate selectivities. It generates bushy join trees that have very low intermediate results.

After the best evaluation order is derived, the algebraic form is mapped into an evaluation plan by a rule-based rewriting system. Our system considers the available access paths (indexes) and the available physical algorithms to generate different plans. During this phase, all alternative plans (for the derived order of operators) are generated and priced and the best plan is selected. Finally, each evaluation plan is translated into an intermediate evaluation code that reflects the signatures of the evaluation algorithms used in λ-DB. This code can then be translated straightforwardly into C++ or interpreted by the λ-DB interpreter.

The λ-DB evaluation engine is built on top of SHORE [Carey et al. 1994]. We use the SDL layer of SHORE exclusively in our implementation, because we believe it is more resilient to change and is easier to use than the Shore Storage Manager. An alternative is to write our own value-added server on top of the storage manager. Much of our effort has been devoted to make the implementation of the evaluation engine as simple as possible without sacrificing performance. This required a very careful design. For example, one of our design requirements was to put all the evaluation algorithms in a library, so that the query evaluation code would consist of calls to these algorithms. This approach sounds obvious enough and simple to implement, but it required some special techniques (described below) borrowed from the area of functional programming.

Algebraic operators, such as $R \bowtie_{R.A=S.B} S$, are intrinsically higher order. That is, they are parameterized by pieces of code (i.e., functions) that specify some of the operation details. For the previous join, the piece of code is the predicate $R.A = S.B$. Most commercial DBMSs evaluate query plans using a plan interpreter. Very few systems actually compile plans into code, and when they do it is only for procedures written in a database

language. If it were to be implemented in C++, the join operator could be specified as

```
Relation join ( Relation x, Relation y, bool (pred) (tuple,tuple) ),
```

where pred is the address of the predicate function. The same holds for any evaluation plan operator, such as the nested-loop join operator. If a query is to be compiled into execution code with no runtime interpretation of predicates, it should make use of higher-order evaluation algorithms. The alternative is to define the evaluation algorithms as kinds of macros to be macroexpanded and individually tailored for each different instance of the evaluation operator in the query plan. For example, the nested-loop join would have to be a function with a "hole" in its body, to be filled with the predicate code. Then the entire code, the evaluation function and the predicate, would be generated and inserted inside the query code. This approach is clumsy and makes the development and extension of the evaluation algorithms very tedious and error-prone.

One very important issue when writing evaluation algorithms is stream-based processing of data (sometimes called iterator-based processing or pipelining). Whenever possible, it is highly desirable to avoid materializing the intermediate data in a secondary storage. The λ-DB system pipelines all evaluation algorithms. Instead of using threads to implement pipelining, as traditionally done in most systems, we developed a special technique borrowed from the area of lazy functional languages. Each of our evaluation algorithms is written in such a way that it returns a piece of data (one tuple) as soon as it constructs one. To retrieve all the tuples, the algorithm must be called multiple times. For example, the pipeline version of the nested-loop join will have the signature

```
tuple nested_loop (Stream sx, Stream sy, bool (pred) (tuple,tuple)),
```

where Stream is a stream of tuples equipped with standard operations, such as first and next. In contrast to the standard nested-loop algorithm, this algorithm exits when it finds the first qualified tuple (that satisfies pred). To pipeline an algorithm, we construct a *suspended stream*, which is a structure with just one component: an embedded function, which, when invoked, calls the evaluation algorithm to construct one tuple. For example, to pipeline our nested-loop algorithm, we construct a suspended stream whose embedded function, F, is defined as tuple F ( ) return nested_loop(s1,s2,pred); where s1 and s2 are the streams that correspond to the join inputs and pred is the address of the predicate function. When a tuple is needed from a suspended stream, its embedded function is called with no arguments to return the next tuple. This approach is a clean and efficient way of implementing pipelining. This type of evaluation resembles lazy evaluation in functional programming languages. Here, though, we provide an explicit control over the lazy execution, which gives a better handle on controlling the data materialization in a secondary storage.

The intermediate evaluation code, which is generated from a query evaluation plan, is a purely functional program that consists of calls to the

evaluation algorithms. The functional parameters of the calls (such as the predicate function of the nested-loop join) are represented as anonymous functions (lambda abstractions). If the intermediate code is compiled into C++, the lambda abstractions are translated into named C++ functions by a simple defunctionalization process. If the intermediate code is interpreted, each call to an evaluation algorithm is executed without much overhead because function addresses are stored into a vector and calls to functions are dispatched in constant time by retrieving the function directly from the vector. Lambda abstractions are interpreted by using the address of the interpreter itself as a parameter and by pushing the body of the lambda abstraction into a special stack. This form of dispatching makes the interpreter very compact (400 lines of C++ code only) and very fast.

Currently, the λ-DB evaluation engine supports table scan of a class extent, index scan, external sorting, block nested-loop join, indexed nested-loop join, sort-merge join, pointer join, unnesting, nesting (group-by), and reduction (to evaluate aggregations and quantifications). The λ-DB OODBMS is available as open source software at

`http://lambda.uta.edu/lambda-DB.html`.

9. PERFORMANCE EVALUATION

To evaluate the performance of our query-unnesting algorithm, we used the λ-DB OODBMS to optimize and run 13 nested OQL queries against a sample database. The database schema is given in Figure 1 and the OQL queries are shown in Figure 13. The experiments were performed under the following assumptions:

First, the queries not unnested by our decorrelation algorithm were evaluated in a naive nested loop fashion: for each step of the outer query, the entire inner query was evaluated. The query-unnesting algorithm introduces outer-joins and outer-unnests to preserve the tuples of the outer query that are not joined with any tuple from the inner query. The outer-joins introduced by query unnesting are evaluated using outer-block-nested-loop joins. An outer-block-nested-loop join uses a fixed-size memory buffer to store a large number of tuples from the outer stream, thus reducing the number of inner stream scans. Similarly, a left outer-join is implemented by marking the joined tuples in the buffer and then by checking for unmarked tuples before the buffer is replaced. We used a buffer of 1,000 tuples in our tests, which means that the inner stream of a join was scanned about 1,000 fewer times than the naive nested-loop implementation.

Second, in contrast to joins, there are very few ways to evaluate the unnest operator. In fact, λ-DB provides only one implementation for the unnest operator: a naive nested loop between each tuple and its collection component being unnested. This case is also true for the outer-unnests. Thus, if there are no joins in a query, then the only improvement introduced by query unnesting would come from moving predicates between the inner and outer queries, which may result in more selective unnest operations.

1. select x: e.name, y: (select c.name from c in e.teaches) from e in Instructors
2. select x: e.name, y: count(e.dept.instructors) from e in Instructors
3. select x: e.name, y: (select x: c.name, y: count(c.has_prerequisites)
   from c in e.teaches) from e in Instructors
4. select x: dn, y: count(partition) from e in Instructors group by dn: e.rank
5. select d.name, c: count(select * from e in d.instructors where e.rank=“professor”)
   from d in Departments order by count(select * from e in d.instructors
   where e.rank=“professor”)
6. select e.name, c: count(e.teaches) from e in Instructors where count(e.teaches)≥4
7. select x, y, c: count(c) from e in Instructors, c in e.teaches
   group by x: e.ssn, y: c.name having x>60 and y>“CSE5330”
8. select x: x, y: count(e) from e in Instructors group by x: count(e.teaches)
   having x>0
9. select x, y, c: count(e) from e in Instructors group by x: count(e.teaches),
   y: (exists c in e.teaches: c.name=“CSE5330”) having x>0
10. select x: x, y: count(e) from d in Departments, e in d.instructors
    group by x: count(e.teaches) having x>0
11. sum(select sum(select e.salary from e in d.instructors) from d in Departments)
12. select e.name, X: (select x from c in e.teaches
    group by x: count(c.has_prerequisites)) from d in Departments, e in d.instructors
13. select e from e in Instructors where for all c in e.teaches:
    exists d in c.is_prerequisite_for: d.name = “CSE5330”

Fig. 13. The benchmark queries.

Since “predicate move around” has already been shown to be effective by others [Levy et al. 1994; Mumick et al. 1990], we decided to do benchmarks over queries with joins by translating outer-joins into outer-block-nested loops. This modification was mainly accomplished by mapping path expressions such as e.dept.instructors in Query 2 into pointer joins between class extents. This optimization, known as materialization of path expressions, is very effective when there are multiple paths with the same prefix. It is also very effective when combined with query unnesting, since query unnesting promotes the pointer joins of the inner queries to the outer query, thus allowing mixing the inner and outer operators.

Third, the group-by queries (queries 7, 8, 9, and 10) were optimized into group-by operations only when the queries were unnested. This step was necessary because the group-by queries in λ-DB are first translated into nested queries (as described in Section 3) and then these nested queries are translated into group-by operators as described in Section 7.4 (using the simplification transformations for query unnesting). We made this transformation in two steps instead of one, since the original group-by query and

Table II. Performance Evaluation of the Query Unnesting Algorithm

Departments/ Instructors/Courses	10/100/50	20/200/100	30/300/150	50/500/200
Query 1:	0.03 / 0.09	0.04 / 0.46	0.14 / 0.99	0.17 / 2.57
Query 2:	0.41 / 0.19	0.93 / 1.08	1.51 / 2.05	2.59 / 6.96
Query 3:	0.05 / 0.19	0.13 / 0.59	0.18 / 1.29	0.43 / 2.70
Query 4:	0.01 / 0.27	0.10 / 1.02	0.03 / 2.42	0.05 / 6.65
Query 5:	0.01 / 0.24	0.03 / 0.78	0.04 / 1.96	0.08 / 7.36
Query 6:	0.01 / 0.15	0.07 / 0.56	0.10 / 1.16	0.20 / 2.53
Query 7:	0.00 / 0.16	0.03 / 2.83	0.07 / 5.56	0.05 / 10.94
Query 8:	0.03 / 10.69	0.06 / 76.44	0.15 / 255.92	0.18 / 745.03
Query 9:	0.04 / 19.53	0.09 / 168.63	0.26 / 590.28	0.42 / *
Query 10:	0.07 / 11.94	0.16 / 92.09	0.24 / 301.30	0.39 / *
Query 11:	0.02 / 0.06	0.07 / 0.12	0.09 / 0.21	0.13 / 0.70
Query 12:	0.19 / 0.17	0.17 / 0.61	0.30 / 1.14	0.74 / 2.96
Query 13:	0.05 / 0.26	0.11 / 1.30	0.16 / 2.59	0.32 / 5.81

the translated nested query are semantically equivalent, and the resulting evaluation plan should not depend on the way the user chooses to express the query. Of course, other systems may choose to do this translation in one step, in which case there would be no differences in measurement in the nested and unnested versions of queries 7 through 10, since both result in the same plan.

Finally, both nested and unnested queries were optimized extensively by the λ-DB optimizer, and the evaluation plans were evaluated in a stream-based fashion.

The test database was created in four different sizes, as indicated by the cardinalities of the class extents for Departments, Instructors, and Courses in Table II. The benchmarks were run on a lightly loaded Linux workstation with a single 800MHz Pentium III, 256MB memory, and an ultra–160/m scsi 10000rpm disk. The buffer pool size of the SHORE client was set to 20MB. The numbers shown in Table II indicate processor time in seconds, calculated using the unix function clock. Each entry in the table has two values, n/m, to indicate the execution time of the query in seconds with (n) and without (m) query unnesting. A star indicates incomplete execution due to a SHORE server crash. From the measurements above, it is clear that query unnesting offers a significant performance improvement for the selected queries.

10. RELATED WORK

There are many proposals for object query algebras [Leung et al. 1993; Danforth and Valduriez 1992; Cluet and Delobel 1992; Beeri and Kornatzky 1990; Pistor and Traunmueller 1986]. In contrast to our algebra, these algebras support multiple bulk operators. But as we have demonstrated in this article, we get enough expressive power with just one operator, namely the monoid homomorphism. Supporting a small number of operators is highly desirable, since the more bulk operations an algebra

supports, the more transformation rules it needs and, therefore, the harder the optimization task becomes.

Our framework is based on monoid homomorphisms, which were first introduced by Breazu-Tannen et al. [1992a; 1992b] and Breazu-Tannen and Subrahmanyam [1991] as an effective way to capture database queries. Their form of monoid homomorphism (also called *structural recursion* over the union presentation—SRU) is more expressive than ours. However, operations of the SRU form require the validation of the associativity, commutativity, and idempotence properties of the monoid associated with the output of this operation. These properties are hard to check by compiler [Breazu-Tannen and Subrahmanyam 1991], which makes the SRU operation impractical. Breazu-Tannen and Subrahmanyam first recognized that there are some special cases where these conditions are automatically satisfied, such as the ext(f) operation (equivalent to hom[$\oplus$, $\oplus$](f) for a collection monoid $\oplus$). In our view, SRU is too expressive, since inconsistent programs cannot always be detected in that form. Moreover, the SRU operator can capture nonpolynomial operations, such as the powerset, which complicates query optimization. In fact, to our knowledge, there is no normalization algorithm for SRU forms in general (i.e., SRU forms cannot be put in canonical form). On the other hand, ext(f) is not expressive enough, since it does not capture operations that involve different collection types and cannot express predicates and aggregates. Our monoid comprehension is the most expressive subset of SRU proposed so far, in which inconsistencies can always be detected at compile time and, more importantly, where all programs can be put in canonical form.

Monad comprehensions were first introduced by Wadler [1990] as a generalization of list comprehensions. Monoid comprehensions are related to monad comprehensions but are considerably more expressive. In particular, monoid comprehensions can mix inputs from different collection types and may return output of a different type. This mixing is not possible for monad comprehensions, since they restrict the inputs and the output of a comprehension to be of the same type. Monad comprehensions were first proposed as a convenient database language by Trinder and Wadler [1989]; Trinder [1991]; and Chan and Trinder [1994], who also presented many algebraic transformations over these forms, as well as methods for converting comprehensions into joins. The monad comprehension syntax was also adopted by Buneman et al. [1994] as an alternative syntax to monoid homomorphisms. The comprehension syntax was used to capture operations that involve collections of the same type, while structural recursion was used for expressing the rest of the operations (such as converting one collection type to another, predicates, and aggregates).

In a previous work [Fegaras 1993], we used an alternative representation for lists, called insert-representation, to construct list values. This method is commonly used in functional programming, mainly because it uses only two constructors, cons and nil, to construct list values, instead of using three: append, singleton, and nil. We found that the append-representation

is superior to the insert-representation when adding extra properties to the append constructor, such as commutativity and idempotence, to define bags and sets.

Our normalization algorithm is influenced by Wong's work on normalization of monad comprehensions [Wong 1993; 1994]. He presented some very powerful rules for flattening nested comprehensions into canonical comprehension forms whose generators are over simple paths. These canonical forms are equivalent to our canonical forms for monoid comprehensions. His work, though, does not address query unnesting for complex queries (in which the embedded query is part of the predicate or the comprehension head), which cannot be unnested without using outer-joins and grouping.

There is an increasing number of proposals on OODB query optimization. Some of them are focused on handling nested collections [Ozsoyoglu and Wang 1992; Colby 1989], others on converting path expressions into joins [Kemper and Moerkotte 1990; Cluet and Delobel 1992], others on unnesting nested queries [Cluet and Moerkotte 1995b; 1995a], while still others are focused on handling encapsulation and methods [Daniels et al. 1991].

Query decorrelation was first introduced in the context of relational queries [Kim 1982; Ganski and Wong 1987; Muralikrishna 1992], mostly in the form of source-to-source transformations. There are a number of proposals lately that instead of rewriting correlated queries into more efficient queries or algebraic forms, offer new evaluation algorithms that perform better than the default nested loop evaluation of embedded queries. Examples of such methods include magic decorrelation [Mumick et al. 1990], which promotes predicates in the inner loop of the nested loop evaluation of correlated queries; and predicate move around [Levy et al. 1994], which moves predicates around the query graph. Our decorrelation algorithm not only moves predicates to the right place, it also intermixes operators between outer and inner queries, thus supporting more possibilities for join permutations—which may lead to better plans.

Our query unnesting algorithm is influenced by the work of Cluet and Moerkotte [1995a; 1995b], which covered many cases of nesting in OODBs, including nested aggregate queries and, more importantly, validated all the transformations. Their work was extended by Claussen et al. [1997] to include universal quantification. Our work proposes a rewriting system for complete unnesting, while their work considers algebraic equalities for some forms of unnesting.

The work of Lin and Ozsoyoglu [1996] addresses nested queries in a different way. Nested OQL queries are first translated into generalized path expressions that have enough expressive power to capture outer-joins and grouping. This translation may be very inefficient at the beginning due to a large number of redundant joins, but an optimization phase exists where path expressions are translated into joins and redundancies are eliminated. But no proof was provided to show that their two-phase unnesting method avoids redundancies in all forms and generates the same quality of plans as those of other unnesting techniques. Furthermore, no proof was given about the correctness and completeness of their method.

However, we believe that when these issues are addressed properly, their technique may turn out to be very effective and practical.

Another approach to query unnesting—which has the same goals as ours—is that of Cherniack and Zdonik [1998a; 1998b]. In contrast to our approach, they used an automatic theorem prover to prove the soundness of their unnesting rewrite rules, while we proved our transformations by hand. Their work is a moderate generalization of Kim's query unnesting techniques for relational queries, while ours covers all forms of query nesting for complex OODB queries. The use of a theorem prover is highly desirable for extensible systems, since it makes the query optimizer very flexible. In particular, with the help of a theorem prover, an optimizer does not require validation in a form of a formal proof each time a new algebraic operator or a new rewrite rule is introduced.

11. CURRENT AND FUTURE WORK

Even though the basic framework for query unnesting has already been developed, there are some issues that we plan to address in future research. The most important is query unnesting in the presence of methods. The problem is not in handling a method invocation that contains a query in one of its arguments; this case can be handled effectively by our unnesting framework.The problem is when the method body itself contains queries. There are a number of proposals for handling this case. One is breaking the encapsulation and unfolding the method body in the place of the call. This approach does not apply to recursive methods and does not support separate compilation very well, since it requires query recompilation when a method changes. Another approach is attaching user-defined optimization rules to every class or method, which requires a substantial effort from programmers. Yet another approach is considering a method as a view and materializing it into a table (known as memoization in programming languages). We are planning to investigate all these approaches more thoroughly and to choose the one that fits our unnesting algorithm best.

We are currently developing a formal framework for optimizing object-oriented queries in the presence of side effects [Fegaras 1999a]. These queries may contain object updates at any place and in any form. We have proposed a language extension to the monoid comprehension calculus to express these object-oriented features and given a formal meaning to these extensions. Our method is based on denotational semantics, which is often used to give a formal meaning to imperative programming languages. The semantics of our language extensions is expressed in terms of our monoid calculus, without the need for any fundamental change to our basic framework. Our initial results suggest that our method not only maintains referential transparency, which allows us to do meaningful query optimization, but it is also practical for optimizing OODB queries, since it allows the same optimization techniques applied to regular queries to be used with minimal changes for OODB queries.

12. CONCLUSION

We have presented a uniform calculus based on comprehensions that captures many advanced features found in modern object-oriented and object-relational query languages. We showed that this calculus is easy to manipulate by presenting a normalization algorithm that unnests many forms of nested comprehensions. The main contribution of this article is an effective framework for optimizing OODB queries. Queries in our framework are first translated into comprehensions and then into a lower-level algebra that reflects many relational DBMS physical algorithms. The algorithm that translates the calculus into the algebra performs query decorrelation by unnesting any form of nested comprehensions. Finally, we reported on an implementation of an ODMG-based DBMS that depends on our optimization framework.

APPENDIX

A. PROOFS OF THEOREMS

THEOREM 3. *The normalization rules in Figure 4 are meaning-preserving.*

PROOF. We prove the correctness of Rule (N8) only, since it is the most difficult to prove (there are similar proofs for the other rules). Even though the proof of Rule (N8) assumes the correctness of the other rules, there is no circularity because the other rules can be proved without a reference to Rule (N8). We define $\mathcal{F}_{\oplus}[\![\bar{q}]\!][\![e]\!]$, where $\bar{q}$ is a sequence of qualifiers and e is a term in the calculus, as follows:

$$\mathcal{F}_{\oplus}[\![\,]\!][\![e]\!] = e \tag{F1}$$

$$\mathcal{F}_{\oplus}[\![x \leftarrow u, \bar{q}]\!][\![e]\!] = \text{hom}[\otimes, \oplus](\lambda x.\mathcal{F}_{\oplus}[\![\bar{q}]\!][\![e]\!])u \tag{F2}$$

$$\mathcal{F}_{\oplus}[\![\text{pred}, \bar{q}]\!][\![e]\!] = \textbf{if } \text{pred} \textbf{ then } \mathcal{F}_{\oplus}[\![\bar{q}]\!][\![e]\!] \textbf{ else } \mathcal{Z}_{\oplus} \tag{F3}$$

where $\otimes$ in Eq. (F2) is the collection monoid associated with u. Note that $\mathcal{F}$ is not a function over values but a function over terms in the calculus. It gets a sequence of qualifiers and a term as input and returns a term as output. It can be easily proved using these equations and Definition 3:

$$\oplus\{\, e \mid \bar{q}, \bar{r} \,\} = \mathcal{F}_{\oplus}[\![\bar{q}]\!][\![\oplus\{\, e \mid \bar{r} \,\}]\!] \tag{W1}$$

First, we prove that

$$\oplus\{\, e \mid v \leftarrow \otimes\{\, e' \mid \bar{r} \,\}, \bar{s} \,\} = \oplus\{\, e \mid \bar{r}, v \equiv e', \bar{s} \,\} \tag{W2}$$

using induction over the number of generators in $\bar{r}$ and using structural induction over the domain of the first generator in $\bar{r}$. If $\bar{r}$ does not contain any generator (i.e., it contains only predicates), then the equation above is true. If the first qualifier of $\bar{r}$ is a predicate, then we can move the predicate at the end of $\bar{r}$. Let $\bar{r} = w \leftarrow u, \bar{t}$, where t has n generators. We assume

that Eq. (W2) is true for a n generators in $\bar{r}$ (induction hypothesis), and we prove it for $\bar{r}$ (induction step). Let u be a collection associated with the monoid $\odot$. Then u is $\mathcal{Z}_{\odot}$, or it is $\mathcal{U}_{\odot}(a)$ for some a, or it is $X \odot Y$ for some X and Y. If $u = \mathcal{Z}_{\odot}$, then both sides of the equation are equal to $\mathcal{Z}_{\oplus}$. If $u = \mathcal{U}_{\odot}(a)$, then

$$\begin{aligned}\oplus\{ e \,[\!]\, v \leftarrow \otimes\{ e' \,[\!]\, w \leftarrow \mathcal{U}_{\odot}(a), \bar{t}\,\}, \bar{s}\,\} &= \oplus\{ e \,[\!]\, v \leftarrow \otimes\{ e' \,[\!]\, w \equiv a, \bar{t}\,\}, \bar{s}\,\} \\ &= \oplus\{ e \,[\!]\, v \leftarrow \otimes\{ e'[a/w] \,[\!]\, \bar{t}[a/w]\}, \bar{s}\,\} \\ &= \oplus\{ e \,[\!]\, \bar{t}[a/w], v \equiv e'[a/w], \bar{s}\,\} \\ &= \oplus\{ e \,[\!]\, w \leftarrow \mathcal{U}_{\odot}(a), \bar{t}, v \equiv e', \bar{s}\,\}\end{aligned}$$

from Rule (N6), Rule (N1), and from the induction hypothesis (assuming that there are no name conflicts between variables). We assume that Eq. (5) is true for X and Y and prove it for $u = X \odot Y$:

$$\begin{aligned}&\oplus\{ e \,[\!]\, v \leftarrow \otimes\{ e' \,[\!]\, w \leftarrow (X \odot Y), \bar{t}\,\}, \bar{s}\,\} \\ &= \oplus\{ e \,[\!]\, v \leftarrow (\otimes\{ e' \,[\!]\, w \leftarrow X, \bar{t}\,\}) \otimes (\otimes\{ e' \,[\!]\, w \leftarrow X, \bar{t}\,\}), \bar{s}\,\} \\ &= (\oplus\{ e \,[\!]\, v \leftarrow \otimes\{ e' \,[\!]\, w \leftarrow X, \bar{t}\,\}, \bar{s}\,\}) \oplus (\oplus\{ e \,[\!]\, v \leftarrow \otimes\{ e' \,[\!]\, w \leftarrow Y, \bar{t}\,\}, \bar{s}\,\}) \\ &= (\oplus\{ e \,[\!]\, w \leftarrow X, \bar{t}, v \equiv e', \bar{s}\,\}) \oplus (\oplus\{ e \,[\!]\, w \leftarrow Y, \bar{t}, v \equiv e', \bar{s}\,\}) \\ &= \oplus\{ e \,[\!]\, w \leftarrow (X \odot Y), \bar{t}, v \equiv e', \bar{s}\,\}\end{aligned}$$

from Rule (N7) and from induction hypothesis. Therefore, the left part of Rule (N8) becomes

$$\begin{aligned}\oplus\{ e \,[\!]\, \bar{q}, v \leftarrow \otimes\{ e' \,[\!]\, \bar{r}\,\}, \bar{s}\,\} &= \mathcal{F}_{\oplus}[\![\bar{q}]\!][\![\oplus\{ e \,[\!]\, v \leftarrow \otimes\{ e' \,[\!]\, \bar{r}\,\}, \bar{s}\,\}]\!] \\ &= \mathcal{F}_{\oplus}[\![\bar{q}]\!][\![\oplus\{ e \,[\!]\, \bar{r}, v \equiv e', \bar{s}\,\}]\!] \\ &= \oplus\{ e \,[\!]\, \bar{q}, \bar{r}, v \equiv e', \bar{s}\,\}\end{aligned}$$

from Eq. (W2). □

THEOREM 7. *The rules in Figure 10 satisfy the following equations:*

$$[\![\oplus\{ e \,[\!]\, \bar{r}\,\}]\!]^{()}_{w} E = \oplus\{ e \,[\!]\, w \leftarrow E, \bar{r}\,\} \tag{TH2}$$

$$\begin{aligned}[\![\oplus\{ e \,[\!]\, \bar{r}\,\}]\!]^{u}_{w} E = \{\, (u', \oplus\{ e \,[\!]\, w \leftarrow E, w\backslash u \neq \text{NULL}, u' \doteq u, \bar{r}\,\}) \\ [\!]\, u \leftarrow \textstyle\prod_{\lambda w.u}(E)\,\} \text{ for } u \neq ()\end{aligned} \tag{TH3}$$

PROOF. Note that for $u = w$, we have $w \backslash u = ()$ and $\prod_{\lambda w.w}(E) = E$, in which case Eq. (TH3) becomes

$$[\![\oplus\{ e \,[\!]\, \bar{r} \}]\!]^{w}_{w} E = \{ (w, \oplus\{ e \,[\!]\, \bar{r} \}) \,[\!]\, w \leftarrow E \} \qquad \text{(TH4)}$$

since w in the inner comprehension is restricted to be equal to u' in the outer comprehension, and according to Property (P1), the generator $w \leftarrow E$ of the inner comprehension can be safely removed. We prove Eqs. (TH2) through (TH4) using structural induction. That is, assuming that all the subterms of a term satisfy the theorem (induction hypothesis), we prove that a term itself satisfies the theorem (induction step). The induction basis is for the terms that do not contain subterms. All rules in Figure 10 are compositional, which means that the translation of a term is expressed in terms of the translation of its subterms. Thus, the structural induction will follow the recursion pattern of the rules in Figure 10. We prove the theorems for Rules (C5), (C6), (C8), (C9), and (C11). The proofs for the other rules are similar. Proof for Rule (C5):

$$\begin{aligned} [\![\oplus\{ e \,[\!]\, p \}]\!]^{()}_{w} E &= \Delta^{\oplus/\lambda w.e}_{\lambda w.p}(E) && \text{from Rule (C5)} \\ &= \oplus\{ e \,[\!]\, w \leftarrow E, p \} && \text{from Eq. (O4)} \end{aligned}$$

which proves Theorem (TH2). Proof for Rule (C6):

$$\begin{aligned} & [\![\oplus\{ e \,[\!]\, v \leftarrow X, \bar{r}, p \}]\!]^{()}_{w} E \\ &= [\![\oplus\{ e \,[\!]\, \bar{r}, p[\overline{(w, v)}] \}]\!]^{()}_{(w, v)} (E \bowtie_{\lambda(w, v).p[(w, v)]} (\sigma_{\lambda v.p[v]}(X))) \\ &= \oplus\{ e \,[\!]\, (w, v) \leftarrow (E \bowtie_{\lambda(w, v).p[(w, v)]} (\sigma_{\lambda v.p[v]}(X))), \bar{r}, p[\overline{(w, v)}] \} \\ &= \oplus\{ e \,[\!]\, (w, v) \leftarrow \{ (w, v) \,[\!]\, w \leftarrow E, v \leftarrow \{ v \,[\!]\, v \leftarrow X, p[v] \}, p[(w, v)] \}, \bar{r}, p[\overline{(w, v)}] \} \\ &= \oplus\{ e \,[\!]\, w \leftarrow E, v \leftarrow X, p[v], p[(w, v)], \bar{r}, p[\overline{(w, v)}] \} \\ &= \oplus\{ w \leftarrow E, v \leftarrow X, \bar{r}, p \} \end{aligned}$$

from Rule (C6), from induction hypothesis (TH2), from Eqs. (O2) and (O1), and by using normalization. This proves Theorem (TH2). Proof for Rule (C8):

$$\begin{aligned} & [\![\oplus\{ e \,[\!]\, p \}]\!]^{u}_{w} E = \Gamma^{\oplus/\lambda w.e/\lambda w.u}_{\lambda w.p/\lambda w.w \backslash u}(E) \\ &= \{ (u[v/w], \oplus\{ e \,[\!]\, w \leftarrow E, w \backslash u \neq \text{NULL}, u[v/w] \doteq u, p \}) \,[\!]\, v \leftarrow \textstyle\prod_{\lambda w.u}(E) \} \\ &= \{ (u', \oplus\{ e \,[\!]\, w \leftarrow E, w \backslash u \neq \text{NULL}, u' \doteq u, p \}) \,[\!]\, u' \leftarrow \textstyle\prod_{\lambda w.u}(E) \} \end{aligned}$$

(from Rule (C8) and from Eq. (O7), which proves Theorem (TH3). Proof for Rule (C9):

$$\begin{aligned} & [\![\oplus\{ e \,[\!]\, v \leftarrow X, \bar{r}, p \}]\!]^{u}_{w} E \\ &= [\![\oplus\{ e \,[\!]\, \bar{r}, p[\overline{(w, v)}] \}]\!]^{u}_{(w, v)} (E \mathbin{=\!\!\bowtie}_{\lambda(w, v).p[(w, v)]} (\sigma_{\lambda v.p[v]}(X))) \\ &= \{ (u', \oplus\{ e \,[\!]\, (w, v) \leftarrow (E \mathbin{=\!\!\bowtie}_{\lambda(w, v).p[(w, v)]} (\sigma_{\lambda v.p[v]}(X))), (w, v) \backslash u \neq \text{NULL}, u' \doteq u, \\ &\qquad \bar{r}, p[\overline{(w.v)}] \}) \,[\!]\, u' \leftarrow \textstyle\prod_{\lambda(w, v).u}(E \mathbin{=\!\!\bowtie}_{\lambda(w, v).p[(w, v)]} (\sigma_{\lambda v.p[v]}(X))) \} \end{aligned}$$

from Rule (C9) and from induction hypothesis (TH3). From Eqs. (O5) and (O2) and after normalization, we get

$$E \mathbin{=\!\!\bowtie}_{\lambda(w,v).p[(w,v)]} (\sigma_{\lambda v.p[v]}(X)) = \{\!\{ (w,v) \;[\!]\; w \leftarrow E, v \leftarrow F \}\!\}$$

where $F = \textbf{if} \wedge\{\!\{ \neg p[(w,v)] \;[\!]\; w \neq \text{NULL}, v \leftarrow X, p[v] \}\!\} \textbf{ then } [\text{NULL}] \textbf{ else } \{\!\{ v \;[\!]\; v \leftarrow X, p[v], p[(w,v)] \}\!\}$. Thus, $\prod_{\lambda(w,v).u}(E \mathbin{=\!\!\bowtie}_{\lambda(w,v).p[(w,v)]} (\sigma_{\lambda v.p[v]}(X)))$ is equal to $\prod_{\lambda(w,v).u}(\{\!\{ (w,v) \;[\!]\; w \leftarrow E, v \leftarrow F \}\!\})$, which is equal to $\prod_{\lambda w.u}(E)$, according to Property (P2). Therefore, we have

$$
\begin{aligned}
&[\![\oplus\{\!\{ e \;[\!]\; v \leftarrow X, \bar{r}, p \}\!\}]\!]^{u}_{w} E \\
&= \{\!\{ (u', \oplus\{\!\{ e \;[\!]\; (w,v) \leftarrow \{\!\{ (w,v) \;[\!]\; w \leftarrow E, v \leftarrow F \}\!\}, v \neq \text{NULL}, w \backslash u \neq \text{NULL}, \\
&\qquad u' \doteq u, \bar{r}, p[\overline{(w,v)}] \}\!\}) \;[\!]\; u' \leftarrow \textstyle\prod_{\lambda w.u}(E) \}\!\} \\
&= \{\!\{ (u', \oplus\{\!\{ e \;[\!]\; w \leftarrow E, v \leftarrow F, v \neq \text{NULL}, w \backslash u \neq \text{NULL}, u' \doteq u, \bar{r}, p[\overline{(w,v)}] \}\!\}) \\
&\qquad [\!]\; u' \leftarrow \textstyle\prod_{\lambda w.u}(E) \}\!\} \\
&= \{\!\{ (u', \oplus\{\!\{ e \;[\!]\; w \leftarrow E, v \leftarrow \{\!\{ v \;[\!]\; v \leftarrow X, p[v], p[(w,v)] \}\!\}, w \backslash u \neq \text{NULL}, \\
&\qquad u' \doteq u, \bar{r}, p[\overline{(w,v)}] \}\!\}) \;[\!]\; u' \leftarrow \textstyle\prod_{\lambda w.u}(E) \}\!\} \\
&= \{\!\{ (u', \oplus\{\!\{ e \;[\!]\; w \leftarrow E, v \leftarrow X, p[v], p[(w,v)], w \backslash u \neq \text{NULL}, u' \doteq u, \bar{r}, \\
&\qquad p[\overline{(w,v)}] \}\!\}) \;[\!]\; u' \leftarrow \textstyle\prod_{\lambda w.u}(E) \}\!\} \\
&= \{\!\{ (u', \oplus\{\!\{ e \;[\!]\; w \leftarrow E, v \leftarrow X, w \backslash u \neq \text{NULL}, u' \doteq u, \bar{r}, p \}\!\}) \;[\!]\; u' \leftarrow \textstyle\prod_{\lambda w.u}(E) \}\!\}
\end{aligned}
$$

since $v \neq$ NULL. Thus we have a proof of Rule (C9). For Rule (C11), we have two cases; when $u = ()$:

$$
\begin{aligned}
&[\![\oplus\{\!\{ e_1 \;[\!]\; \bar{s}, p(\otimes\{\!\{ e_2 \;[\!]\; \bar{r} \}\!\}) \}\!\}]\!]^{()}_{w} E \\
&\qquad = [\![\oplus\{\!\{ e_1 \;[\!]\; \bar{s}, p(v) \}\!\}]\!]^{()}_{(w,v)}([\![\otimes\{\!\{ e_2 \;[\!]\; \bar{r} \}\!\}]\!]^{w}_{w} E) \\
&\qquad = \oplus\{\!\{ e_1 \;[\!]\; (w,v) \leftarrow ([\![\otimes\{\!\{ e_2 \;[\!]\; \bar{r} \}\!\}]\!]^{w}_{w} E), \bar{s}, p(v) \}\!\} \\
&\qquad = \oplus\{\!\{ e_1 \;[\!]\; (w,v) \leftarrow \{\!\{ (w, \otimes\{\!\{ e_2 \;[\!]\; \bar{r} \}\!\}) \;[\!]\; w \leftarrow E \}\!\}, \bar{s}, p(v) \}\!\} \\
&\qquad = \oplus\{\!\{ e_1 \;[\!]\; w \leftarrow E, \bar{s}, p(\otimes\{\!\{ e_2 \;[\!]\; \bar{r} \}\!\}) \}\!\}
\end{aligned}
$$

(from Rule (C11), from hypotheses (TH2) and (TH4), and by using normalization); and when $u \neq ()$:

$$
\begin{aligned}
&[\![\oplus\{\!\{ e_1 \;[\!]\; \bar{s}, p(\otimes\{\!\{ e_2 \;[\!]\; \bar{r} \}\!\}) \}\!\}]\!]^{u}_{w} E \\
&= [\![\oplus\{\!\{ e_1 \;[\!]\; \bar{s}, p(v) \}\!\}]\!]^{u}_{(w,v)}([\![\otimes\{\!\{ e_2 \;[\!]\; \bar{r} \}\!\}]\!]^{w}_{w} E) \\
&= \{\!\{ (u', \oplus\{\!\{ e_1 \;[\!]\; (w,v) \leftarrow ([\![\otimes\{\!\{ e_2 \;[\!]\; \bar{r} \}\!\}]\!]^{w}_{w} E), (w,v) \backslash u \neq \text{NULL}, u' \doteq u, \bar{s}, p(v) \}\!\}) \\
&\qquad [\!]\; u' \leftarrow \textstyle\prod_{\lambda(w,v).u}([\![\otimes\{\!\{ e_2 \;[\!]\; \bar{r} \}\!\}]\!]^{w}_{w} E) \}\!\} \\
&= \{\!\{ (u', \oplus\{\!\{ e_1 \;[\!]\; (w,v) \leftarrow (\{\!\{ (w, \otimes\{\!\{ e_2 \;[\!]\; \bar{r} \}\!\}) \;[\!]\; w \leftarrow E \}\!\}), (w,v) \backslash u \neq \text{NULL},
\end{aligned}
$$

$$u' \doteq u, \bar{s}, p(v) \}) \,[\!]\, u' \leftarrow \textstyle\prod_{\lambda(w,v).u}(\{ (w, \otimes\{ e_2 \,[\!]\, \bar{r} \}) \,[\!]\, w \leftarrow E \}) \}$$
$$= \{ (u', \oplus\{ e_1 \,[\!]\, (w, v) \leftarrow (\{ (w, \otimes\{ e_2 \,[\!]\, \bar{r} \}) \,[\!]\, w \leftarrow E \}), (w, v)\backslash u \neq \text{NULL},$$
$$u' \doteq u, \bar{s}, p(v) \}) \,[\!]\, u' \leftarrow \textstyle\prod_{\lambda w.u}(E) \}$$
$$= \{ (u', \oplus\{ e_1 \,[\!]\, w \leftarrow E, w\backslash u \neq \text{NULL}, u' \doteq u, \bar{s}, p(\otimes\{ e_2 \,[\!]\, \bar{r} \}) \}) \,[\!]\, u' \leftarrow \textstyle\prod_{\lambda w.u}(E) \}$$

from Rule (C11), from induction hypotheses (TH3) and (TH4), from Property (P2), and by using normalization. □

Corollary 1. The rules in Figure 10 are meaning-preserving. That is,

$$[\![\oplus\{ e \,[\!]\, \bar{r} \}]\!]_{()}^{()}\{()\} = \oplus\{ e \,[\!]\, \bar{r} \}.$$

PROOF. This corollary is a consequence of Eq. (TH2) of Theorem 7 (for $w = ()$ and $E = \{()\}$). □

REFERENCES

BEECH, D. 1993. Collections of objects in SQL3. In *Proceedings of the Nineteenth International Conference on Very Large Data Bases* (VLDB '93, Dublin, Ireland, Aug.). Morgan Kaufmann Publishers Inc., San Francisco, CA, 244–255.

BEERI, C. AND KORNATZKY, Y. 1990. Algebraic optimization of object-oriented query languages. In *Proceedings of the Third International Conference on Database Theory* (ICDT '90, Paris, France, Dec.), S. Abiteboul and P. C. Kanellakis, Eds. Springer Lecture Notes in Computer Science. Springer-Verlag, New York, NY, 72–88.

BELLANTONI, S. AND COOK, S. 1992. A new recursion-theoretic characterization of the polytime functions (extended abstract). In *Proceedings of the 24th Annual ACM Symposium on Theory of Computing* (STOC '92, Victoria, B. C., Canada, May 4–6), R. Kosaraju, M. Fellows, A. Wigderson, and J. Ellis, Chairs. ACM Press, New York, NY, 283–293.

BREAZU-TANNEN, V., BUNEMAN, P., AND NAQVI, S. 1992a. Structural recursion as a query language. In *Proceedings of the Third International Workshop on Database Programming Languages: Bulk Types & Persistent Data* (DBPL3, Nafplion, Greece, Aug. 27–30), P. Kanellakis and J. W. Schmidt, Eds. Morgan Kaufmann Publishers Inc., San Francisco, CA, 9–19.

BREAZU-TANNEN, V., BUNEMAN, P., AND WONG, L. 1992b. Naturally embedded query languages. In *Proceedings of the 4th International Conference on Database Theory* (LNCS 646, Berlin). Springer-Verlag, New York, NY.

BREAZU-TANNEN, V. AND SUBRAHMANYAM, R. 1991. Logical and computational aspects of programming with sets/bags/lists. In *Proceedings of the 18th International Colloquium on Automata, Languages and Programming* (Madrid, July 8–12), J. L. Albert, B. R. Artalejo, and B. Monien, Eds. Springer Lecture Notes in Computer Science. Springer-Verlag, New York, NY, 60–75.

BUNEMAN, P., LIBKIN, L., SUCIU, D., TANNEN, V., AND WONG, L. 1994. Comprehension syntax. *SIGMOD Rec. 23*, 1 (Mar.), 87–96.

BUNEMAN, P., NAQVI, S., TANNEN, V., AND WONG, L. 1995. Principles of programming with complex objects and collection types. *Theor. Comput. Sci. 149*, 1 (Sept. 18), 3–48.

CAREY, M. AND DEWITT, D. 1996. Of objects and databases: A decade of turmoil. In *Proceedings of the 22nd International Conference on Very Large Data Bases* (VLDB '96, Bombay, Sept.). Morgan Kaufmann Publishers Inc., San Francisco, CA, 3–14.

CAREY, M. J., DEWITT, D. J., FRANKLIN, M. J., HALL, N. E., MCAULIFFE, M. L., NAUGHTON, J. F., SCHUH, D. T., SOLOMON, M. H., TAN, C. K., TSATALOS, O. G., WHITE, S. J., AND ZWILLING, M. J. 1994. Shoring up persistent applications. *SIGMOD Rec. 23*, 2 (June), 383–394.

CATTELL, R., ED. 2000. *The Object Data Standard: ODMG 3.0*. Morgan Kaufmann, San Mateo, CA.

CHAN, D. AND TRINDER, P. 1994. Object comprehensions: A query notation for object-oriented databases. In *Proceedings of the 12th British Conference on Databases* (Guildford, UK, July). Morgan Kaufmann Publishers Inc., San Francisco, CA, 55–72.

CHERNIACK, M. AND ZDONIK, S. 1998a. Changing the rules: Transformations for rule-based optimizers. *SIGMOD Rec. 27*, 2, 61–72.

CHERNIACK, M. AND ZDONIK, S. 1998b. Inferring function semantics to optimize queries. In *Proceedings of the 24th International Conference on Very Large Data Bases*. 239–250.

CLAUSSEN, J., KEMPER, A., MOERKOTTE, G., AND PEITHNER, K. 1997. Optimizing queries with universal quantification in object-oriented and object-relational databases. In *Proceedings of the 23rd International Conference on Very Large Data Bases* (VLDB '97, Athens, Greece, Aug.). 286–295.

CLUET, S. AND DELOBEL, C. 1992. A general framework for the optimization of object-oriented queries. *SIGMOD Rec. 21*, 2 (June 1), 383–392.

CLUET, S. AND MOERKOTTE, G. 1995a. Efficient evaluation of aggregates on bulk types. Tech. Rep. 95-05. Aachen University of Technology, Aachen, Germany.

CLUET, S. AND MOERKOTTE, G. 1995b. Nested queries in object bases. In *Proceedings of the Fifth International Workshop on Database Programming Languages* (Gubbio, Italy, Sept.).

COLBY, L. S. 1989. A recursive algebra and query optimization for nested relations. *SIGMOD Rec. 18*, 2 (June), 273–283.

DANFORTH, S. AND VALDURIEZ, P. 1992. A FAD for data intensive applications. *IEEE Trans. Knowl. Data Eng. 4*, 1 (Feb.), 34–51.

DANIELS, S., GRAEFE, G., KELLER, T., MAIER, D., SCHMIDT, D., AND VANCE, B. 1991. Query optimization in revelation, an overview. *Data Eng. 14*, 2 (June), 58–62.

DEUX, O. 1990. The story of O2. *IEEE Trans. Knowl. Data Eng. 2*, 1, 91–108.

EISENBERG, A. AND MELTON, J. 1999. SQL1999, formerly known as SQL3. *SIGMOD Rec. 28*, 1 (Mar.), 131–138.

FEGARAS, L. 1993. Efficient optimization of iterative queries. In *Proceedings of the International Workshop on Database Programming Languages* (New York, NY). Morgan Kaufmann Publishers Inc., San Francisco, CA, 200–225.

FEGARAS, L. 1998a. A new heuristic for optimizing large queries. In *Proceedings of the Ninth International Conference on DEXA* (Vienna, Aug.). Springer-Verlag, Heidelberg, Germany, 726–735.

FEGARAS, L. 1998b. Query unnesting in object-oriented databases. *SIGMOD Rec. 27*, 2, 49–60.

FEGARAS, L. 1999a. Optimizing queries with object updates. *J. Intell. Inf. Syst. 12*, 219–242.

FEGARAS, L. 1999b. VOODOO: A visual object-oriented database language for ODMG OQL. In *Proceedings of the First ECOOP Workshop on Object-Oriented Databases* (Lisbon, Portugal, June). 61–72.

FEGARAS, L. AND MAIER, D. 1995. Towards an effective calculus for object query languages. In *Proceedings of the 1995 ACM SIGMOD International Conference on Management of Data* (SIGMOD '95, San Jose, CA, May 23–25), M. Carey and D. Schneider, Eds. ACM Press, New York, NY, 47–58.

FEGARAS, L., MAIER, D., AND SHEARD, T. 1993. Specifying rule-based query optimizers in a reflective framework. In *Proceedings of the Conference on Deductive and Object-Oriented Databases* (Phoenix, AZ, Dec.). Springer-Verlag, Vienna, Austria, 146–168.

FEGARAS, L., SRINIVASAN, C., RAJENDRAN, A., AND MAIER, D. 2000. λ-db: An odmg-based object-oriented dbms. In *Proceedings of the ACM SIGMOD Conference on Management of Data* (SIGMOD '2000, Dallas, TX, May). ACM Press, New York, NY.

GANSKI, R. A. AND WONG, H. K. T. 1987. Optimization of nested SQL queries revisited. *SIGMOD Rec. 16*, 3 (Dec.), 23–33.

KEMPER, A. AND MOERKOTTE, G. 1990. Advanced query processing in object bases using access support relations. In *Proceedings of the 16th International Conference on Very Large Data Bases* (VLDB, Brisbane, Australia, Aug. 13-16), D. McLeod, R. Sacks-Davis, and H. Schek, Eds. Morgan Kaufmann Publishers Inc., San Francisco, CA, 290–301.

KIM, W. 1982. On optimizing an SQL-like nested query. *ACM Trans. Database Syst. 7*, 3 (Sept.), 443–469.

LEIVANT, D. 1993. Stratified functional programs and computational complexity. In *Proceedings of the 20th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages* (POPL '93, Charleston, SC, Jan. 10–13), S. L. Graham, Ed. ACM Press, New York, NY, 325–333.

LEUNG, T., MITCHELL, G., SUBRAMANIAN, B., VANCE, B., VANDENBERG, S., AND ZDONIK, S. 1993. The AQUA data model and algebra. In *Proceedings of the International Workshop on Database Programming Languages* (New York, NY). Morgan Kaufmann Publishers Inc., San Francisco, CA, 157–175.

LEVY, A. Y., MUMICK, I. S., AND SAGIV, Y. 1994. Query optimization by predicate move-around. In *Proceedings of the 20th International Conference on Very Large Data Bases* (VLDB'94, Santiago, Chile, Sept.). VLDB Endowment, Berkeley, CA, 96–107.

LIBKIN, L., MACHLIN, R., AND WONG, L. 1996. A query language for multidimensional arrays: design, implementation, and optimization techniques. *SIGMOD Rec. 25*, 2, 228–239.

LIN, J. AND OZSOYOGLU, Z. M. 1996. Processing OODB queries by O-algebra. In *Proceedings of the Fifth International Conference on Information and Knowledge Management* (CIKM '96, Rockville, MD, Nov. 12–16), M. T. Özsu and K. Barker, Eds. ACM Press, New York, NY, 134–142.

MAIER, D. AND VANCE, B. 1993. A call to order. In *Proceedings of the Twelfth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems* (PODS, Washington, DC, May 25–28), C. Beeri, Chair. ACM Press, New York, NY, 1–16.

MEIJER, E., FOKKINGA, M., AND PATERSON, R. 1991. Functional programming with bananas, lenses, envelopes and barbed wire. In *Proceedings of the Fifth ACM Conference on Functional Programming Languages and Computer Architecture* (Cambridge, MA, Aug. 26–30), J. Hughes, Ed. Springer Lecture Notes in Computer Science. Springer-Verlag, New York, NY, 124–144.

MUMICK, I. S., FINKELSTEIN, S. J., PIRAHESH, H., AND RAMAKRISHNAN, R. 1990. Magic is relevant. *SIGMOD Rec. 19*, 2 (June), 247–258.

MURALIKRISHNA, M. 1992. Improved unnesting algorithms for join aggregate SQL queries. In *Proceedings of the 18th International Conference on Very Large Data Bases* (Vancouver, B.C., Aug.). VLDB Endowment, Berkeley, CA, 91–102.

OZSOYOGLU, Z. AND WANG, J. 1992. A keying method for a nested relational database management system. In *Proceedings of the Eighth International Conference on Data Engineering* (Tempe, AZ, Feb.).

PEYTON JONES, S. L. 1987. *The Implementation of Functional Programming Languages*. Prentice-Hall, New York, NY.

PIERCE, B. C. 1991. *Basic Category Theory for Computer Scientists*. Foundations of Computing. MIT Press, Cambridge, MA.

PISTOR, P AND TRAUNMUELLER, R. 1986. A database language for sets, lists and tables. *Inf. Syst. 11*, 4 (Oct.), 323–336.

SELINGER, P. G., ASTRAHAN, M. M., LORIE, R. A., AND PRICE, T. G. 1979. Access path selection in a relational database management system. In *Proceedings of the ACM SIGMOD International Conference on Management of Data* (SIGMOD '79, Boston, MA, May 30–June 1). ACM Press, New York, NY, 23–34.

STEENHAGEN, H. J., APERS, P. M. G., AND BLANKEN, H. M. 1994. Optimization of nested queries in a complex object model. In *Proceedings of the Fourth International Conference on Extending Database Technology: Advances in Database Technology* (EDBT '94, Cambridge, UK, Mar. 28–31), M. Jarke, J. Bubenko, and K. Jeffery, Eds. Springer Lecture Notes in Computer Science. Springer-Verlag, New York, NY, 337–350.

THOMPSON, S. 1998. *Haskell: The Craft of Functional Programming*. Addison-Wesley, Reading, MA.

TRINDER, P. 1992. Comprehensions, a query notation for DBPLs. In *Proceedings of the Third International Workshop on Database Programming Languages: Bulk Types & Persistent Data* (DBPL3, Nafplion, Greece, Aug. 27–30), P. Kanellakis and J. W. Schmidt, Eds. Morgan Kaufmann Publishers Inc., San Francisco, CA, 55–68.

TRINDER, P. AND WADLER, P. 1989. Improving list comprehension database queries. In *Proceedings of on TENCON'89* (Bombay, Nov.). 186–192.

WADLER, P. 1990. Comprehending monads. In *Proceedings of the 1990 ACM Symposium on LISP and Functional Programming* (Nice, France, June 27–29), G. Kahn, Chair. ACM Press, New York, NY, 61–78.

WONG, L. 1993. Normal forms and conservative properties for query languages over collection types. In *Proceedings of the Twelfth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems* (PODS, Washington, DC, May 25–28), C. Beeri, Chair. ACM Press, New York, NY, 26–36.

WONG, L. 1994. Querying nested collections. Ph.D. Dissertation. University of Pennsylvania, Philadelphia, PA.

Received: November 1998; revised: July 2000; accepted: October 2000