Architecture of the Orion Database System

Liza Blaney

October 28, 1996

Slide 1

Three implementations of ORION:

ORION-1 single processor, single user, multitask

ORION-1SX LAN-based client/server $\Longleftarrow$

ORION-2 distributed object-management

Major Subsystems

Slide 2

- Message Handler
- Object Subsystem
  - schema
  - query optimization
  - versions
  - composite objects
  - multimedia objects
- Transaction Subsystem
  - concurrency
  - recovery
- Storage Subsystem - persistence

Message Handler

Slide 3

- Orion is a message-based system
- single-user vs. multi-user
- hardware dependent
- dependent on underlying operating system process support

Storage Subsystem

Slide 4

Objects need to be managed in two different formats:

- a disk-format in the page buffer
- an in-memory format in the object buffer

Storage subsystem consists of:

- page buffer manager - fix/unfix pages
- object buffer manager - Resident Object Table (ROT)
- access manager - transfers objects between page and object formats/buffers
- storage manager - fetching/writing/logging

Storage Subsystem Issues

Slide 5

- dual-buffer management
- cache validation
- object identifier syntax
- disk space management
- clustering: physical storage segment
- object directory: hashing (vs. B-tree)
- object storage format
- secondary indexing

Secondary Indexing

Slide 6

- class hierarchy index (vs. single class)
- space/performance tradeoffs
- single attribute, class hierarchy indices

Domain	*Key*
primitive class	primitive value
user-defined class	UID of an instance of the domain class

Transaction Subsystem

Slide 7

- serializable, atomic transactions
- concurrency: locking and deadlock detections
- recovery:
  - logging
  - no archival dumping for hard-crash recovery

Object Subsystem

Slide 8

- dynamic schema evolution
  - 2 types of evolution:
    - class definition
    - class hierarchy
  - deferred update (vs. immediate update)
- query optimization and processing
  - single target class
  - class-hierarchy rooted at a target class
  - no support for extended query model

Query Optimization Issues

Slide 9

- Query Execution Plan
  - order to traverse the classes: the target class and the nested domain classes on which search prediacates are specified
  - technique to join an instance of one class and an instance of the domain of an attribute of that class
  - access path: all instances of a class or a secondary index
- Parse Tree - 3 types of nodes:
  - attribute node: target class or an attribute of a class
  - connective node: boolean
  - comparator node: predicate

Query Processing Issues

Slide 10

- single buffer vs. dual buffer
- single-class and nested class
- client and server queries