

# Lazy Type Changes in Object-oriented Database

Shan Ming Woo and Barbra Jaiskov

MIT Lab. for Computer Science

December 1999

# Background

- w Behavior of OODB apps compose of behavior of persistent obj
- w Behavior of objects governed their types
- w *Type change* needed to update OODB apps
- w *How to execute type change?*

# Requirements

- w A type change may affect types
- w An *upgrade* consists of a set of retype changes
- w Upgrades are ordered
- w Execution of an upgrade have to be atomic w.r.t. app transactions to prevent type errors

# Na ve Execution

- w Step 1: shut down the database
- w Step 2: transform objects
- w Step 3: restart the database
- w Drawback *database availability is su*
- w Not suitable for large databases and mission-critical databases
- w Solution *lazy type changes*

# Lazy Type Changes

- w Objects transformed, lazy, just before use
- w Database availability affected
- w Workload of an upgrade:
  1. Spread effectively over time
  2. Distributed among apps
- w Suitable for all data classes in large-scale and mission-critical

# Theory

w Type change  $\mathcal{C}_T = \langle T, T, f^T \rangle$

Pre-type

Post-type

Transform  
function

w Upgrade:  $\mathcal{U} \langle n, \{ \mathcal{C}_i \} \rangle$

Serial #

Set of type changes

# Theory (cont.)

w Transform functions:

- n Act on one object at a time
- n Preserve object identity.  
object references survive changes
- n Should not modify any data

w Upgrades have to be complete i.e.  
should not affect other types

# Implementation Design (1)

w Based on Thor  
[ECOOP99]

n Distributed client/server  
OODBMS

n Optimistic concurrency  
control

n Servers store objects,  
validate transactions

n Clients cache objects,  
operate on cached  
objects on behalf of apps<sup>n</sup>

n Object identity partly  
location dependent

n Objects in client cache  
indexed by a resident  
object table (ROT)

n Each ROT entry stores a  
dispatch vector pointer  
and a field pointer for  
object

n Pointers in ~~empty~~  
entry are null

n Pointers in ~~full~~  
entry are up-to-date



# Implementation Design (2)

## w Single Server

- n Upgrades stored at the server and pushed to clients
- n Objects transformed by clients before used by apps
- n Objects transformations are regarded as modifications

- n Invariant: *All full ROT entries represent up-to-date objects.*
- n At receipt of new upgrade, client
  - w aborts running transaction if it used affected objects
  - guarantee atomicity
  - w scans ROT and empties affected entries
  - preservation invariant

# Implementation Design (3)

- w Single server continued
  - n Special client cooperate with garbage collector to transform rarely used objects
  - n Upgrades complete in order: when all objects affected by the oldest upgrade has been transformed, it is discarded
  - n Problem: objects change sizes across transformations
- n Size decreases---
  - w overwrite original object
- n Size increases---
  - w find space on the same server page:
    - w If succeed, update page offset table
    - w Otherwise, write to another page and replace original object with surrogate
  - n Reference through surrogate is shortcut when the referring object is modified

# Implementation Design (4)

## w Multiple servers

- n A master server stores the master copy of all upgrades

- n Upgrades pushed from master server to other servers

- n Each server sends upgrades to its clients and each client processes upgrades as in the single server implementation

- n Problem: need to maintain consistency when a client talks to multiple servers---an upgrade may arrive at one server before another

- n Client acts as relay to restore consistency: ea commit request is tagged with serial number of newest upgrade at client

# Conclusion

w Lazy type changes preserve database availability and can be efficiently implemented essential for large-scale mission-critical databases

## Future work

- w Project focused on persistent objects: how avoid recompiling applications? Universal application framework?
- w How to extend non-database environment