




Database Versioning

David Landén
Per-Magnus Olsson


08-12-09 1



Overview

- Motivation
- Common operations
- Example: Gemstone
- Example: O₂
- Exercises


08-12-09 2



Motivation

- Two major reasons for versioning:
 - Data change and desire to access historical values.
 - Schema change.


08-12-09 3



General Considerations

- When to convert?
 - Immediate – change everything now.
 - Deferred – change as object is used
- References to deleted objects?


08-12-09 4



Common Operations

- Add/remove/modify/rename attribute
- Add/remove/modify/rename class
- Make class S a superclass of C
- Remove inheritance between S and C

08-12-09 5



Example: GemStone

- Object-oriented database management system in mid 80's,
- GemStone's successors are used today in finance and many other areas.
- Influenced by Smalltalk
- Influenced J2EE

08-12-09 6

Philosophy

- Allow modifications that are
 - Useful
 - Well understood
 - Have reasonable implementation
- → Not all operations are supported

08-12-09

7

Authorization and Concurrency

- All objects belong to a segment
- Authorization per segment, not per object
- Access to object does not imply access to instance variables
- Access conflicts checked at commit time
- Read-only transactions always succeeds
- Changes visible to other sessions after commit
- Changes made by others visible after current transaction ends

08-12-09

8

Variables and Constraints

- Indexable – directly accessible through index (c.f. id number)
- Named – accessible through name.
- Anonymous – not accessible through name or index.
- Named or anonymous are constrained, the value of the variable is constrained to be a kind of a given class.

08-12-09

9

Class Level Operations

- Removing a class unless
 - instances exists
- Adding a class
 - Leaf nodes without restrictions
 - Interior nodes requires specification of superclass and subclass
 - New instance variables must be added separately

08-12-09

10

Class Level Operations, cont.

- Making a class indexable allowed unless
 - subclass already indexable and indexing the superclass makes the variable constraints inconsistent.
- Making a class not indexable
 - allowed for all instances
 - not propagated to subclasses
 - Exception: subclass must be indexable of superclass is indexable

08-12-09

11

Variable Level Operations

- Adding named instance variable allowed unless
 - name clash with existing instance variable in class
 - variable exists in subclass
 - variable propagated to subclasses
- Removing named instance variable allowed unless
 - variable exists in superclass
 - not propagated to subclasses

08-12-09

12

Variable Level Operations, cont

- Renaming of variables allowed unless
 - variable inherited
 - name clash occurs with new name

08-12-09

13

Variable Constraints

- Specialization
 - New constraint is subclass of old constraint
- Generalization:
 - New constraint is superclass of old constraint
- Specialization allowed unless
- Generalization allowed unless

08-12-09

14

O₂

- O₂ has Objects and types:
 - A type is defined recursively from atomic types.
 - Class methods coded in O₂C or C++.
 - Attaching an object or type to name in a schema make it persistent.
 - Encapsulation at: class, application and schema level.

08-12-09

15

O₂

- Schema modifications:
 - Adding/deleting attributes to a class.
 - Redefining the structure of a class.
 - 10 different schema manipulation primitives:
 - creation of a new class
 - creation of new attribute,
 - deletion of attributes
 - ...

08-12-09

16

O₂

Database updates:

- Conversion functions:
 - Default
 - User defined.
- Supports both deferred and immediate conversion.

08-12-09

17

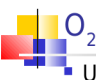
O₂

Default conversion function example:

```
create schema Car showroom;
class Vendor type tuple ( name: string,
                        address: tuple ( city: string,
                                        street: string,
                                        number: real),
                        sold_cars: list( Car ) )
end;
class Car type tuple ( name: string,
                    price: real,
                    horse_power: integer )
end;
```

08-12-09

18




User-defined conversion function example:

```

begin modification in class Car;
delete attribute horse_power;
create attribute kW : integer;
conversion functions:
conversion function mod_kW (old : tuple(name:string, price:real,
horse_power:real)) in class Car
{
self->kW = round( old.horse_power / 1.36 );
};
end modification;

```

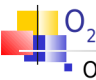
08-12-09 19



Conversion function types:

- Simple:
 - Uses only object local information.
- Complex:
 - Uses other objects in the conversion.


08-12-09 20



Object Migration functions:

- Single object change:
 - Uses a migrate() method.
 - Only downwards in subclasses.
- Class change:
 - Migration functions:

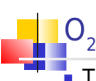
08-12-09 21



Implementation:

- Based on deferred database update.
- Complex conversion functions complicates the updates.
- Solution is to save all versions of objects, (screend).
- An object has its current value and screened value.
- Conversion and migrate functions can access the screened value.


08-12-09 22



The O₂ schema manager keep track on the evolution of schemas.

- updates the schema state for each modification.
- Each class in a schema is represented with a class descriptor, a meta class in O₂.

08-12-09 23



O₂ – Meta Objects

```

class Meta_class type tuple (
...
type : Meta_type,
properties : list (Meta_property),
...
cur_tid : integer,
history : list (Meta_history_entry),
...
end;

```

08-12-09 24

O₂ – Meta History

```
class Meta_history_entry type tuple (
  tid : integer,
  type : Meta_type,
  ex_type : Meta_type,
  struct : list(Meta_property_entry),
  cf : Meta_conversion,
  mf : Meta_migration)
end;
```

08-12-09

25

O₂ – Meta Property

```
class Meta_property_entry type tuple (
  pid : integer,
  sch_state : integer,
  offset : integer,
  type : Meta_type,
  status : {existing, screened})
end;
```

08-12-09

26

O₂ – Meta Conversion

```
class Meta_conversion type tuple (
  next_state : integer,
  function : Meta_binary)
end;
```

08-12-09

27


O₂ – Meta Migration

```
class Meta_migration type tuple (
  sch_state : integer,
  function : Meta_binary)
end;
```

08-12-09

28

O₂ – Example



08-12-09

29

Deferred update Algorithm in O₂

When accessing an object o:

- If the object conforms to its last class definition in the schema.
 - Then the object can be used.
 - If not, the corresponding history entry descriptor is found for the object.

08-12-09

30



The history entry descriptor is examined:

- If there is a migration function.
 - If o is not already migrated use the migration function.
- If there is a conversion function
 - If there is a conversion function use it.
- If there is no migration or conversion function.
 - Use default conversion.

08-12-09

31




Problem with complex conversion:

- When updating an object to time t_n :
 - And accessing other objects in a conversion function could lead to a problems.
 - If the conversion functions was specified in t_m , it could only access other objects of t_m .
 - So other objects must be updated to t_m , when access by this conversion function at t_n . ($t_m < t_n$)

08-12-09

32



$f_1: j = B.j * 5;$
 $f_2: k = j + B.m;$

Questions:
1) An object (instance of A) conforming to t_0 , describe how it is updated at t_1 , using the deferred algorithm.
2) An object (instance of A) conforming to t_0 , describe how it is updated at t_1 , using the deferred algorithm.

08-12-09

33