The *MultiView* OODB View System: Design and Implementation *

Harumi A. Kuno and Elke A. Rundensteiner
Dept. of Elect. Engineering and Computer Science
Software Systems Research Laboratory
The University of Michigan, 1301 Beal Avenue
Ann Arbor, MI 48109-2122
e-mail: kuno@umich.edu, rundenst@eecs.umich.edu

Keywords: View materialization, incremental view maintenance, subjectivity, object-oriented database views, data warehousing, classification

Views are an established technique for restructuring and repartitioning the format of data, classes, and schemata so that applications can customize shared data objects without affecting other applications' perceptions of the data. The *MultiView* system is one of the first OODB systems to support dynamic and updatable materialized object-oriented database views. *MultiView* is fully functional and is being used for a number of projects. In this paper, we describe our system's architecture, the services it provides, and the decisions we made during our implementation. Although the GemStone system we chose for our implementation base offers many features that greatly aided our implementation, it does not support several key object-model properties that are critical for the realization of our design principles. These fundamental properties include multiple classification, dynamic object-restructuring, and the ability to make dynamic changes to the schema. We hence adopted a flexible and powerful technique known as object-slicing to construct the *MultiView* object model. This object-slicing architecture successfully addresses our requirements. The *MultiView* system is distinguished by a number of unique features, including the incorporation of virtual classes into the global schema as first-class database citizens, support for capacity-augmenting views (virtual classes that add new extrinsic properties or behavior), view materialization strategies that take advantage of object-oriented modeling features, and a graphical interface that is tailored to provide easy access to the *MultiView* system. The resulting system preserves all of the capabilities of the underlying GemStone OODB while providing support for dynamic and updatable materialized object-oriented views.

*This work was supported in part by the NSF RIA grant #IRI-9309076, NSF NYI grant #IRI 94-57609, and the University of Michigan Faculty Award Program. We would also like to thank our industrial sponsors, in particular, AT&T, Illustra, and Intel. Harumi Kuno is grateful for support from the NASA Graduate Student Researchers Program.

1. Introduction

MOTIVATION AND BACKGROUND.

Views provide logical data independence, and offer a means by which data can be repartitioned and restructured in format. Applications can use views to customize shared data objects, even adding new extrinsic properties and behavior, without affecting other applications' perceptions of the data. Views have been established as an effective method for virtually restructuring data, classes, and schemata so as to meet the needs of specific applications or users; for integrating heterogeneous and distributed systems; for security and access rights restriction; and for achieving interoperability by hiding the idiosyncrasies of component systems to be integrated into one unified, yet federated system [1, 4, 5, 6, 17]. Because views allow a database object to behave differently depending on the context in which it appears, views are a recognized technique for extending a database system with support for subjectivity.

Many research areas can profit from the flexibility provided by object-oriented views. In addition to the traditional use of views for facilitating the sharing of data by applications, current researchers are examining new areas that can specifically benefit from object-oriented views. Barsalou et al. use object-based views to integrate object-oriented and relational databases while preserving update capabilities [4]. Ra and Rundensteiner utilize views to provide transparent schema evolution, preserving existing views through schema change [33]. Because the object-oriented paradigm offers a more powerful model for integration than the relational one, several papers discuss the integration of heterogeneous data repositories via object-oriented views [5, 4, 17].

While a number of researchers have begun to study view mechanisms with regard to object-oriented databases (OODBs) [1, 15, 24, 40, 35, 37], little work has been done regarding implementation issues related to object-oriented views. Furthermore, of the few research papers that discuss implementations of OODB view systems, most support only limited functionality at this time. There are several outstanding issues that must be addressed for the successful implementation of an object-oriented view system. First, creating views in an object-oriented model is not a simple transfer of the relational view solution to the object-oriented model. Much of the functionality typically provided by relational databases must be re-evaluated in the context of this new technology—for example, how to overcome the view update problem of the relational view mechanism, and how to utilize the complexity of the object-oriented data model for view definition (such as behavioral customization, view hierarchy manipulation, and property inheritance among view classes). Second, and more critically, to the best of our knowledge, no existing commercial OODB system supports all of the properties we identify as critical for the realization of fully-functional object-oriented views. These properties are quite fundamental, including features such as multiple classification, dynamic object-restructuring, and dynamic changes to the schema. This implies that either the data models of existing OODB systems must be extended, or other solutions must be found to address this problem.

THE *MULTIVIEW* APPROACH.

Here at the University of Michigan, we have an NSF-funded project developing the *MultiView* view management system [35]. To the best of our knowledge, *MultiView* is the first to support updatable materialized views [26, 27]. In the context of the *MultiView* project, we provide solutions to the issues outlined above. In order to facilitate the implementation of our system, we chose to use the commercial GemStone OODB as our base [1]. However, although GemStone offers many features that greatly aided our implementation, neither it nor any other existing commercial OODB system supports several of the key object-model properties that are necessary for the realization of our design principles. We therefore developed an object-slicing representational model, a flexible and powerful technique, that addresses these deficiencies. Using the object-slicing representational model, we were able to construct the *MultiView* object model, which provides all the features required for our view system, on top of GemStone. We ran an experimental evaluation of our system to determine the overhead that can be attributed to our choice of an object-slicing representation paradigm (in terms of both storage costs and performance). Our experiments confirm that the overhead is mixed, ranging from performance gains of 50% to increases in execution time of 200%, depending on the characteristics of the specific access patterns.

Our *MultiView* system is currently fully functional, provides clean object-oriented characteristics, supports dynamic updatable materialized views, and is being used for a number of projects. Another unique feature of *MultiView*, besides its support for updatable incrementally materialized views, is that *MultiView* is the only implemented OODB view system to treat virtual classes as first-class database citizens. We integrate both base and virtual classes into a unified global schema. Virtual classes in *MultiView* participate in the actual inheritance hierarchy and thus behave just like base classes. Furthermore, virtual classes in *MultiView* can independently define attributes and methods in addition to those derived from base classes, i.e., they are data capacity-augmenting views.

In previous papers we have described individual components of *MultiView*. [35] introduced the founding principles of the *MultiView* view specification methodology as a three-step process. [36] discusses *MultiView*'s strategy for the generation of view schemata, and [37] presents the classification algorithm by which new virtual classes are integrated into the global schema. In [27], we introduce some algorithms for incremental propagation of updates for maintaining materialized views in *MultiView*. In this paper we now present a unified overview of the *MultiView* project, focusing primarily on issues related to its system implementation. In particular, we outline the motivation and execution of our design decisions, system architecture, and the *MultiView* system's capabilities. The implementation solutions we describe are portable to other object-oriented systems.

SYSTEM OVERVIEW.

We exploit the unique features of the object-oriented paradigm in the design and implementation of the *MultiView* system. As shown in Figure 1, the *MultiView* system consists of a nested-layer architecture.

- The innermost layer of the figure represents the underlying GemStone database system kernel. GemStone provides many of the basic functionalities for our system, including persistence, database programming language support, and transaction management.
- The second layer illustrates how we overlay an object-slicing representation on top of the GemStone architecture in order to implement the fundamental object model needed for view support. We identified features that must be provided by

an object model for the flexible support of object-oriented views according to pure object-oriented principles. Like all the other OODB systems we surveyed, the GemStone OODB underlying our system does not support all of the properties required by our model. We hence designed the object-slicing representational layer, which resolves the significant differences between the underlying system and our object model. Object-slicing allows us to provide all of our required object model properties, including multiple classification, multiple inheritance, the dynamic restructuring of objects and classes, and dynamic changes to the class hierarchy.

- The third layer of the architecture represents *MultiView*'s support for materialized views and virtual schemata. Our implementation provides a Smalltalk interface that supports the creation of virtual classes and the automatic integration of the base and virtual classes into a consistent and correct global schema, the creation and customization of view schemata, the entry and update of data, and the retrieval of data via user queries.
- Finally, *MultiView*'s graphic user interface provides access to the system in an intuitive and user-friendly fashion. The *MultiView* graphic user interface displays global and view schemata in visually intuitive graphic windows, facilitates the creation of new virtual classes and view schemata, and provides an interface for data manipulation and querying class extents. The current version of the *MultiView Viewer*, or short, the *MultiViewer*, is written in Tk/Tcl using C.

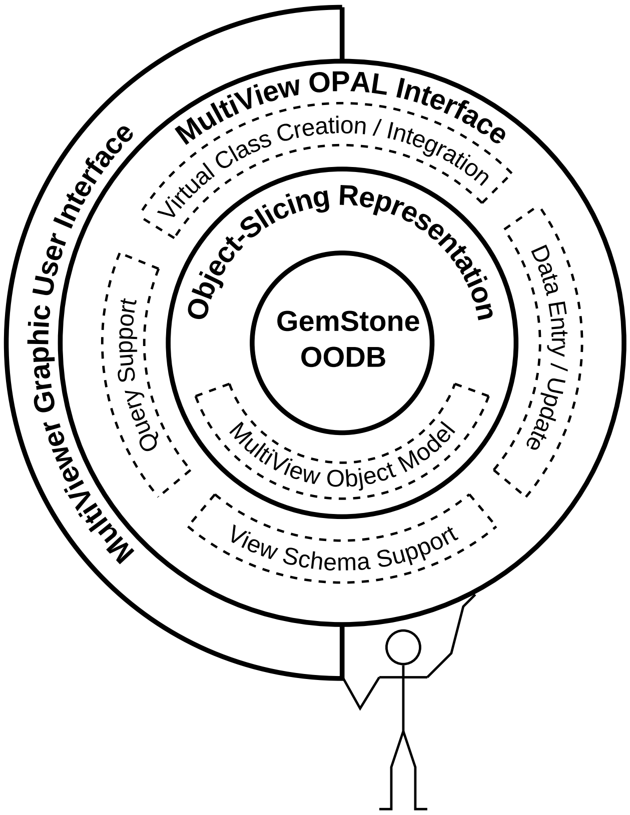

FIG. 1. The Nested Layer Architecture of the *MultiView* System Design.

As represented by the anthropomorphic stick-figure below the diagram of system components, users can access the *MultiView* system either through the graphic user interface or through *MultiView*'s Smalltalk interface. The *MultiView* system as described in this paper is fully functional. We have used *MultiView* to implement the University of Wisconsin's OO7 benchmark suite schema, populated the database with OO7's 10,000 object parts-assembly example, and have tested *MultiView* using both OO7 benchmark queries and queries of our own design. A summary of our performance evaluation, as well as of usage of the *MultiView* prototype in other projects, will be presented in later sections of this paper.

OUTLINE OF THIS PAPER.

In the remainder of this paper, we describe the object model, techniques, and algorithms underlying our implementation of *MultiView*. Our work should be of interest to researchers studying view systems. In addition, because many aspects of our experience, for example our use of an object-slicing technique, can be applied to the general support of object model features such as multiple classification and type-specific behavior, other implementors seeking to support subjectivity will also benefit from this work. We begin with the goals and object model of the *MultiView* project in Section 2. We present our approach of realizing object model features required for the support of our goals in Section 3, then discuss each of the components of our implementation in Sections 4 through 6. We discuss related work in Section 7. Finally, we conclude in Section 8 with a summary of the implementation status of *MultiView*, and a discussion of our contributions and future work.

2. Object Model of *MultiView*

2.1 *Goals of the* MultiView *Project*

The purpose of the *MultiView* project is to build, implement, and evaluate a system for OODB view support. The goals of the project include the following:

- Users should be able to create customized virtual classes at any time.
- Users should be able to query and update both base and virtual classes.
- Users should be able to create and modify customized virtual schemata at any time.

- Virtual classes and schemata should be first-class citizens of the database.

By first-class database citizen, we mean that virtual classes and schemata should look and feel like base classes and schemata. Virtual and base classes should be fully integrated into a global class hierarchy in terms of both type and extent. This implies that both base and virtual classes should fully participate in the inheritance scheme in such a way that either type of class can act as a point of inheritance for properties, and also that the extents of both virtual and base classes should be kept consistent with the class hierarchy. We associate both types and extents with our classes because this information must be maintained in any case in order to correctly populate virtual classes. (For example, consider a hide class that is derived from a select class.)

In addition, *MultiView* supports *capacity-augmenting* virtual classes. First introduced by Bertino [6], a capacity-augmenting virtual class is a virtual class that extends existing data with instance variables that are not derived from the source classes of the virtual class. Finally, the user should be able to select both base and virtual classes from the global schema at any time and add them to view schemata. The *MultiView* system should automatically integrate these selected classes into a consistent and correct view schemata.

In the following subsections, we give a description of an object model that fulfills the *MultiView* project goals. We first formalize our basic concepts and principles, then extend this initial description to include virtual classes and schemata and materialized views.

2.2 Basic Concepts and Principles

An **object instance** (or short, **object**) represents an entity. Anything with distinct existence in objective or conceptual reality can be represented as an object. Let O be an infinite set of object instances. Each object $O_i \in O$ consists of state (the **instance variables** or attributes of the object) and behavior (the **methods**, or messages, to which the object can respond). Because our model is object-oriented and assumes full encapsulation, any modification to an instance variable of an object in our system must be accomplished by means of an **accessing method**. An **accessing method** is a method that modifies or retrieves an instance variable. Together, the **methods** and **instance variables** of an object are referred to as its **properties.** In addition, each object has a unique system-generated value-independent **object identifier**, which makes it possible to distinguish between equality and identity, to share sub-objects among complex objects, and to perform updates on common sub-objects.

Objects are grouped into sets of similar objects that share a common structure and behavior. The term **class** denotes this specification of a common structure and behavior. Let C be the set of all classes in a database. A **class**, $C_i \in C$, has a unique class name, a **type**, and a set membership (known as the class's **extent**), as defined below [2]. We use the term **type** to indicate the set of applicable property functions shared by all object-instances of the class. The set of property functions defined for class C_i is denoted as **properties**(C_i).

Definition. *A* **type** *is a tuple,* **<InstanceVariables, Methods>**, *where InstanceVariables is the set of attributes (instance variables) possessed by the type and Methods is the set of* **methods** *defined by the type. We refer to the the type associated with a class,* $C_i \in C$ *by* **type**(C_i).

Figure 2 shows the class that *MultiView* would generate based on the class definition given in Figure 3, creating a new class named *Person* that has instance variables (and accompanying accessing methods) to store a Person's birthyear and name. As shown in Figure 3, the *Person* class's type contains the instance variables `birthyear` and `name`, retrieval methods `birthyear` and `name`, and assignment methods `birthyear:` and `name:`. (In addition to a Smalltalk-like command-level interface, the current implementation of *MultiView* also supports its own graphic user interface, which can be used for view-related tasks.)

Definition. *For two classes* C_i *and* $C_j \in C$, C_i *is called a* **subtype** *of* C_j, *denoted by* $C_i \preceq C_j$ *if and only if* (**properties**(C_i) $\supseteq$ **properties**(C_j)). C_j *is called a* **supertype** *of* C_i *if and only if* C_i *is a subtype of* C_j. *All properties defined for a supertype are* **inherited** *by its subtypes.*

Definition. *The* **extent** *(also called* **set-membership**) *of a class* $C_i \in C$ *is the collection of object instances that belong to that class. The membership of an object instance, o, in the extent of a class,* C_i, *is denoted by* $o \in C_i$. *We refer to the set of objects that are implemented as* **direct instances** *of* C_i *(objects whose most specific type is* C_i*) as* **LocalExtent**(C_i), *or short,* **extent**(C_i). *We call the collection of all objects that possess a class's type the* **GlobalExtent**(C_i).

For example, we can create a new instance of the Person class and assign values to its birthyear and name instance variables (Figure 4). This instance is automatically added to the extent of the *Person* class.

Definition. *For two classes* C_i *and* $C_j \in C$, C_i *is called a* **subset** *of* C_j, *denoted by* $C_i \subseteq C_j$, *if and only if* $(\forall o \in O)((o \in C_i) \Rightarrow (o \in C_j))$. *Thus* $\mathbf{GlobalExtent}(C_i) = \bigcup \mathbf{LocalExtent}(C_j)\ \forall C_j \subseteq C_i$. C_j *is called a* **superset** *of* C_i, *denoted by* $C_j \supseteq C_i$, *if and only if* $C_i \subseteq C_j$.

Definition. *For two classes* C_i *and* $C_j \in C$, C_i *is called a* **subclass** *of* C_j, *denoted by* C_i is-a C_j, *if and* *only if* $(C_i \subseteq C_j)$ *and* $(C_i \preceq C_j)$. C_i *is a* **direct subclass** *of* C_j *if* $\nexists C_k \in C$ *s.t.* $k \neq i \neq j$, C_i *is a* **subclass** *of* C_k, *and* C_k *is a* **subclass** *of* C_j. *Similarly,* C_j *is called a* **superclass** *of* C_i *if and only if* C_i is-a C_j, *and* C_j *is a* **direct superclass** *of* C_i *if* $\nexists C_k \in C$ *s.t.* $k \neq i \neq j$, C_j *is a* **superclass** *of* C_k, *and* C_k *is a* **superclass** *of* C_i.

Definition. *An* **object schema** *is a rooted directed acyclic graph* $G = (V, E)$, *where* V *is a finite set of vertices and* E *is a finite set of directed edges. Each element in* V *corresponds to a class* $C_i \in C$, *while* E *corresponds to a binary relation on* $V \times V$ *that represents all direct* is-a *relationships between all pairs of classes in* V. *In particular, each directed edge* $e \in E$ *from* C_1 *to* C_2, *denoted by* $\mathrm{edge}(C_1, C_2)$, *represents the relationship* C_1 is-a C_2. *Two classes,* $C_i, C_j \in C$, *share a common property if and only if they inherit it from the same superclass. The designated root node, representing the* **Object** *class, is a class defined to have a global extent of all database object instances and an empty type description.*

```
run
BaseClass createSubclass: #Person
    supers: #[ Root ]
    instVarNames: #( #birthyear #name )
    constraints: #[ #[birthyear, Integer],
                    #[name, String] ]
%

run
Person compileAccessingMethodsFor:
             #( #birthyear #name).
%
```

FIG. 2. Code to create a *Person* class.

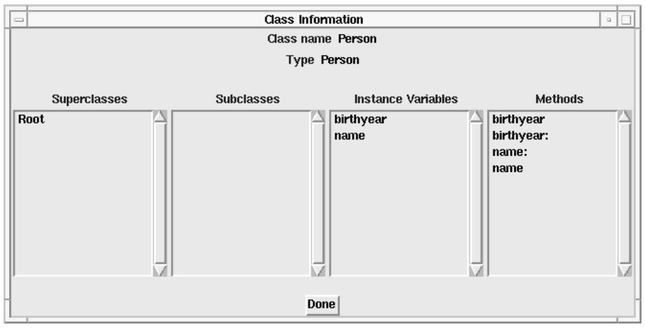

FIG. 3. The *Person* class has instance variables and accessing methods for `name` and `birthyear`.

FIG. 4. Using the *MultiView* graphic interface to create a new instance of the *Person* class.

```
run
BaseClass createSubclass: #Student
    supers: #[ Person ]
    instVarNames: #( #gpa )
    constraints: #[ #[gpa, Number] ]
%
run
Student compileAccessingMethodsFor: #( #gpa ).
%

run
BaseClass createSubclass: #Staff
    supers: #[ Person ]
    instVarNames: #( #title )
    constraints: #[ #[title, String] ]
%
run
Staff compileAccessingMethodsFor: #( #title ).
%
```

FIG. 5. Code defining *Student* and *Staff* as subclasses of the *Person* class.

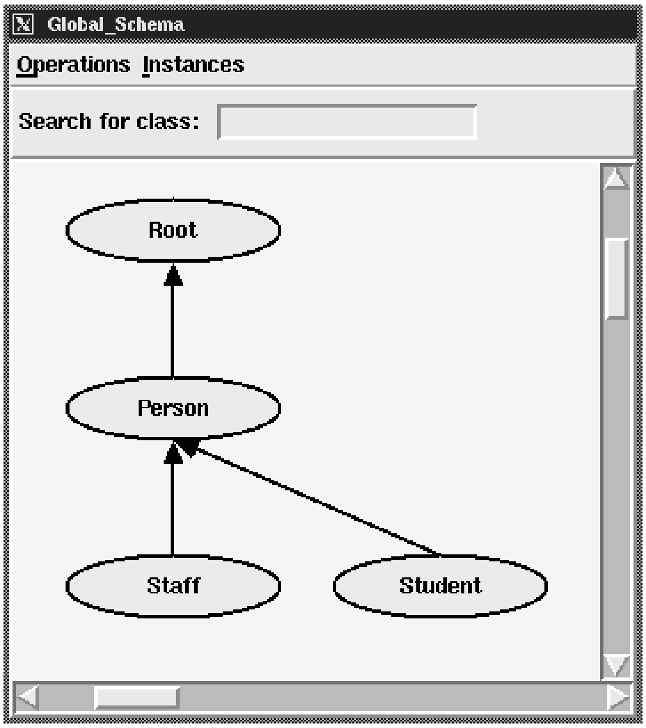

FIG. 6. The new classes are automatically integrated into the Global Schema.

hide	syntax	<*virtual*-class>) := (<*source*-class>) **createHideClass:** [<new-class-name>] **hideMethods:** [<prop-functions>])
	semantics	**type**(<*virtual*-class>) := {p ∈ P | p∈**properties**(<*source*-class>) ∧ p∉<prop-functions>} **extent**(<*virtual*-class>) := **extent**(<*source*-class>)
	class rels	<*source*-class> ⪯ <*virtual*-class> <*source*-class> ⊆ <*virtual*-class> <*source*-class> *is-a* <*virtual*-class>
refine	syntax	<*virtual*-class> := (<*source*-class>) **createRefineClass:** [<new-class-name>] **withProperties:** [<prop-function-defs>]
	semantics	**type**(<*virtual*-class>) := {p∈P | p∈**properties**(<*source*-class>) ∨ p∈<prop-function-def>} **extent**(<*virtual*-class>) := **extent**(<*source*-class>)
	class rels	<*virtual*-class> ⪯ <*source*-class> <*virtual*-class> ⊆ <*source*-class> <*virtual*-class> *is-a* <*source*-class>
select	syntax	<*virtual*-class> := (<*source*-class>) **createSelectClass:** (<*new-class-name*>) **query:** (<predicate>)
	semantics	**type**(<*virtual*-class>) := **type**(<*source*-class>) **extent**(<*virtual*-class>) := {o∈O | o∈<*source*-class> ∧ <predicate>(o)=true}
	class rels	<*virtual*-class> ⪯ <*source*-class> <*virtual*-class> ⊆ <*source*-class> <*virtual*-class> *is-a* <*source*-class>
union	syntax	<*virtual*-class> := (<*source*-class1>) **createUnionClassWith:** (<*source*-class2>) **named:** (<*new-class-name*>)
	semantics	**type**(<*virtual*-class>) := **type**(<*source*-class1>) ⊓ **type**(<*source*-class2>) **extent**(<*virtual*-class>) := {o∈O | o∈<*source*-class1> ∨ o∈<*source*-class2>}
	class rels	<*source*-class1> ⪯ <*virtual*-class> ∧ <*source*-class2> ⪯ <*virtual*-class> <*source*-class1> ⊆ <*virtual*-class> ∧ <*source*-class2> ⊆ <*virtual*-class> <*source*-class1> *is-a* <*virtual*-class> ∧ <*source*-class2> *is-a* <*virtual*-class>
intersect	syntax	<*virtual*-class> := (<*source*-class1>) **createIntersectClassWith:** (<*source*-class2>) **named:** (<*new-class-name*>)
	semantics	**type**(<*virtual*-class>) := **type**(<*source*-class1>) ⊔ **type**(<*source*-class2>) **extent**(<*virtual*-class>) := {o∈O | o∈<*source*-class1> ∧ o∈<*source*-class2>}
	class rels	<*virtual*-class> ⪯ <*source*-class1> ∧ <*virtual*-class> ⪯ <*source*-class2> <*virtual*-class> ⊆ <*source*-class1> ∧ <*virtual*-class> ⊆ <*source*-class2> <*virtual*-class> *is-a* <*source*-class1> ∧ <*virtual*-class> *is-a* <*source*-class2>
difference	syntax	<*virtual*-class> := (<*source*-class1>) **createDifferenceClassWith:** (<*source*-class2>) **named:** (<*new-class-name*>)
	semantics	**type**(<*virtual*-class>) := **type**(<*source*-class1>) **extent**(<*virtual*-class>) := {o∈O | o∈<*source*-class1> ∧ o∉<*source*-class2>}
	class rels	<*virtual*-class> ⪯ <*source*-class1> <*virtual*-class> ⊆ <*source*-class1> <*virtual*-class> *is-a* <*source*-class1>
join	syntax	<*virtual*-class> := (<*source*-class1>) **createJoinWith:** (<*source*-class2>) **named:** (<*new-class-name*>) **withMethods:** (<*aMethodArray*>) **joinAtt1:** (<*property*>) **joinAtt2:** (<*property*>)
	semantics	**type**(<*virtual*-class>) := newly-generated type **extent**(<*virtual*-class>) := newly-generated extent
	class rels	<*virtual*-class> *is-a* <*Root*>

FIG. 7. The Object Algebra Operators: Syntax, Semantics and Class Relationships.

Figure 6 shows the result of executing the code shown in Figure 5 in *MultiView*, thereby defining the *Student* and *Staff* classes as subclasses of our original *Person* class. *MultiView* will update the schema incorporating each new class into the global schema.

2.3 Views in OODBs

Now we extend the basic definitions given above to include the concept of virtual classes and virtual schemas. In the relational model, a **view,** or **virtual table**, is defined to be a **named, stored query**. Similarly, our object-oriented model permits the application of a query operator to a **source class** (or classes) that restructures the source class's type and/or extent membership in order to form a **virtual class** with a type and extent derived from its source class(es).

2.3.1 View-defining Queries Let Q be the set of all possible queries. We constrain a **query** $Q_i \in Q$ used to define a virtual class to correspond to a single algebra operation. Figure 7 displays the view-forming queries currently supported by our model. If a complex query is specified by nesting algebra operators, then each intermediate algebra operator generates a separate virtual class [35].

Definition. *A* **query** $Q_i \in Q$ *includes the following components:* <**QueryType,SourceClasses, TypeModifiers, MemberPredicates, ValuePredicates**>, *where QueryType* $\in$ *{Select, Hide, Refine, Union, Intersect, Difference}; SourceClasses* $\in C$ *is an ordered list of the classes from which the class is derived (that is, the classes to which the query is applied); TypeModifiers defines the derivation relationship between the type of the virtual class and that of the source class from which it is derived; MemberPredicates is an ordered list of the class membership requirements that apply to instances of the virtual class, listed in the same order as the classes in SourceClasses to which they correspond (e.g., a difference query is associated with the constraint that qualifying objects be members of one source lass and not be members of the other source class); and ValuePredicates is the set of simple value predicates that constrain the extent of the virtual class. We refer to the query that defines a virtual class,* $VC_i \in VC$, *as* **query**(VC_i).

```
run
Person createSelectClass:
                    #YoungPerson
  query:

   [ :a | a birthyear > 1975 ].
%
```

FIG. 8. *MultiView* definition of a new *YoungPerson* virtual class.

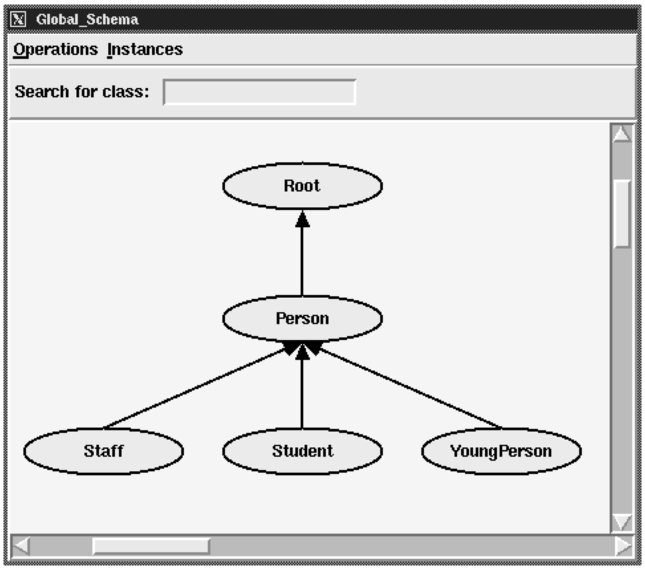

FIG. 9. *MultiView* automatically integrates the new virtual class into the schema.

Currently we support the object algebra shown in Figure 7 for the derivation of virtual classes. These queries allow us to determine the methods, instance variables, and extent of the virtual classes. Now suppose we were to use a *select* query upon class *Person* to define a new virtual class, *YoungPerson*: **createSelectClass:** `YoungPerson` **query:** `[ :person | person birthyear > 1975]`. We can derive all aspects of the new virtual class from this query. The new class's name will be *YoungPerson*, its type will have the same properties as that of the *Person* class, **SourceClasses**(*YoungPerson*) will contain only the *Person* class, **MemberPredicates**(*YoungPerson*) = ($\in$ *Person*), and **ValuePredicates**(*YoungPerson*) will contain only the simple predicate `(birthyear > 1975)`. After the *YoungPerson* class has been created, *MultiView* will then integrate the new virtual class into the class hierarchy and add the appropriate edge $(YoungPerson, Person)$ to the object schema's edge set. The new object schema is shown in Figure 9 (right).

2.3.2 Base Classes, Virtual Classes, and View Schemata

Let C be the set of all classes in the database. C is partitioned into two sets – BC is the set of all base classes in the database, and VC is the set of all virtual classes in the database. **Base** classes are defined during the initial schema definition. **Virtual** classes can be defined and added dynamically to the schema throughout the lifetime of the database. They are defined by the application of a query operator to a *source class* (or classes) that restructures the source class's type and/or extent membership.

Just like our original class definition, each database class (whether virtual or base) $C_i \in C$ has a unique class name, a type description, and a set membership. In addition, in order to accommodate the support of virtual classes, each class C_i also maintains a set of all virtual classes $VC_j \in VC$ s.t. $C_i \in$ **SourceClasses**(VC_j). We use the corresponding notation VC_j **derived-from** C_i to characterize the relationship $C_i \in$ **SourceClasses**(VC_j).

Let VC be the set of all virtual classes. Like members of the set of base classes, a **virtual class** $VC_i \in VC$ has a unique class name, a type description, and a set of the virtual classes derived from VC_i. In addition, because the set membership of VC_i is **derived** using its derivation query, VC_i also includes the query from which it is derived, denoted by **query**(VC_i), as defined in Definition 2.3.1.

Definition. *We define the* **Derived-from Sub-Graph** *of a class* $C_i \in C$ *to be a schema* $DS(C_i) = (DV, DE)$ *where each class* $DV_j \in DV$ *corresponds to a class that is either directly or indirectly derived from* C_i, *and each edge* $DE_k \in DE$ *represents a direct derivation relationship between two classes in* DV:

1. C_i *is the root of the DS.*
2. $DV \subseteq V$*,*
3. $\forall V_j \in V$ *s.t.* V_j *corresponds to virtual class* $VC_j \in VC$*,* $DV_j \in DV$ *if and only if* VC_j **derived-from** C_i *or* $\exists V_k \in DV$ *s.t.* VC_j **derived-from** VC_k*.*
4. *Each edge* $DE_i = \langle DV_i, DV_j \rangle \in DE$ *corresponds to a* **direct derivation relationship** *between two classes* C_i *and* $C_j \in C$ *with corresponding nodes* DV_i *and* $DV_j \in DV$*, meaning that* C_j **derived-from** C_i*.*

In Definition 4, we define an object schema as a database schema containing all of the classes in the database. We now extend this definition to define a **global schema** as a database schema containing all of the database classes, both base and virtual, integrated into a consistent generalization hierarchy; a **base schema** as a database schema containing all base classes; and a **view schema** as a database schema containing a user-selected subset of the classes (either base or virtual) from the global schema. A view management system must support a flexible classification mechanism in order to maintain this global class hierarchy. For example, the system must be able to make dynamic changes to the class hierarchy, possibly inserting a new class between two existing classes without affecting the types or extents of previously-existing classes in the hierarchy. We have proposed elsewhere [38, 35] algorithms and techniques by which the global class hierarchy can be maintained.

Virtual classes are often simply called "views," but we interpret an object-oriented view to imply a complete **view schema**, which is a user-selected subgraph of the global schema. At any time during the lifetime of the database, users can create new view schemata or add base and/or virtual classes from the global schema to existing view schemata. Because all classes are fully integrated into the global schema, *MultiView* easily maintains the correct subsumption relationships between classes in view schemata.

Definition. *Given a global schema* $GS = (V, E)$*, a* **view schema***,* VS*, is defined to be a schema* $VS = (VV, VE)$ *with the following properties:*

1. VS *has a unique view identifier denoted* $< VS >$*,*
2. $VV \subseteq V$
3. $VE \subseteq$ transitive-closure(E)

Because view schemata have all the properties of the global schema, and because classes can participate in view schemata under assumed names, different applications can use their own views of a shared database. For example, we can create a new view schema, named *YoungCultureView*, and select the *YoungPerson* and *Student* classes to participate in that schema under the names *GenerationXer* and *Student*, as shown in Figure 10. The *MultiView* system will automatically build the class hierarchy for the new view schema. Figure 11 shows the the resulting view schema.

2.3.3 View Materialization

Typically, the contents of a view are not stored, but are instead derived using the stored query defining the view. The term **view materialization** means that results of the query that defines a view are actually maintained in the view's extent, as opposed to being computed on demand. Materialized views have already been established in the relational context as offering improved query times. We define

```
run
   YoungCultureView
       addVC: YoungPerson
       withName: #GenerationXer
%

run
   YoungCultureView
       addVC: Student
       withName: #Student
%
```

FIG. 10. Code to add classes to the *YoungCulture* virtual schema.

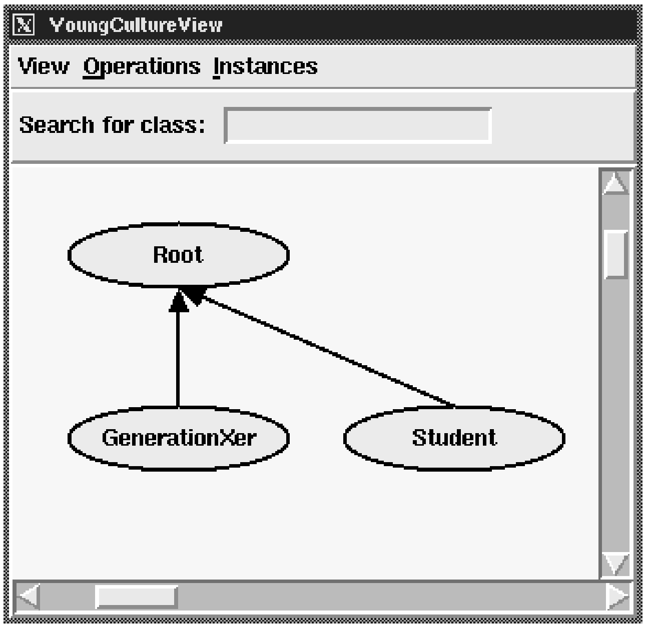

FIG. 11. The *YoungCultureView* schema contains a selected set of classes organized into a schema by *MultiView*.

a **materialized virtual class** as a virtual class that stores either copies of or references to data instances in its extent rather than computing its extent upon access. For example, if the *YoungPerson* virtual class in Figure 9 were materialized, then it would have an extent associated with it that would contain the set of *Person* objects with a birthyear greater than 1975. The primary problem of view materialization is that the extent of a materialized virtual class must be maintained in a consistent state with respect to updates by the database system. For example, should an instance of the *Person* class have its birthyear instance variable updated from a value of '1980' to '1960', then that instance must be removed from the extent of the *YoungPerson* virtual class.

3. *MultiView* Design Principles

In the following sections we describe and discuss our system design and its components, as illustrated by the concentric circles of Figure 1. We begin in this section by reviewing the properties required for the support of our view model's features. We next review the properties provided by the underlying GemStone system. Finally, we compare the required properties to those that are provided.

3.1 *Properties Required for View Support*

MultiView's object model includes some fundamental design principles. These principles require the underlying database system to support certain properties. Figure 12 outlines the design principles and the properties needed to support them.

- **A class has a type and an extent.** Objects should possess both virtual and base types. For example, if an object qualifies for membership in two select classes, then it should be a member of both—regardless of whether or not any subsumption relationship exists between the select classes.

Design Principles \ Required Properties	multiple classification	multiple inheritance	dynamic class restructuring	dynamic object restructuring	flexible inheritance system
class has type and extent	✓			✓	
dynamic class creation		✓	✓		✓
integrated class hierarchy	✓	✓		✓	✓
virtual class first-class database citizen	✓	✓	✓	✓	✓
objects gain and lose types dynamically	✓			✓	

FIG. 12. Properties required for view support

The system must therefore support *multiple classification* and the *dynamic restructuring of objects.*

- **Users should be able to create virtual classes dynamically, customizing both type and extent.** If the new virtual class's type is a customization of an existing class's type, then the correct position in the global type hierarchy might be between two existing classes. The underlying inheritance implementation must therefore be flexible enough to support the *dynamic reclassification* of database classes. Furthermore, if the new virtual class is to serve as the point of inheritance for a new property, that property could be relocated to the new class. Our system must therefore support *dynamic class restructuring.* This dynamic class restructuring includes support for the possible migration of state as properties are moved from class to class.
- **Both virtual and base classes should be integrated into a correct and consistent hierarchy.** As detailed in [37], this integration requires both *multiple classification* and *multiple inheritance.* For example, if a virtual class is formed by applying a union query to two base classes, then the newly-formed union class should be classified as a superclass of both base classes, which should continue to be subclasses of any other classes from which they inherit properties. In addition, because new classes will have to be integrated dynamically, the *inheritance implementation must be flexible* and the underlying system must support *dynamic class restructuring.*
- **Virtual classes should be treated as first-class database citizens.** This means that virtual classes should participate in the upwards-inheritance mechanism of the schema, that base classes inherit from virtual ones, and that the system should support *capacity-augmenting virtual classes.*
- **Objects should be able to gain or lose types dynamically.** This feature requires support for the *dynamic restructuring of objects* because an object could gain or lose state depending on which types it possesses. As a consequence, an object may simultaneously be a member of multiple classes, regardless of whether or not any subsumption relationships exist between the classes.

3.2 *GemStone Kernel*

As illustrated by the innermost circle of Figure 1, *MultiView* is built on top of Servio Corporation's GemStone OODBMS using GemStone's Smalltalk-like OPAL programming environment. We chose to use the GemStone OODBMS rather than implement *MultiView* from scratch because GemStone provides a rich object-oriented data model with supporting tools. Despite significant differences between the GemStone and *MultiView* data models, GemStone offers key features that proved extremely useful in the implementation of *MultiView.* Besides the typical database functionalities, such as persistence, database programming language support, and transaction management, features of GemStone include:

- GemStone provides automatic, system-maintained object identity.
- GemStone treats everything in the system, including code blocks and classes, as objects.
- GemStone offers a number of programming language interfaces, such as C, C++, and Smalltalk, which facilitate the development and integration of a graphic interface.
- GemStone permits access to the source code for most methods, whether system or user defined.

3.3 *Differences Between* MultiView *and GemStone*

Our implementation made a number of extensions to the GemStone system that were necessary in order to support views. For example, because GemStone does not maintain explicit extents to collect all instances of a type, which is needed for the specification of select virtual classes, *MultiView* extends the GemStone concept of class to include an extent. We add system methods to automatically add objects to the appropriate extents upon creation, and maintain the extents of virtual classes upon updates.

More significantly, the goals of *MultiView* require that the underlying object model support certain properties, as discussed in Section 3.1. Many of these properties are not supported by GemStone's data model. The most important of these properties are:

- GemStone does not support *multiple classification*, which is a required characteristic for view support in *MultiView*.
- GemStone does not support *multiple inheritance*, which is necessary both in order to preserve our single-point-of-inheritance property and also the participation of virtual classes as first-class citizens in an upward-inheritance mechanism.
- *Class restructuring* in GemStone is severely restricted for classes with instances, which would prevent the manipulation of the schema hierarchy necessary for the integration of virtual and base classes into a single global schema.
- GemStone supports object restructuring only in terms of migrating objects from one class to another. Furthermore, an object can only belong to one class, and migration can only take place between two classes that share a class history.
- GemStone's native inheritance schema does not support the independent restructuring of class inheritance relationships that is necessary if virtual classes are to participate fully in the inheritance schema.

Multiple classification is particularly necessary in a capacity-augmenting view system, because an object can be an instance of multiple virtual classes (as well as its base class) regardless of whether or not any subsumption relationships exist between those classes. While regular view systems (i.e., that do not support capacity-augmenting views) also must permit an object to be an instance of multiple virtual classes (in addition to its base class), note that virtual classes do not carry any additional stored data – and it is thus trivial to make the object a transient member of the virtual classes on access. This is no longer sufficient for capacity-augmenting views. To the best of our knowledge, current OODB systems do not support multiple classification — with the exception of IRIS [10], which is a functional database system that uses a relational database as a storage system, to store data from one object across many relations[3]. Inspired by this approach we have developed an object-slicing technique to address these advanced features in the context of the OODB technology [32]. In the following section, we describe our object-slicing technique.

4. Object-Slicing Architecture

We used a flexible and powerful technique called object-slicing to construct the *MultiView* object model on top of GemStone. In object-slicing, a real-world object corresponds to a hierarchy of **implementation objects** (one for each class whose type the object possesses) linked to a **conceptual object** (used to represent the object-itself) rather than associating one implementation with each conceptual object as is commonly assumed in OODB systems [28].

For example, Figure 13 depicts a schema composed of two base classes, *Cat* and *HouseCat*, and two virtual classes, *HeavyCat* (derived from a selection query upon the *Cat* class) and *DietingHeavyCat* (derived by refining *HeavyCat* to add a new instance variable, `diet`). The right-hand side of the figure illustrates a single instance of the *HouseCat* class with instance variables `height = 9"`, `weight = 17lbs`, and `owner = Fred`. The instance fulfills the membership requirements of *HeavyCat* (weight > 15 lbs), and thus it belongs to both the *HeavyCat* and the *DietingHeavyCat* virtual classes. Note that the state of the object is maintained by the implementation objects of the classes defining each property. Because the instance belongs to *DietingHeavyCat*, it now carries a data value for the `diet` attribute.

We formalize the object-slicing paradigm as follows. Let $O_i \in O$ be a user-defined database object. In the object-slicing model, O_i is represented using two kinds of objects: a single conceptual object and one implementation object for each class $C_j \in C$ to which the object belongs. We refer to object O_i's conceptual object as $O_{i_{conc}}$ and refer to the implementation object representing object O_i's membership in class C_j as $O_{i_{impl_{C_j}}}$. Each implementation object $O_{i_{impl_{C_j}}}$ is an object instance of the database class $C_j \in C$ it represents. Conceptual objects are object instances of a special system class, *ConceptualObject*, rather than of a user-defined class. Because a single real-world object is now represented using multiple database objects, we define a number of functions to operate on objects in the model, including object creation, equality comparison, etc.

Definition. *A* **conceptual object** *consists of a tuple, <OID, implementationObjects>, where OID is the unique, system-generated object-identifier of the conceptual object, and implementationObjects is the set of* **implementation objects** *that are linked to the conceptual object. An* **implementation object** *is a 4-tuple* <**OID,oid,class,state**> *where OID is the object-identifier of the linked conceptual object, oid is the object-identifier of the implementation object itself*[4]*, class is the class of which the implementation object is a direct instance, and state corresponds to the values of the local instance variables stored for the given object.*

4.1 Implementation of Required Properties Using Object-Slicing

The object-slicing architecture offers the advantages of a flexible and powerful implementation base that can readily be configured to support all of the properties we require for the support of the *MultiView* model.

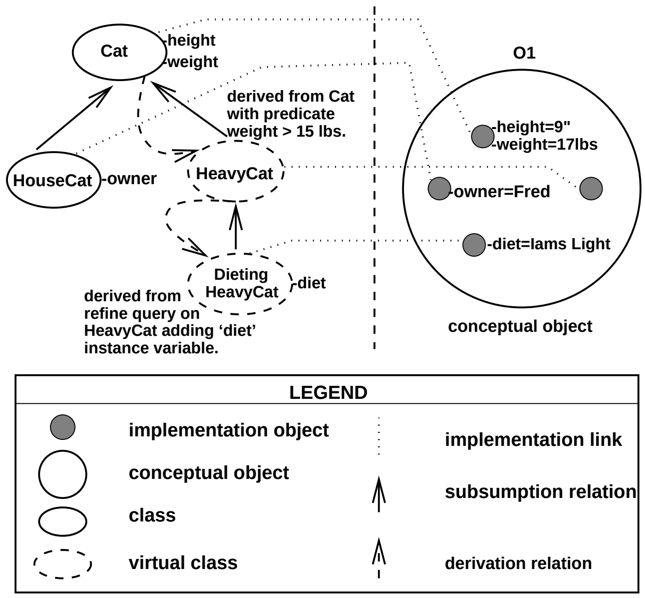

FIG. 13. Example of object-slicing.

MULTIPLE CLASSIFICATION.

The object-slicing approach implements multiple classification by using implementation objects to represent an object's membership in multiple classes. This technique of representing class membership via implementation objects is quite flexible. For example, an object-slicing system can dynamically reclassify an object from being the instance of one class ($C1$) to becoming that of another class ($C2$), by linking the object instance to an implementation object of the class $C2$ and discarding that of the class $C1$. The effect of reclassification under object-slicing is minimal. No new classes are created using this approach, nor is the type of any existing class affected by the reclassification. Furthermore, because objects are compared using the oid of their conceptual objects, the multiply-classified object's identity is not compromised.

CUSTOMIZING INHERITANCE USING OBJECT-SLICING.

If an object O_i possesses an implementation object $O_{i_{impl_{C_j}}}$ for some class $C_j \in C$, then O_i must also possess implementation objects for all classes C_k s.t. C_j *is-a* C_k. Thus the implementation objects associated with a given conceptual object will mirror the structure of the class hierarchy. Object-slicing intrinsically includes its own inheritance mechanism. Let there be an implementation object $O_{i_{impl_{C_i}}}$, $O_i \in O, C_i \in C$. Suppose the method $m_i \in M$ were to be invoked upon the implementation object $O_{i_{impl_{C_i}}}$ (which is in the extent of class C_i), and method m_i is not defined locally in the type of class C_i (**type**(C_i)). In this case, the implementation object $O_{i_{impl_{C_i}}}$ will delegate the method m_i to its associated conceptual object $O_{i_{conc}}$, which will then conduct a search "upwards" through **SuperclassHierarchy**(C_i). If method m_i is not found in the type of some class $C_k \in$ **SuperclassHierarchy**(C_i), then an error is returned. Otherwise, if method m_i is found in the type of a class C_k, then the method is invoked upon the appropriate implementation object $O_{i_{impl_{C_k}}}$, which can be located using the function **ImplLink**$(O_{i_{conc}}, C_k)$[5]. Objects of user-defined classes in an object-slicing representation use this object-slicing inheritance mechanism.

Figure 14 shows the process that takes place if one of the implementation objects receives a message for a method that is not locally defined in the implementation object's class. The left side of Figure 14(a) shows a class hierarchy composed of five classes. The right side of Figure 14(a) represents a conceptual object that belongs to all of the classes in the hierarchy. The smaller circles within the object each represent an implementation object for a specific class whose type the object possesses. Assume that an implementation object for the *PedigreedCat* class receives a message for method "color," which is an access method for instance variable "color" defined in *Cat* (Figure 14(c)). Because the receiving implementation object (which is of type *PedigreedCat*) does not understand the method "color," it passes the request to its conceptual object. The conceptual object then searches upwards in the *PedigreedCat* class's superclass hierarchy, and identifies that the *PedigreedCat* class inherits the "color" method from the *Cat* class. The conceptual object then sends the message to its *Cat* implementation object. Because both the instance variable and the accessing method are locally defined at *Cat*, the *Cat* implementation object contains state for the instance variable "color" and can respond to message "color." It thus recognizes the message and performs the appropriate operation, returning the value of "color" to the user (Figure 14(c)).

DYNAMIC RESTRUCTURING OF OBJECTS AND CLASSES.

Recall that our data model supports a *single point of inheritance* characteristic for database properties, meaning that if two classes C_1 and C_2 share some common property then they both must have inherited it from some lowest common superclass (LCS). Not only does this approach support true upwards inheritance of method code for both base and virtual classes, but it also avoids the possible maintenance and storage problems that could be caused by the duplication of code. This approach is compatible with object-slicing, because the point of inheritance corresponds to the class that maintains all implementation objects that have values for the particular instance variable.

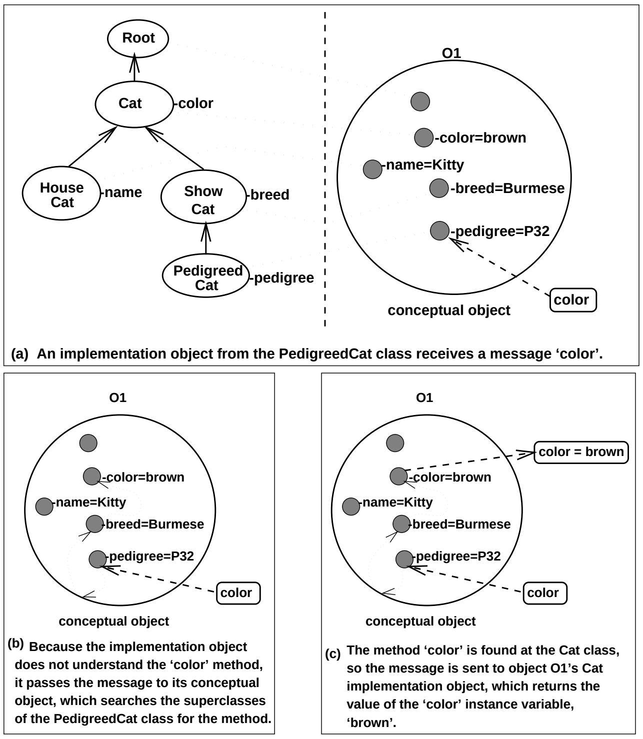

(a) An implementation object from the PedigreedCat class receives a message 'color'.

(b) Because the implementation object does not understand the 'color' method, it passes the message to its conceptual object, which searches the superclasses of the PedigreedCat class for the method.

(c) The method 'color' is found at the Cat class, so the message is sent to object O1's Cat implementation object, which returns the value of the 'color' instance variable, 'brown'.

FIG. 14. Inheritance Under the Object-Slicing Model

In order to preserve uniform upwards inheritance, it may be necessary to *promote* methods and/or instance variables from a subclass to a new superclass during the classification of a new virtual class into the global schema [6]. We call this promotion *property migration*. This will ensure that the property will be located above all the classes that inherit that property. When methods or instance variables are moved from an existing class to a new superclass, the system performs the following tasks: First the new superclass is extended to include the migrating (to be locally defined) methods and/or instance variables. Next, the migrating methods and instance variables are removed from their original class. Finally, each implementation object instance of the original class is *split* into an implementation object of the modified original class and an implementation object of the new superclass [7]. These new implementation objects are linked together by their corresponding conceptual objects. Below we present a result that limits the effort required for property relocation.

Property 1. *If two classes C_i, $C_j \in C$ share some common property p_l, then they must ultimately have inherited it from the same superclass; i.e., there must exist a lowest common superclass in the class lattice $C_k \in C$ s.t. C_i* is-a *C_k, C_j* is-a *C_k, and $p_l \in$* **properties(C_k)**.

This property follows from our classification algorithm, which places each newly-created virtual class into the global class hierarchy in a consistent and correct way [37].

Proposition 2. *Only two classes will be involved in any property migration caused by virtual class integration.*

PROOF.

Suppose that there exist two distinct source classes, *Class A* and *Class B*, and that there are two methods, "a" and "b", respectively located in *Class A* and *Class B*, as shown in Figure 16. Because we do not allow cycles, the relationship between *Class A* and *Class B* cannot be ambiguous – class *Class A* can be a superclass of class *Class B* (in which case *Class B* would inherit method "a" from *Class A*), as shown in Figure 16(a), or class *Class B* can be a superclass of class *Class A* (in which case *Class A* would inherit method "b" from *Class B*), as shown in Figure 16(b), but one class cannot be both superclass and subclass of the other, Figure 16(c). Therefore, by Property 1, at least one class of *Class A* and *Class B* must not have inherited any methods from the other class. Because the only way that a class could include method "a" or "b" in its type

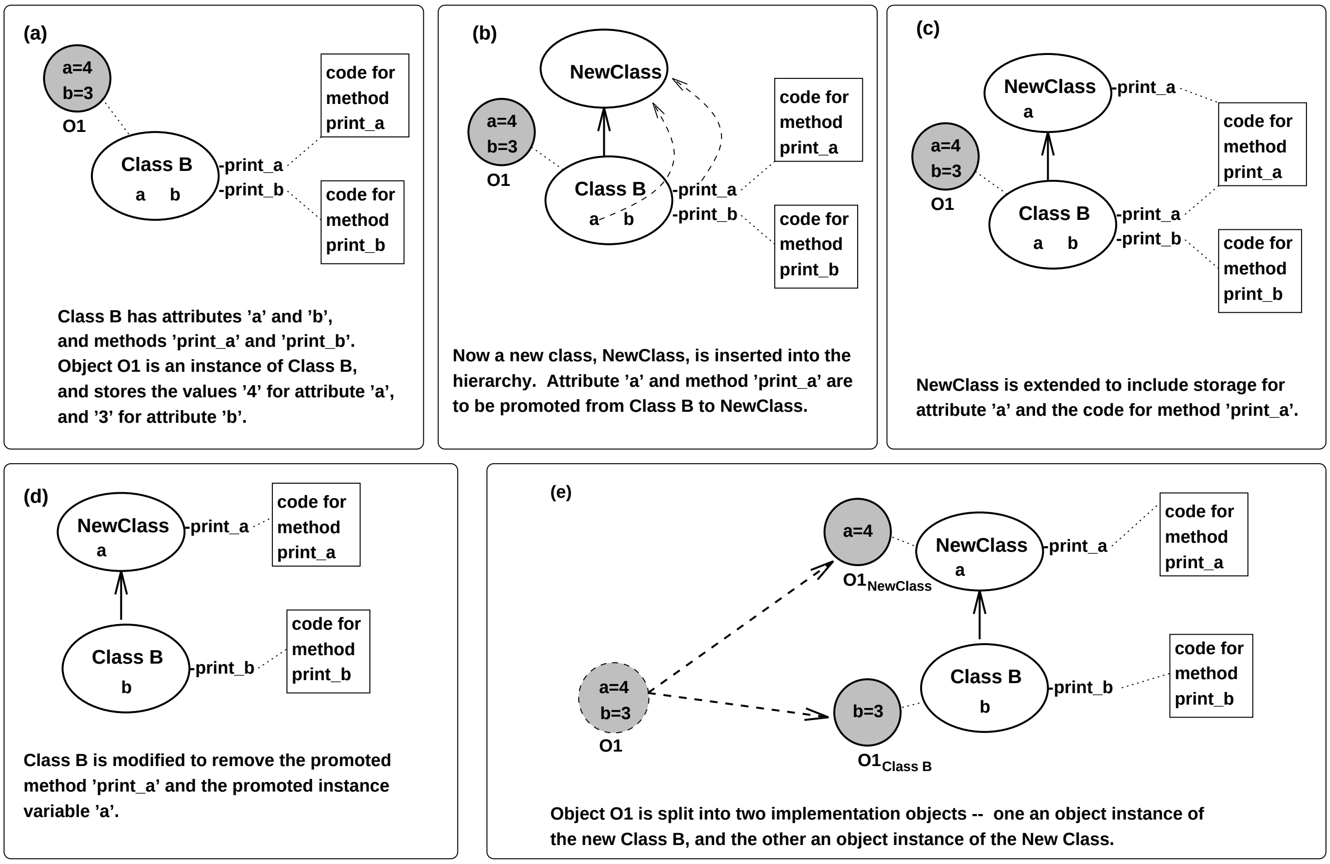

FIG. 15. Attribute a and method print_a are moved from class B to a new superclass.

would be to have inherited them from class *Class A* or *Class B*, and because *Class A* and *Class B* cannot both inherit from each other, at most one of classes *Class A* and *Class B* could include both methods “a” and “b”.

Now suppose, contrary to fact, that there were a new class such that two methods, “a” and “b”, from different source classes, *Class A* and *Class B*, were to be relocated in the new class. Property 1 states that every method is located in exactly one place; thus both of the source classes would now be subclasses of the new class.

If both source classes are subclasses of the new class, then each of them must inherit both “a” and “b” from the new class. However, we already established that it is not possible that both *Class A* and *Class B* include both methods “a” and “b”. Therefore, there cannot be a new class such that two methods from different source classes are relocated in the new class.

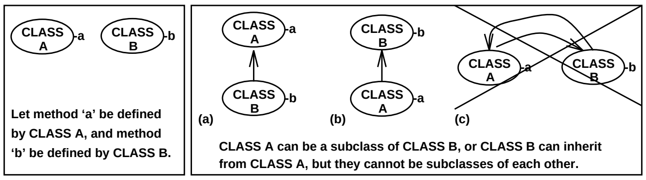

FIG. 16. Cycles cannot exist in the class hierarchy.

OBJECT-SLICING AND VIEW MATERIALIZATION.

The object-slicing technique, combined with the characteristic of *single point of inheritance* and *classification*, lends itself well to materialization, because (1) it elegantly avoids the need to duplicate data for materialized classes and (2) any update to an object will take place at a location determined by the property involved regardless of the source of the update request. Since object-slicing requires an object to possess implementation objects for any class to which it belongs, an object-slicing view implementation is effectively materialized in that the implementation objects representing an object's membership in various virtual classes are actual object-instances of those virtual classes. The object-slicing technique is thus subject to the intrinsic problem of view materialization – how to update materialized views so as to keep their extents consistent with the rest of the database. This problem and its solutions are addressed by the *MultiView* system's update model, which is described in Section 5.

EVALUATION OF OBJECT-SLICING'S IMPACT ON PERFORMANCE

Extending an existing DBMS with object-slicing techniques necessarily involves the overhead of additional data structures, maintenance costs, and processing time. Although object-slicing is a known technique that is being utilized for view systems [26], schema evolution [33], and role systems [11], to our knowledge no other work has been done evaluating the costs of object-slicing. Consequently, we were motivated to perform some initial experiments evaluating the relative costs and benefits of adopting the object-slicing techniques. The complete results of these experiments are presented in [23].

In evaluating the performance of databases, I/O operation time typically dominates CPU operation time. Consequently, an evaluation of object-slicing must consider the effect of object-slicing on I/O time. One major variable for calculating I/O time is the number of objects that can be stored in a disk block, known as the *blocking factor* (bf). Because the object-slicing model represents any given object using a conceptual object and some number of implementation objects, the object-slicing architecture inherently requires more storage to store an object's state than would be needed in a traditional architecture. However, despite the extra overhead, the space needed to store individual implementation objects can be smaller than that needed to store conventional objects, and thus the blocking factor for storing the implementation objects in the extent of a class under object-slicing can be superior to that for storing the entire instances of the same class under a conventional architecture. Our experiments confirm that due to the superior blocking factor of object-slicing can actually improve query performance in the case of queries upon local properties.

However, another significant factor in database performance is data access costs—specifically, accessing data items sequentially provides superior performance to accessing data items with random block accesses. Retrieval of inherited attributes in object-slicing can involve significant overhead. Because under object-slicing the implementation objects belonging to a given class are by default clustered together, a traversal to an implementation object in order to retrieve an inherited value may require a random block access (thereby resulting in a page fault) and impair performance. We have designed and carried out an extensive experimental study evaluating clustering techniques to ameliorate this cost [23].

5. *MultiView* View System

MultiView supports customized virtual classes and customized virtual schemata in such a way that the system's class-restructuring capabilities are powerful and flexible, and that the virtual classes and virtual schemata look and feel like the actual database. The *MultiView* system automates as many of the database maintenance tasks as possible, including the integration of new classes into the global schema, the construction of the class relationships in view schemata, and the maintenance of class extents. The *MultiView* OPAL interface preserves all the functionalities of the original GemStone OPAL system, while providing new view capabilities. The subsections of the third circle of Figure 1 illustrate the four major functions performed by the *MultiView* view system:

Virtual class creation / integration. *MultiView* users can customize existing type structures and object sets (class extents) by deriving virtual classes via object-oriented queries. *MultiView* will then automatically integrate the new virtual class with existing classes into a single consistent global schema graph, maintaining relationships between base and virtual classes.

View schema support. *MultiView* users can specify (at any time) view classes (both base and virtual) from the augmented global schema for membership to a particular view. *MultiView* will then construct arbitrarily complex view schemata composed of these view classes.

Queries. Because *MultiView* is completely compatible with the underlying GemStone system, any set-based query that can be performed using GemStone's OPAL interface can also be performed by *MultiView*. In addition, *MultiView* also supports object algebra queries over class extents and query-by-form via the graphic interface (See Section 6).

Data manipulation. *MultiView* supports updates on both base and virtual classes. Because *MultiView* supports materialized views, the system incrementally maintains materialized virtual class extents in the face of the entry and update of data objects (using efficient update propagation algorithms).

5.1 *Virtual Class Creation and Integration*

In this subsection, we focus on the implementation of the first task—the specification and classification of virtual classes in *MultiView*.

VIRTUAL CLASS CREATION.

MultiView provides a number of class methods for virtual class maintenance that implement the query operators shown in Figure 7. These class methods automatically create the new virtual class, generate its type (migrating all appropriate properties), construct its extent (preserving all property values), and integrate the new class into the global class hierarchy.

VIRTUAL CLASS CLASSIFICATION.

In *MultiView*, all virtual classes are automatically integrated into one global generalization hierarchy. Such complete classification offers a number of advantages, including (1) the relationships between base and derived classes are made explicit (which aids query processing and view materialization strategies); (2) because virtual classes fully participate in the inheritance schema, property functions and object interfaces can be shared; and (3) it facilitates the formation of arbitrarily complex view schema graphs. Both the classification algorithm and its proof of correctness are presented in [37]. Classification in *MultiView* is automatic and compulsory. That is to say, every time a class is added to a *MultiView* database, the system automatically integrates it into the inheritance hierarchy. *MultiView* makes classification the responsibility of the system rather than requiring users to create their own *is-a* arcs manually because we want the global class hierarchy to be unique, correct, and consistent.

There are two steps to the classification algorithm: first the class hierarchy is prepared by the addition of any necessary intermediate classes, then the virtual class is placed into the global schema. The *MultiView* view creation process includes the invocation of methods which use deterministic algorithms to perform the automatic integration of virtual classes into the global schema.

In [37] we described the problem of how sometimes there may be no correct location for the placement of a new virtual class in an existent global schema graph – in order to allow for both the full inheritance invariant as well as subsumption. The current implementation of the *MultiView* system includes this classifier, further described in [37], which generates one or more intermediate classes in order to guarantee these properties. This solution is both necessary and sufficient to guarantee the closure of the resulting class hierarchy. That is to say, integrating the newly generated classes into the global schema will never cause the generation of additional new classes.

Once the class hierarchy has been prepared by the insertion of any necessary intermediate classes, the new virtual class can be inserted into the global schema by identifying all direct *is-a* relationships between the new class and the other classes using a single depth-first downwards traversal algorithm. We have shown elsewhere that the resulting global schema incorporates the virtual class in a consistent and most efficient manner [37].

5.2 View Schemata

As proven in [36], once the global schema integration has been achieved, the tasks of specifying and constructing view schemata can be reduced to simple graph algorithms. *MultiView* provides a simple graphical interface for the definition of view schemata (See Section 6. The user can create any number of view schemata, and can add both virtual and base classes to a view schema as *view classes* (see Figure 11). The *MultiView* view management module automatically computes the proper hierarchical relationships for classes in a view schema. View classes can have different names in different view schemata. Each view schema is associated both with an access control list of GemStone users who are allowed to access the view, and with a password that controls access to this access control list.

5.3 Query Capabilities

The *MultiView* OPAL interface is fully OPAL compatible, and thus users can employ all queries supported by the native GemStone system. In addition, because the *MultiView* model associates an extent with each class, users can explicitly perform queries upon class extents. *MultiView* class extents are implemented as a subclass of the OPAL *Set* class, and thus all *Set* methods can be invoked upon class extents, including iteration, collection, and selection. Note that because inserting or deleting an object from a class's extent can have an impact upon the extent content of other classes, we overrode the adding and removing methods to perform the necessary update propagation. *MultiView*'s query capabilities exceed GemStone's, as the *MultiView* object algebra supports declarative queries across both base and virtual classes. Section 6 presents examples of *MultiView*'s query functionalities as captured by the *MultiView* graphic interface, called the *MultiViewer*.

5.4 Data Manipulation and Associated View Maintenance

MultiView supports a suite of update operations, as well as the underlying view update propagation strategies necessary to keep materialized classes consistent [26, 27]. The update operators *MultiView* supports include creation, deletion, addition of a type, removal of a type, and modification of an instance variable.

Since *MultiView* maintains materialized views, the effects of updates must be propagated to all classes affected by it. For example, if an object gains the type of a class that has a union class derived from it, then one direct effect of that update would be that the object gains the type of the union class (i.e., a new implementation object of the union class is created and linked to the updated object).

We identify specific techniques by which we can exploit unique features of the object-oriented paradigm to optimize updatable materialized views [26, 27]. Our update algorithms are incremental, and perform selective notification to the set of classes determined to be directly-affected by a given update. We use inherent object-oriented features, such as the integration of classes into a generalization hierarchy, encapsulation, and membership materialization, both to facilitate the identification of the classes that are directly-affected by an update and also to optimize the propagation to the classes derived from classes directly-affected by the update. Our update optimization process performs two tasks:

1. First, we identify the set of classes that are directly-affected by the update, i.e., the set of classes whose set-membership changes as a direct result of the update.
2. We propagate to the classes derived from the directly-affected classes in an optimized manner, so as to avoid both propagation to classes that turn out not to be affected by the update and also self-cancelling update propagation.

IDENTIFYING DIRECTLY-AFFECTED CLASSES.

In order to be able to efficiently identify the set of classes directly-affected by modification updates, we introduce a "registration" service by which virtual classes can register their interest in specific properties and be notified upon modification of those properties. We are aided to this end by our data model's injunction that each property must reside at a single class that defines the property and acts as its **point of inheritance**. All classes whose membership may be affected by modifications to an attribute are thus registered with the attribute's **point of inheritance**. We associate *triggers* with the accessing methods to propagate updates to the registered classes if and only if the relevant property's value is modified. The principles of our registration strategy are:

- When a virtual class is created, its predicate is parsed and the class is *registered* with each property involved in the query specification. Each registration is sent to the class that serves as the point of inheritance for the property involved. Because all classes in our model are arranged in a generalization hierarchy, we optimize our registration process to take advantage of subsumption relationships. The *registration* structures can thus be distributed along the generalization hierarchy, as illustrated in Figure 17. Distributing the *registration* structures allows us to avoid evaluating the update in terms of classes that are obviously unaffected by the update (i.e., classes derived from classes to which the object did not belong either before or after the update).
- When an object is updated, the update method [8] triggers a *notification* function that propagates the update to all virtual classes that have *registered* an interest in the value of the updated property.
- When a virtual class has been informed, via the *notification* class method, that an object in which it has a potential interest has been modified, it then evaluates the update. Should the update warrant that the object gain, keep, or lose the type of the class, then the class will send a message notifying those dependent upon it of the change. Otherwise it does not propagate the change across its derivation hierarchy and the update propagation terminates.

ELIMINATING SELF-CANCELLING UPDATE PROPAGATION

The propagation process could involve updates that cancel out each other's effects. We call this the *self-cancelling update propagation* problem because it could lead to repetitive evaluations of the object in the context of a single class, sometimes with self-cancelling results. For example, Figure 18(a) presents a schema vulnerable to the self-cancelling problem. Figure 18(b) illustrates how an update to the *weight* of an instance of the *Cat* class could cause it to first gain, then lose, the type of the *DenseCat* class. From this example it follows that a self-cancelling propagation problem could extend to a potentially huge set of classes to the detriment of efficiency.

In order to overcome this problem, *MultiView* incorporates the following strategy. First, we assign a unique **derivation number**, denoted by DerNum(), to each class in the database, indicating its global derivation depth in the schema. Using this derivation number framework, we propagate changes to affected classes using a breadth-first traversal by order in the derivation hierarchy. Because a derived class will only be processed after all of its source classes have been processed, we can guarantee that no class will be processed more than once. This solution avoids the problem of self-cancelling update propagations by ensuring that if DerNum(C_k) $<$ DerNum(C_l), then a modified object's membership in C_l can be evaluated only after the evaluation of the object's membership in C_k has been completed. Figure 18(c) illustrates how this strategy avoids the pitfall illustrated in Figure 18(b).

IDENTIFYING BRANCH TERMINATION CONDITIONS

Under certain conditions, propagation to branches of the derivation or generalization hierarchy can be terminated. For example, propagation to a branch of the derivation hierarchy can be terminated if the object does not qualify to belong to any of the entrance points into the branch either before or after the update. That

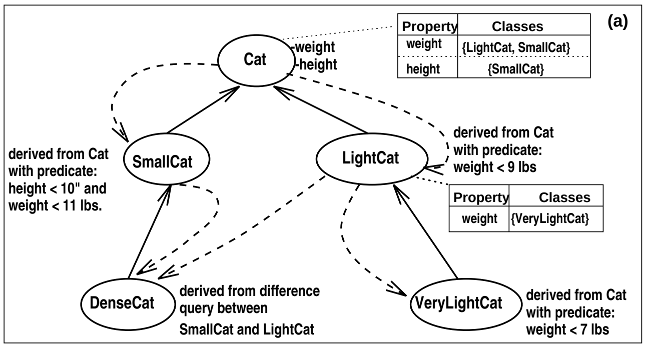

FIG. 17. Registrations use both point of inheritance and subsumption.

is to say, given an updated object O_i and a virtual class VC_j, if O_i does not qualify to belong to VC_j either before or after the update, then O_i does not qualify to belong to any virtual class VC_k s.t. VC_k is in the **Derived-from Subgraph** of VC_j and $\nexists VC_l$ s.t. O_i qualifies to belong to VC_l either before or after the update and VC_k is in the **Derived-from Subgraph** of VC_l. For example, because the *VeryLightCat* class is subsumed by the *LightCat* class, if an instance of the *Cat* class shown in Figure 18 were to change the value of its *weight* instance variable from 12 to 10 lbs, then our update algorithm should determine that because the instance remains unqualified for membership in the *LightCat* class there is no need to evaluate the instance's status in the *VeryLightCat* class.

A detailed performance evaluation confirming the benefit of these and other optimization strategies for view materialization can be found elsewhere [25].

6. *MultiView* Graphic User Interface

Although the GemStone OODB provides its own native graphic programming tool, named GeODE, GeODE is not an appropriate tool for *MultiView* users for a variety of reasons. First of all, GeODE does not reflect the *MultiView* object model. For example, GemStone's object model is single-inheritance, and thus does not recognize the multiple inheritance schemata of *MultiView*. Similarly, because the GemStone model does not associate extents with classes, GeODE does not support the concept of class extents and as a consequence does not support queries upon class extents. More importantly, GeODE does not provide support for the *MultiView* virtual class operators and view schema capabilities. (In addition, GemStone, Inc. has recently been phasing out the GeODE tool suite.) We therefore built a custom graphic user interface for the *MultiView* system.

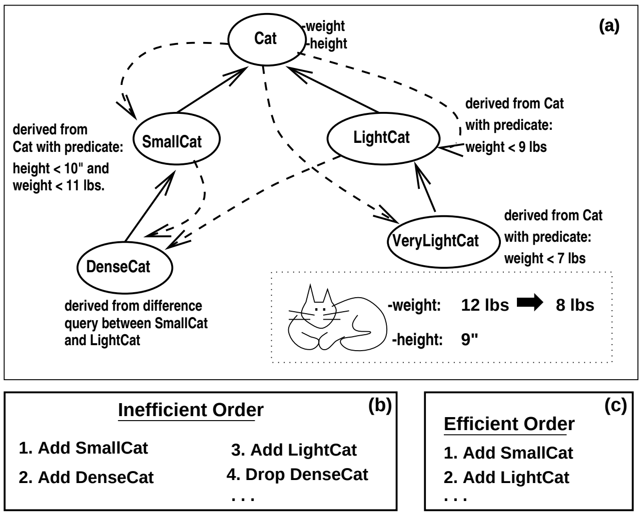

FIG. 18. An example schema where propagation order can lead to self-cancelling propagation.

For our interface, we created a customized implementation of the Tcl Windowing Shell ("wish") with extensions to interface with both *MultiView* and the underlying GemStone database system. We call our graphic interface *MultiViewer*. The *MultiViewer* extensions allow the manipulation of database objects using Tcl language constructs.

The *MultiViewer* provides convenient access to basic GemStone functions such as logging into and out of the GemStone database, sending OPAL code to the database and accessing the objects returned, and committing changes to the database structures, In addition, the *MultiViewer* provides a graphic interface to the *MultiView* database, reflecting *MultiView*-specific constructs such as:

- Global and view schema display.
- Virtual class and view schema operations.
- Access to class properties specific to the *MultiView* model, such as class extents and virtual class derivation information.

6.1 Global and View Schemata

The *MultiViewer* displays class hierarchies using a directed graph representing inheritance relationships. The graph layout is determined by an algorithm designed to minimize the number of edge crossings [44]. This results in a clean-looking directed graph that simplifies visualization of the generalization relationships by grouping similar classes together on the display.

Schema graphs can be of arbitrary size; if a schema graph is too large to fit on the physical display, it can be browsed using the scroll bars below and to the right of the graph display. Each schema graph includes a search text widget that can be used to find a specified class and to position it in the center of the display.

6.2 Virtual Class and View Schema Operations

MultiView operations are available through the schema windows' menus. Classes can be selected for participation in *MultiView* operations by clicking on them with the mouse. For example, given the base schema depicted in Figure 19, we create a new class *PropCreature* that is defined as the union of the *Horta* and *CrystalEntity* classes. *MultiView* automatically integrates the new virtual class into the global schema (Figure 20).

Users can create or modify a view schema at any time, populating it with selected base and virtual classes, to create a customized interface into the database (Figure 21). To users, these view schemata look and feel just like the base schema. View schemata thus serve as customized interfaces to the database.

6.3 Queries and Data Entry

The *MultiViewer* allows a user to interact with object instances without using native database syntax. This is especially useful when the person responsible for entering or updating records is not a computer specialist. The *MultiViewer* generates a user interface, called a form, for any object based on constraints established by the class's creator. The exact way that a class or object presents itself is completely configurable.

OBJECT FORMS FOR DATA DISPLAY, BROWSING, AND EDITING.

The basic interface to an object is a collection of graphical entries, one for each attribute defined by the object's class. The values displayed in the fields will update dynamically as changes are made in the database. The user can change any field to modify the associated object. Figure 22 shows a basic object form for a *Person* class object. *MultiViewer* builds the interface to an object dynamically by querying the object's class about which of its properties are externally modifiable. The class responds with a set of ordered tuples specifying, for each value, the name to be presented to the user, the domain class, and the accessing methods that can retrieve and set the value.

MultiView provides appropriate interface forms for many internal built-in GemStone classes, such as Integer, Strings, Set, Boolean, etc. For example, Booleans are represented in the *MultiViewer* by toggle switches representing the value's current state. The system is extensible in that additional primitive representation forms can easily be defined as needed.

For more complex classes, the *MultiViewer* will query the class about its interface form by interacting with the *MultiView* system and dynamically constructing an interface within the original form. For example, suppose a class called *Person* has a field *BestFriend*, which is constrained to be another *Person*. When editing a *Person*, the *MultiViewer* will build a form consisting of various *tools*, say, *Name*, *Age* and *Gender*. It will also display the person's *BestFriend* as a subform, containing its own fields *Name*, *Age*, *Gender* and *BestFriend*. To prevent infinite recursion, each subform is not actually built and displayed until activated by the user. Figure 23 shows a form that displays the instance indirectly referenced by an object, namely the *Person* "Jimmy" has a best friend named "Homer."

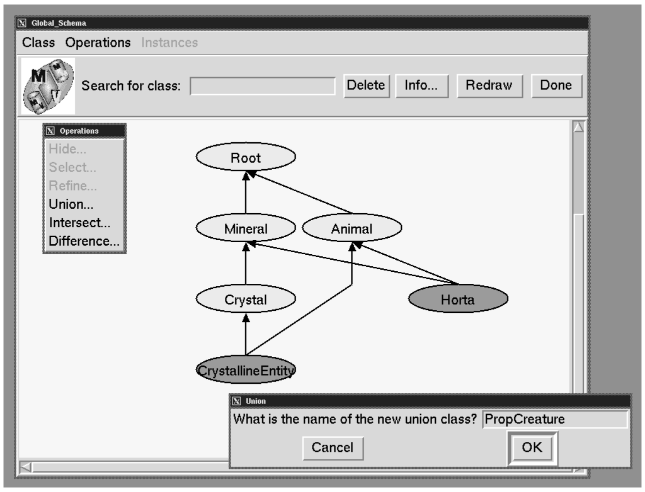

FIG. 19. Applying a union query to create a virtual class named *PropCreature* as the union of *Horta* and *CrystallineEntity*.

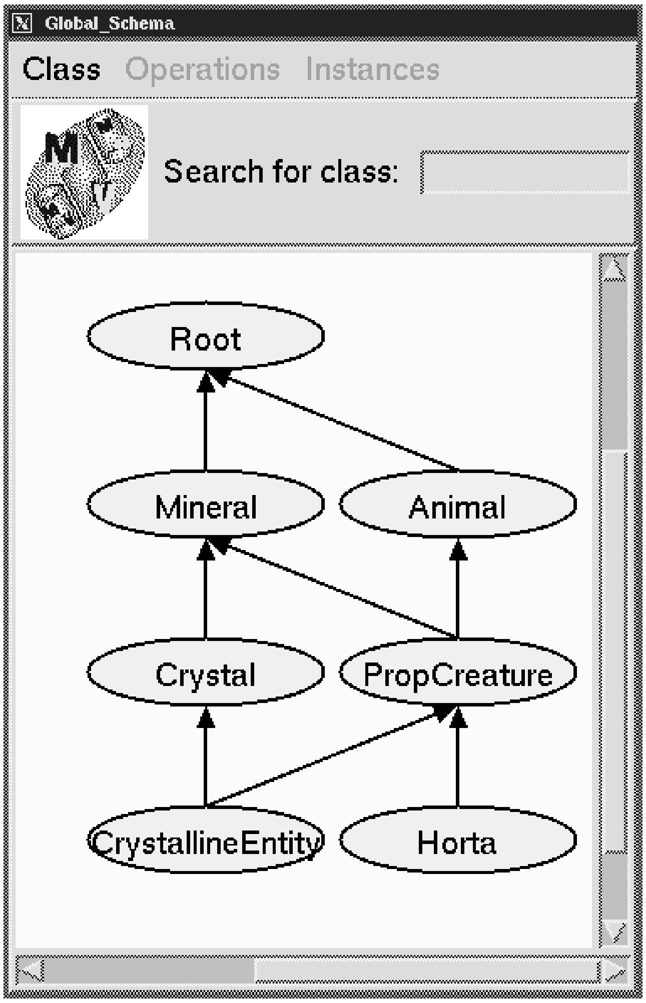

FIG. 20. The augmented global schema contains both base and virtual classes.

SELECTION FORMS FOR QUERYING.

In addition to the ability to browse and edit object instances, forms can be used to formulate queries. The *MultiViewer* particularly facilitates the specification of selection predicates. These can be used to filter a large database for objects with interesting properties, or in the creation of virtual classes. A selection form is created in much the same fashion as an object form, but the fields themselves are different. For each field, several tests customized to the data type of the attribute can be performed. These query forms support the specification of nested queries involving implicit joins. Namely, the user can base the selection predicate on attributes of other objects. By displaying the *BestFriend* subform and interacting with its fields, it is simple to select, for example, the set of all *Persons* whose *BestFriend* is a teenager.

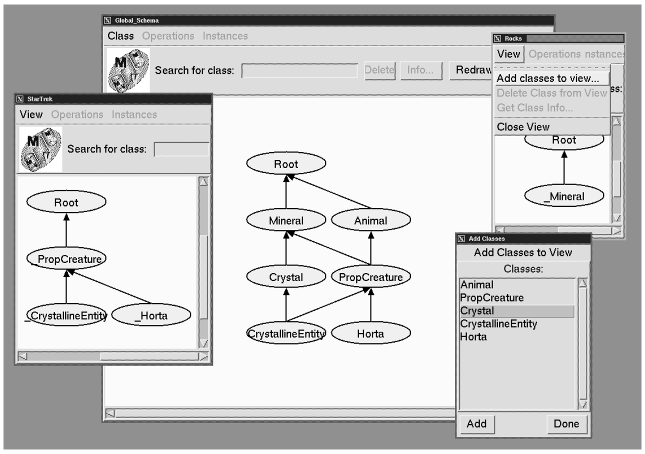

FIG. 21. View schemata look and feel like a base schema.

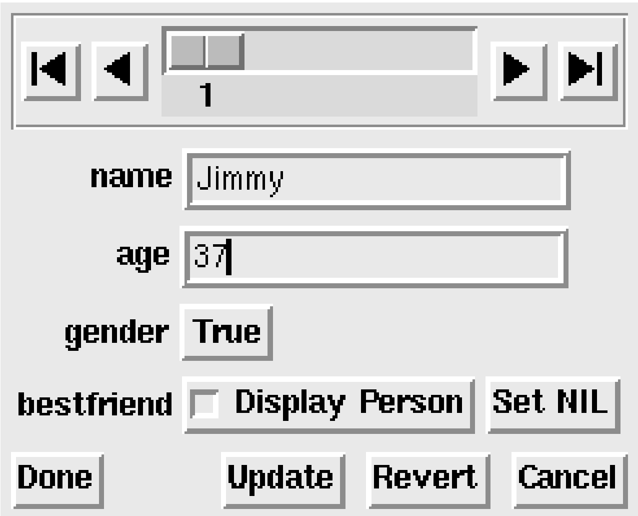

FIG. 22. Object Form for Person class objects.

CONFIGURABILITY.

If *MultiViewer*'s form generation mechanism is not sufficiently flexible for an application's purposes, the creator of an object can design a custom interface to its externally visible attributes (e.g., for multi-media display). Because *MultiViewer* is written in Tcl, an interpreted language, it can easily be extended with dynamic user-interface code. This would allow a *MultiViewer* user to view pictures or use tools custom-built for another user's class without having to build or install any software packages.

7. Related Work

RELATIONAL VERSUS OBJECT-ORIENTED VIEWS.

View mechanisms have been extensively studied for the relational model [7, 12, 2, 16]. Our research on view materialization in OODBs borrows several techniques from the relational arena. Blakeley et al. test modified tuples to see if they fulfill view predicates, thereby detecting irrelevant and autonomously computable updates [7, 8]. This resembles our solution of filtering irrelevant updates by exploiting the generalization hierarchy and the derivation structure. The system provided by [9] performs incremental view maintenance using production rules that are triggered by update operations. Similarly, we override generic-update operations with type-specific update operators for virtual classes. Although we can benefit from this work, there are some signifi-

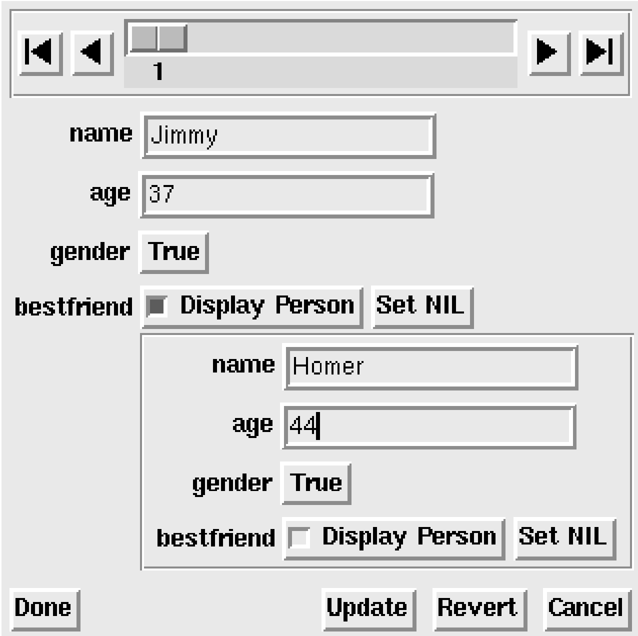

FIG. 23. Object Form with expanded subform.

cant differences between relational and object-oriented views.

- Tables in relational databases (RDBs) are arranged in a "flat" fashion in the schema. There is no information about the subset or subtype relationships contained between tables in the schema. Both virtual and base classes in the *MultiView* model are arranged in an *integrated generalization hierarchy*. *MultiView* users can thus easily perceive the relationships between OODB classes. These explicit class relationships also facilitate *MultiView*'s special optimization techniques for update propagation.
- Views in RDBs refer to virtual tables that are formed by applying query operations to tables. Because an OODB schema is organized into a generalization hierarchy, OODB views include both virtual classes (formed by applying query operations to individual classes) and virtual schema (formed by applying query operations to a base or view schema).
- Attributes in RDBs are identified as the names of columns in tables. If two relational tables each have an attribute with the same name, there is no guarantee that the two attributes refer to the same property. Each property in the *MultiView* model is defined at a single class and using property inheritance shared by other classes. Database users can thus unambiguously identify whether or not two properties represent the same instance variable or method. In addition, when a predicate term in a virtual class's query invokes a given property, we can identify exactly which class defines that property, enabling us to *localize the effects of updates*.
- The values of relational attributes are updated using generic queries, and thus update procedures must be implemented for the general case. On the other hand, because OODBs support *encapsulation*, updates to attributes are localized to specific accessing methods that are themselves defined at a single identifiable class. We can embed our update propagation methods into the accessing methods for the properties referenced by the predicates of materialized virtual classes. The accessing methods then serve as *update triggers*.
- Attributes in RDBs have simple values. Although the simple value can be a foreign key referring to a row in another table, RDB attributes cannot contain sets or other complex data structures. Instance variables in OODBs, on the other hand, can contain references to actual objects or sets of objects. These *aggregation relationships* of OODBs can be used in view formation queries.
- Data in RDBs exists as values in tables. Thus when virtual tables in RDBs are materialized, the materialized tables typically contain *copies of data values* from the base tables [9]. On the other hand, data in OODBs exists as objects that can be uniquely identified and then accessed using unique object identifiers (oids). Thus when virtual classes in OODBs are materialized, the materialized classes can contain references to the relevant objects rather than value duplicates (copies) of the original objects. This not only saves space, but also significantly simplifies the view update problem. Furthermore, regardless of which properties are "hidden" by a "project" query, objects belonging to a virtual class can be easily associated with their base incarnations. In fact, this provides us with a basis for the arbitrary restructuring of an object's look and feel without compromising its identity.

OBJECT-ORIENTED VIEW APPROACHES.

In recent years, several proposals for object-oriented view systems have appeared in the literature [1, 6, 15, 29, 40, 43]. Readers may benefit from Kim and Kelley's overview chapter on OODB view support in [21]. Most previous work regarding view systems for OODBs focuses on view formation to the exclusion of view incorporation. Note that our approach of providing for the integration of virtual classes into a single unified global schema is distinct from others found in the literature. Existing approaches either: (1) require the user to specify explicitly the relationship between a virtual class and existing base classes [19, 45]; or (2) relate a virtual class only with its direct source class via a subclass/superclass relationship [40]; or (3) simply relate a virtual class with its source class via a *derived-from* relationship [6], or (4) with the root of the schema [15, 20].

The first approach is vulnerable to potential consistency problems, since the users might introduce an inconsistency in the schema graph by inserting is-a arcs between two classes not related by a subclass relationship. A solution of verifying the correctness of the relationship in essence would have to provide a capability similar to the automatic classification approach advocated in our system, namely, a means of automatically computing the *subsumes* relationships between pairs of classes. The second approach is prone to misrepresenting the subclass relationships normally represented in a class hierarchy, in particular, because a derived class may not be *is-a* related to its immediate source class. It would at best result in a partial, hence less informative, classification of class extents. The third approach ignores the issue of determining subclass relationships by introducing a parallel *derived-from* relationship hierarchy, which is not very informative in terms of relating different classes and their type descriptions. Note that in all other approaches given above, one would of course also maintain this *derived-from* relationship by keeping the class derivation query (which will be used to recompute the population of the virtual class, whenever needed). Finally, the last approach completely ignores the issue of classification, thus resulting in a flat class structure.

Most of the current proposals for view management systems have not yet been implemented. O2

Views [29] [39], based on the work of Abiteboul and Bonner [1], is the first and only commercial implementation of an object-oriented view management system, realized. The O2 Views approach does include the integration of view classes into a view schema, but rather than supporting a global class hierarchy and migrating properties, it instead daisy-chains views and "bases."

Scholl et al. [40, 41] have developed an object-preserving algebra to define virtual classes and thus achieve updatable views. Their system, named Cocoon, has been implemented on top of a nested relational model. The *MultiView* object-algebra is similar in flavor to Cocoon's. However, Cocoon does not support the automatic generation of complete view schemata, the implementation of capacity-augmenting views, nor the incremental support of *materialized* views.

Others define view schemata through the manipulation of the object schema graph rather than solely by query languages. Tanaka et al. [45] propose that view schemata be defined by manually manipulating the edges in the global schema graph. Kim also uses DAG rearrangement for view schema definition [19]. Such DAG manipulation approaches must deal with the issues of (1) possibly introducing inconsistencies into the view schema due to human error and of (2) unintentionally modifying the semantics of a virtual class due to side effects of graph manipulation.

Rather than handle an object as being of a single fixed type, a *MultiView* object is distributed among multiple object-slicing implementation objects. In this way we overcome the stringent constraint that each object belongs to one and only one most-specific type—which is an unreasonable assumption for view systems. The Iris functional database system resembles our system in that, being built on top of a relational engine, it distributes data over several relational tables [10]. Iris does not support view mechanisms, and does not address issues of classification, inheritance for virtual classes, etc. We can compare the implementation objects of our object-slicing strategy to the oids Kifer et al. generate for complex objects returned by queries [18]. Kifer et al. provide an *id-function* that provides a correspondence between objects in the base set and the objects returned by the queries. We, on the other hand, override comparison and object identity operators so that all the implementation objects of a database objects are explicitly associated with the oid of the conceptual object.

Only a few published papers address issues of view materialization in OODBs. Hardwick and Downie provide a view materialization model in which updates are propagated by use of change files, representing histories of design sessions [13]. However, they duplicate objects (including identifiers) for virtual classes rather than merely storing references to objects [13]. Konomi et al. address maintaining consistency for a particular type of join class formed along an existing path in the aggregation graph [22]. Our work instead focuses on the exploitation of the structure of the schema hierarchy and derivation dependency graph in order to reduce update propagation.

ROLE MODELING SYSTEMS.

MultiView uses an object-slicing mechanism to address the object model requirements underlying the support of object-oriented views. These requirements include multiple classification and dynamic object migration. The flexibility offered by the object-slicing approach naturally lends itself to implementing role systems: object-slicing's implementation objects can easily be adapted to represent the various roles of objects in a role system [11]. Although our object-slicing techniques can be compared to mechanisms used in *role modeling* approaches [11, 30], no other research in the literature discusses the application of the object-slicing paradigm regarding the support of object-oriented views. In *role modeling* systems, objects dynamically gain and lose multiple interfaces (a.k.a. *roles*) throughout their lifetimes. These roles can be compared to the implementation objects of an object-slicing implementation, in that both permit objects to belong to multiple classes and change types dynamically. In some sense, accessing an object through one of its implementation objects is like accessing an object while it is playing one of its roles. However, role systems and views systems have different goals. Role systems strive to increase the flexibility of objects by enabling them to dynamically change types and class membership at the level of individual object-instances. Such changes are done explicitly by user request, and on an object-by-object basis. View systems, on the other hand, enable users to restructure the types and class membership of complete classes—based on content-based queries.

In [42], Sciore proposed an object specialization approach, in which a real world entity is modeled by multiple objects arranged in an object hierarchy. These object hierarchy objects inherit from each other, enabling each individual entity object to decide its own inheritance hierarchy. Although our implementation objects resemble object hierarchy objects in that they inherit from each other, objects in our implementation always conform to the existing global class hierarchy. This single-point-of-inheritance allows us to optimize view update propagation.

Unlike many role systems, which allow object hierarchies to exist independently from class hierarchies [42], objects in our model always conform to the existing global class hierarchy. In short, if an object possesses an implementation object of a given class's type, it must also possess an implementation object for every class that is a superclass of that given type. This achieves an efficient and uniform inheritance scheme. Also, unlike many role systems, in our implementation conceptual objects can be associated with at most one implementation object of a given type [11]. Role systems generally do not deal with the issues of virtual class derivation, classification, nor with method promotion.

The role system proposed by Gottlob et al. was implemented using techniques similar to object-slicing [11]. This system, like ours, is implemented in Smalltalk by overriding the `doesNotUnderstand:` method. Since they target role support, their system does not provide any support for views, and thus does not address issues related to view management.

The *Deputy Mechanisms* proposed by Peng and Kambayashi unify the concepts of object views, roles, and migration in the form of deputy objects and deputy classes [31]. In the deputy mechanism paradigm, view objects are treated as roles of database objects. The deputy mechanism object model is probably the work most closely related to *MultiView*. In particular, similarities of the two systems include: Both support capacity-augmenting views, dynamic classification, and multiple classification. Neither duplicate the state of an object when representing it as a virtual object. And finally, both support update propagation to materialized virtual/deputy classes.

There are some significant differences between the two systems, however. *MultiView* treats virtual classes as first-class database citizens, and thus integrates virtual and base classes into a unified global schema, whereas Deputy classes are not integrated into the global schema. Virtual classes in *MultiView* participate in the actual inheritance hierarchy and thus behave just like base classes. *MultiView* supports the automatic generation of view schemata composed of selected base and virtual classes, while the Deputy system does not support view schemata. Although both the Deputy and *MultiView* systems are based on Smalltalk, *MultiView* supports multiple inheritance. Deputy objects and source objects have independent object-identifiers, whereas a view object and a source object in *MultiView* share a conceptual object-identifier in addition to individual implementation object-identifiers. Finally, *MultiView* supports optimized incremental materialized view maintenance algorithms that exploit the integrated class hierarchy structure.

SUBJECTIVITY.

Although *MultiView* is a view system, its capabilities can be compared to those required by recent subjectivity technology. A *subject* is a collection of related classes that form a subjective viewpoint [14]. *MultiView*'s view schemata could be used to implement such collections. In addition, we can compare *MultiView*'s support for customized interfaces for objects via virtual class types to subjective systems' multiple interfaces to objects. *MultiView* also facilitates the sharing of data objects amongst diverse applications. However, because the *MultiView* effort has focussed on providing comprehensive view features, it currently does not provide other features commonly associated with subjectivity, such as explicit active database capabilities, roles that are distinct from the class hierarchy, or support for customized concurrency control.

8. Conclusions

In this paper, we have provided an overview of the *MultiView* system—one of the first implemented OODB systems to support dynamic and updatable materialized object-oriented database views. In particular, we have presented the motivation and execution of our design decisions, system architecture, and the *MultiView* system's capabilities. The implementation solutions we describe are general and thus should be applicable to other object-oriented systems. To summarize, the *MultiView* system is distinguished by a number of unique features, which include:

- *MultiView* integrates virtual and base classes into a unified global schema, resulting in effective code reuse as well as sharing of object storage between base and virtual objects.
- Virtual classes in *MultiView* participate in the actual inheritance hierarchy and thus behave just like first-class database citizens (i.e., base classes).
- *MultiView* supports capacity-augmenting views (virtual classes can define additional attributes and methods that are not derivable from the existing schema).
- *MultiView* automatically generates view schemata composed of user-selected base and virtual classes.
- *MultiView* includes optimized incremental view maintenance algorithms for materialized views that exploit the integrated class hierarchy structure.

During our initial design stage, we identified a set of key object-model properties that are critical for the realization of an object-oriented view system. To the best of our knowledge no existing commercial OODB supports all of these properties. Our implementation of the *MultiView* system uses the object-slicing representational technique [23] to provide these necessary object model properties, which include multiple classification, object migration, and dynamic re-classification.

The current version of *MultiView* uses the 4.0.1 version of GemStone and was developed using the OPAL interface. The database is built on a Sun Sparc 10 running SunOS 4.1.3 with 48 megabytes of memory. *MultiView*'s implementation consists of approximately 50 classes and 250 class and instance methods, all of which are defined using around 11,000 lines of code. The code used to implement object-slicing makes up about 5% of the total code. The graphic interface was implemented separately, using Tcl/Tk.

We have performed several studies to evaluate the *MultiView* system and its embedded algorithms. Our experiments demonstrate the benefit of techniques used by our update propagation algorithms [25, 27]. Namely, we show that the potential problems of self-cancelling propagation and of early branch condition termination are indeed handled by our solution and result in significant performance gains. In a related paper, we explore the use of clustering to ameliorate the overhead of the object-slicing technique employed by our implementation [23].

Because it is one of the only implemented object-oriented view systems to date, *MultiView* is currently being used in a number of projects at the University of Michigan. For example, the *Transparent Schema Evolution* project, funded by an NSF Young Investigator Award, uses *MultiView* as an implementation base to develop a suite of schema evolution tools for interoperability of diverse applications via a shared, yet evolving object repository by preserving existing database interfaces through schema change [33, 34].

In the future, we plan to extend our current work to examine the area of distributed views. We also want to study issues related to the support of deferred updates and multiple (batched) updates for materialized views.

Acknowledgments

A thank you goes to Young-Gook Ra, who helped design and implement the *MultiView* object-slicing architecture. Thank you and appreciation also to Pedro Jose Marron, Daryl Scott Cantrell, Matthew Jones and Victor Ng, who implemented *MultiView*'s graphic interface. Besides the persons listed above, we also would like to thank the other members of our UM Database Group in particular Viviane Taube, Stacie Hibino, Woody Huang, and Nauman Chaudhry for valuable discussions and suggestions throughout the development of the *MultiView* system.

Notes

1. GemStone is a registered trademark of GemStone Systems, Inc.
2. We associate both **type** and **extent** with our concept of **class**. Although there is no general agreement on whether or not classes in OODBs should incorporate their own extents rather than requiring users to maintain their own collections of class-instances, several systems follow this philosophy, including Orion and the system proposed by H. J. Kim [19]. Furthermore, the proposed ODMG standard [3] recently formulated by several key OODB vendors also follows this approach.
3. Most OODBs typically represent an object as a chunk of contiguous storage determined at object creation time. They thus adhere to the invariant that *an object belongs to exactly one class only* — and indirectly also to all the class's superclasses.
4. Each implementation object by default possesses its own object identifier. However, because the implementation object serves as an interface for a specific conceptual object, the object-identifier of the conceptual object supersedes that of the implementation object for most practical purposes, such as determining object-equality.
5. If a method with selector m_i is found in more than one class in **SuperclassHierarchy**(C_i), then the user is prompted to **cast** O_i into a non-ambiguous implementation object.
6. The methods used to set or retrieve an instance variable's value are called that instance variable's *accessing methods*. Accessing methods are always located at the same class as the instance variables for which they are defined, and thus when an instance variable migrates from one class to another, that instance variable's accessing methods must make the same migration.
7. For example, we could swap the identities of two implementation objects to give the implementation object of the modified existing class the same object identifier as the original implementation object.
8. For the sake of simplicity, we currently make the assumption that instance variables reside at the same location as their update methods.
9. Note that some materialized RDB views may be implemented using techniques such as view indices that resemble membership materialization.

References

1. S. Abiteboul and A. Bonner. Objects and views. *SIGMOD*, pages 238–247, 1991.

2. R. Agrawal and H. V. Jagadish. Materialization and incremental update of path information. In *IEEE International Conference on Data Engineering*, pages 374–383, 1989.

3. T. Atwood, R. Cattell, J. Duhl, G. Ferran, and D. Wade. The ODMG object model. *Journal of Object Oriented Programming*, pages 64–69, June 1993.

4. T. Barsalou, N. Siambela, A. M. Keller, and G. Wiederhold. Updating relational databases through object-based views. *SIGMOD*, pages 248–257, 1991.

5. E. Bertino. Integration of heterogeneous data repositories by using object-oriented views. *International Workshop on Interoperability in Multidatabase Systems*, pages 22–29, April 1991.

6. E. Bertino. A view mechanism for object-oriented databases. In *3rd International Conference on Extending Database Technology*, pages 136–151, March 1992.

7. J. Blakeley, P. Larson, and F. Tompa. Efficiently updating materialized views. *SIGMOD*, pages 61–71, 1986.

8. J. A. Blakeley, N. Coburn, and P-A Larson. Updating derived relations: Detecting irrelevant and autonomously computable updates. *ACM Transactions on Database Systems*, 14(3):369–400, September 1989.

9. S. Ceri and J. Widom. Deriving production rules for incremental view maintenance. In *International Conference on Very Large Data Bases*, pages 577–589, September 1991.

10. D.H. Fishman. Iris: An object oriented database management system. In *ACM Transactions on Office Information Systems*, volume 5, pages 48–69, January 1987.

11. G. Gottlob, M. Schrefl, and B. Rock. Extending object-oriented systems with roles. manuscript, 1993.

12. E. N. Hanson. A performance analysis of view materialization strategies. *SIGMOD*, pages 440–453, 1987.

13. M. Hardwick and B. R. Downie. On object-oriented databases, materialized views, and concurrent engineering. In *Proceedings of the 1991 ASME International Computers for Engineering Conference and Exposition*. Engineering Databases: An Engineering Resource, 1991.

14. W. Harrison, H. Ossher, and H. Mili. Subjectivity in object-oriented systems workshop summary. In *OOPSLA Workshop*, 1995.

15. S. Heiler and S. B. Zdonik. Object views: Extending the vision. In *IEEE International Conference on Data Engineering*, pages 86–93, 1990.

16. H. Jakobsson. On materializing views and on-line queries. In *ICDT*, pages 407–420, October 1992.

17. M. Kaul, K. Drosten, and E. J. Neuhold. Viewsystem: Integrating heterogeneous information bases by object-oriented views. In *IEEE International Conference on Data Engineering*, pages 2–10, February 1990.

18. M. Kifer, W. Kim, and Y. Sagiv. Querying object-oriented databases. *SIGMOD*, pages 393–402, 1992.

19. H. J. Kim. *Issues in Object-Oriented Database Systems.* PhD thesis, University of Texas at Austin, May 1988.

20. W. Kim. A model of queries in object-oriented databases. In *Proceedings of the International Conference on Very Large Databases*, pages 423–432, August 1989.

21. W. Kim. *Modern Database Systems: the Object Model, Interoperability, and Beyond.* ACM Press, 1995.

22. S. Konomi, T. Furukawa, and Y. Kambayashi. Super-key classes for updating materialized derived classes in object bases. In *International Conference on Deductive and Object-Oriented Databases*, pages 310–326, July 1993.

23. H. A. Kuno, Y. G. Ra, and E. A. Rundensteiner. The object-slicing technique: A flexible object representation and its evaluation. Technical Report CSE-TR-241-95, University of Michigan, 1995.

24. H. A. Kuno and E. A. Rundensteiner. Implementation experience with building an object-oriented view management system. Technical Report CSE-TR-191-93, University of Michigan, Ann Arbor, 1993.

25. H. A. Kuno and E. A. Rundensteiner. Incremental update propagation algorithms for materialized object-oriented views in *MultiView*. Technical Report CSE-TR-246-95, University of Michigan, Ann Arbor, June 1995.

26. H. A. Kuno and E. A. Rundensteiner. Materialized object-oriented views in *MultiView*. In *ACM Research Issues in Data Engineering Workshop*, pages 78–85, March 1995.

27. H. A. Kuno and E. A. Rundensteiner. Using object-oriented principles to optimize update propagation to materialized views. In *To be published in ICDE'96*, 1996.

28. J. Martin and J. Odell. *Object-Oriented Analysis and Design.* Prentice-Hall, Inc., 1992.

29. O2 Technology. *O2 Views User Manual*, version 1 edition, December 1993.

30. M. P. Papazoglou. Roles: A methodology for representing multifaceted objects. In *International Conference on Database and Expert Systems Applications*, pages 7–12. Springer-Verlag, 1991.

31. Z. Peng and Y. Kambayashi. Deputy mechanisms for object-oriented databases. In *IEEE International Conference on Data Engineering*, pages 333–340, March 1995.

32. Y. G. Ra, H. A. Kuno, and E. A. Rundensteiner. A flexible object-oriented database model and implementation for capacity-augmenting views. Technical Report CSE-TR-215-94, University of Michigan, 1994.

33. Y. G. Ra and E. A. Rundensteiner. A transparent object-oriented schema change approach using view schema evolution. In *IEEE International Conference on Data Engineering*, pages 165–172, March 1995.

34. Y. G. Ra and E. A. Rundensteiner. A transparent schema evolution system based on object-oriented view technology. *To be accepted in IEEE Transactions on Knowledge and Data Engineering*, 1995.

35. E. A. Rundensteiner. *MultiView*: A methodology for supporting multiple views in object-oriented databases. In *18th VLDB Conference*, pages 187–198, 1992.

36. E. A. Rundensteiner. Tools for view generation in OODBs. In *International Conference on Information and Knowledge Management*, pages 635–644, November 1993.

37. E. A. Rundensteiner. A classification algorithm for supporting object-oriented views. In *International Conference on Information and Knowledge Management*, pages 18–25, November 1994.

38. E. A. Rundensteiner and L. Bic. Set operations in new generation data models. *IEEE Transactions on Knowledge and Data Engineering*, 4:382–398, August 1992.

39. C. Souza dos Santos, S. Abiteboul, and C. Delobel. Virtual schemas and bases. In *International Conference on Extending Database Technology (EDBT)*, 1994.

40. M. H. Scholl, C. Laasch, and M. Tresch. Updatable views in object-oriented databases. In *Proceedings of the Second DOOD Conference*, December 1991.

41. M. H. Scholl and H. J. Schek. Survey of the cocoon project. *Objektbanken fur Experten*, October 1992.

42. E. Sciore. Object specialization. *ACM Transactions on Information Systems*, pages 103–122, April 1989.

43. J. Shilling and P. Sweeney. Three steps to views: Extending the object-oriented paradigm. In *OOPSLA*, pages 353 – 361, October 1989.

44. K. Sugiyama, S. Tagawa, and M. Toda. Methods for visual understanding of hierarchical system structures. *IEEE Transactions on Systems, Man, and Cybernetics*, pages 109–125, February 1981.

45. K. Tanaka, M. Yoshikawa, and K. Ishihara. Schema virtualization in object-oriented databases. *IEEE International Conference on Data Engineering*, February 1988.