

Object-Oriented Databases

Design and Implementation: OMS Avon

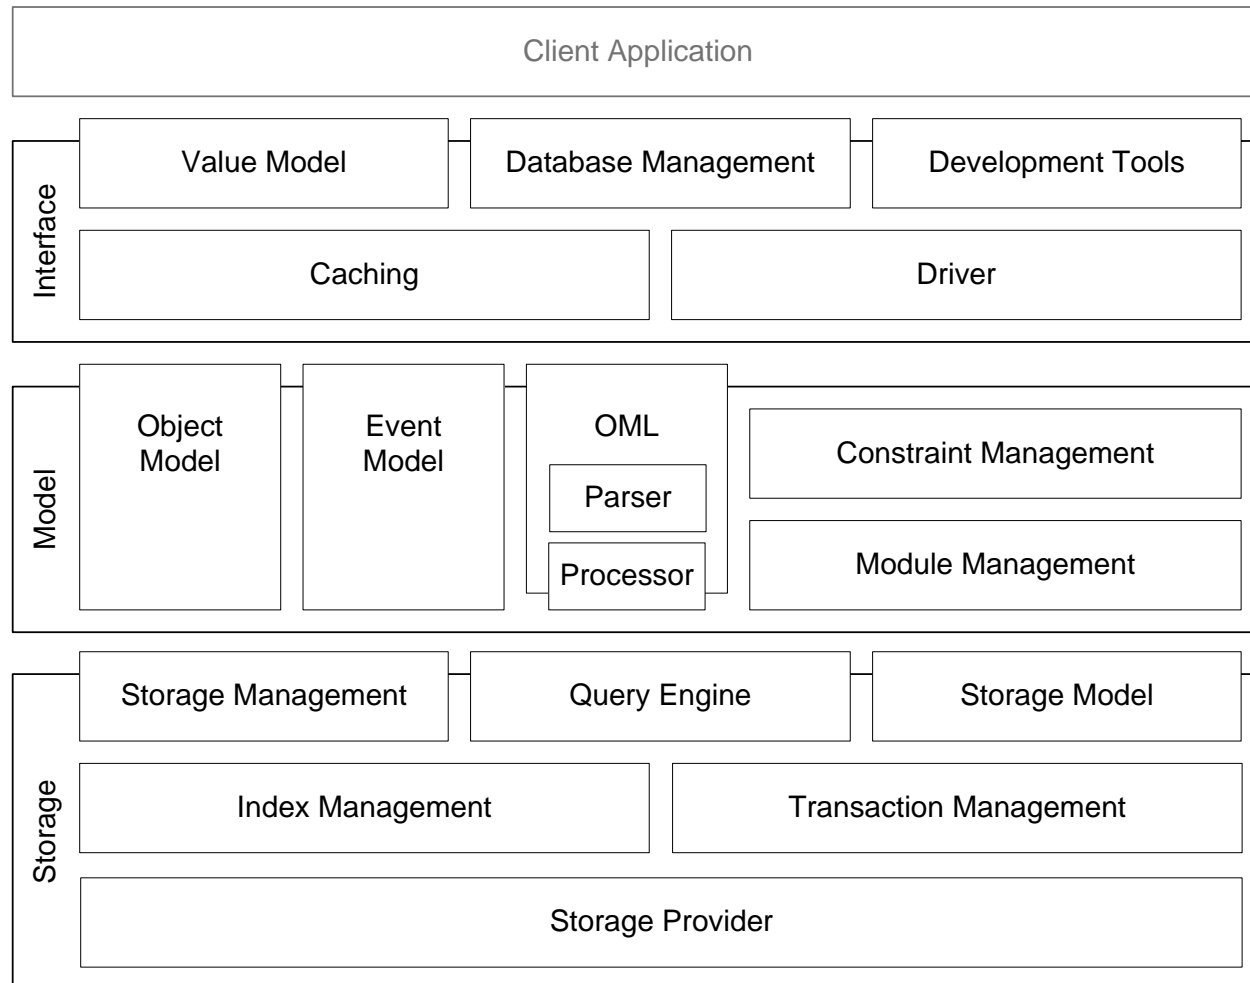
- Architecture
- Storage, Model and Interface Layer
- Database Modules



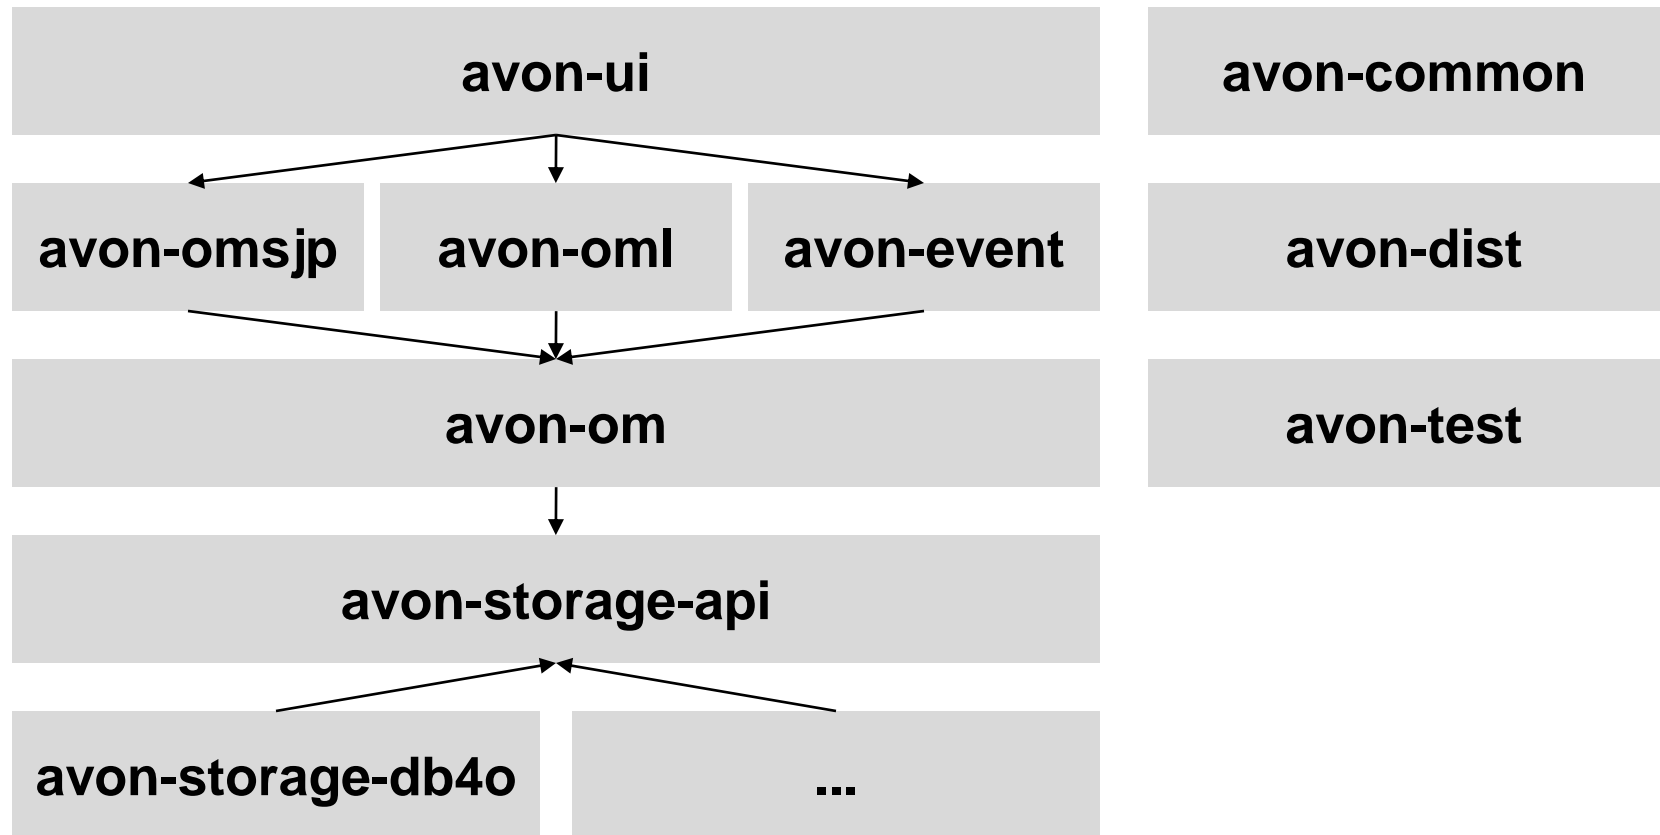
OMS Avon

- Java implementation of the OM data model and OML
- Storage layer
 - manages low-level storage
 - package `ch.ethz.globis.avon.storage`
- Model layer
 - implements functionality associated with the OM data model
 - package `ch.ethz.globis.avon.om`
 - package `ch.ethz.globis.avon.oml`
- Interface layer
 - provides a high-level application programming interface
 - package `ch.ethz.globis.avon.omsjp`

OMS Avon Architecture



OMS Avon Project Modules



OMS Avon Storage Layer

- Storage interface based on type and information units
 - type units provide metadata
 - information units store data
- Application programming interfaces for
 - create, retrieve, update and delete operations
 - schema evolution
 - creating and managing indexes
 - low-level query operators
 - transactions, concurrency control and recovery
- Extent value handles manage bulk values
- Various storage providers

Storage Layer Concepts

«interface» Storage

```
+createBaseType(TransactionHandle, Identifier, BuiltInType)
+createRecordType(TransactionHandle, Identifier, Identifier[])
+createExtentType(TransactionHandle, Identifier, BulkType, Identifier)
+createObjectType(TransactionHandle, Identifier, Identifier[])
+createExtent(TransactionHandle, Identifier): ExtentValueHandle
+getExtentValue(TransactionHandle, Identifier): ExtentValueHandle
+deleteExtent(TransactionHandle, Identifier)
+createObject(TransactionHandle): Identifier
+dressObject(TransactionHandle, identifier, Identifier)
+getAttributeValues(TransactionHandle, Identifier, Identifier): Object[]
+setAttributeValues(TransactionHandle, Identifier, Identifier, Object[])
+stripObject(TransactionHandle, identifier, Identifier)
+deleteObject(TransactionHandle, Identifier)
...
```

«interface» TransactionHandle

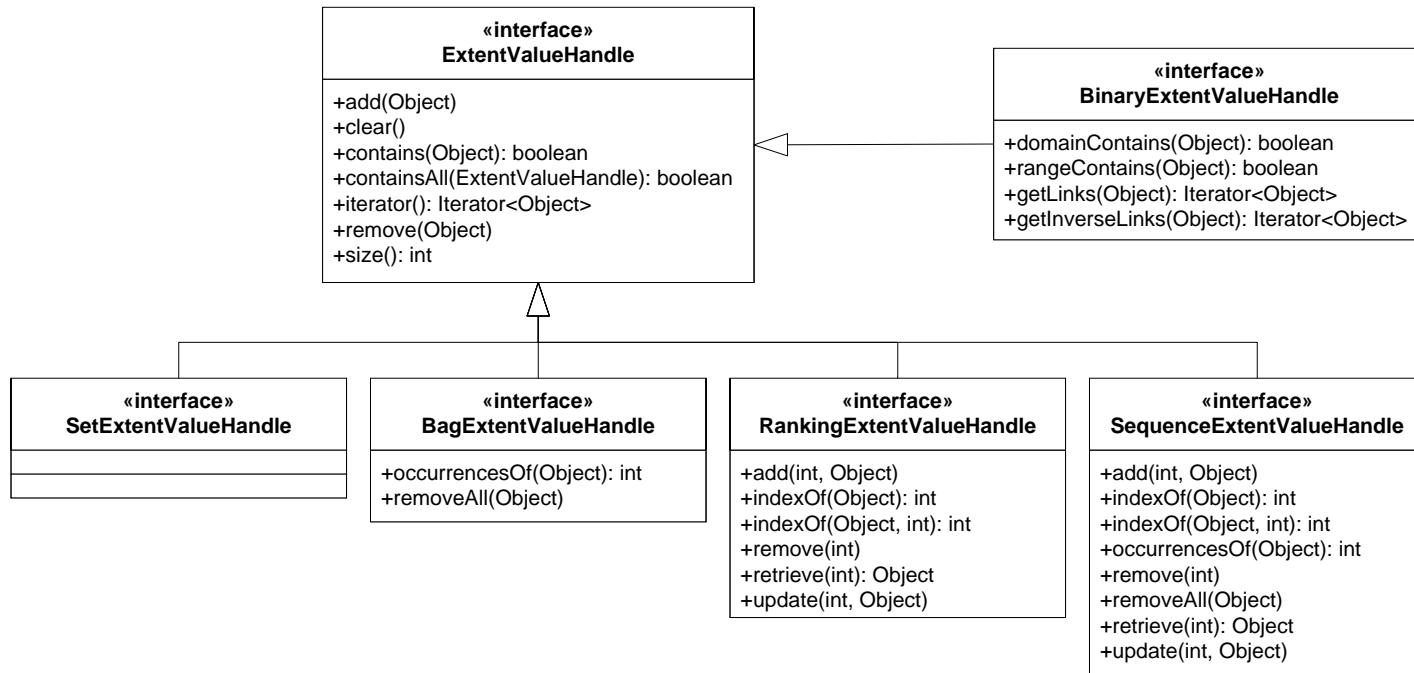
```
+abort()
+commit()
+isAborted(): boolean
+isCommitted(): boolean
+isRunning(): boolean
```

«interface» Identifier

«interface» Algebra

```
+attributeAccess(TransactionHandle, ExtentValueHandle, Identifier, int): ExtentValueHandle
+average(TransactionHandle, ExtentValueHandle): double
+closure(TransactionHandle, ExtentValueHandle): ExtentValueHandle
+composition(TransactionHandle, ExtentValueHandle, ExtentValueHandle): ExtentValueHandle
+count(TransactionHandle, ExtentValueHandle): int
+difference(TransactionHandle, ExtentValueHandle, ExtentValueHandle): ExtentValueHandle
+division(TransactionHandle, ExtentValueHandle, ExtentValueHandle): ExtentValueHandle
+domain(TransactionHandle, ExtentValueHandle): ExtentValueHandle
+domainRestriction(TransactionHandle, ExtentValueHandle, ExtentValueHandle): ExtentValueHandle
+domainSubtraction(TransactionHandle, ExtentValueHandle, ExtentValueHandle): ExtentValueHandle
+extent(TransactionHandle, Identifier): ExtentValueHandle
+flatten(TransactionHandle, ExtentValueHandle): ExtentValueHandle
+intersection(TransactionHandle, ExtentValueHandle, ExtentValueHandle): ExtentValueHandle
+inverse(TransactionHandle, ExtentValueHandle): ExtentValueHandle
+map(TransactionHandle, ExtentValueHandle, Function): ExtentValueHandle
+maximum(TransactionHandle, ExtentValueHandle): Object
+minimum(TransactionHandle, ExtentValueHandle): Object
+nest(TransactionHandle, ExtentValueHandle): ExtentValueHandle
+product(TransactionHandle, ExtentValueHandle, ExtentValueHandle): ExtentValueHandle
+range(TransactionHandle, ExtentValueHandle): ExtentValueHandle
+rangeRestriction(TransactionHandle, ExtentValueHandle, ExtentValueHandle): ExtentValueHandle
+rangeSubtraction(TransactionHandle, ExtentValueHandle, ExtentValueHandle): ExtentValueHandle
+reduce(TransactionHandle, ExtentValueHandle, Function): Object
+selection(TransactionHandle, ExtentValueHandle, Predicate): ExtentValueHandle
+sum(TransactionHandle, ExtentValueHandle): double
+union(TransactionHandle, ExtentValueHandle, ExtentValueHandle): ExtentValueHandle
+unnest(TransactionHandle, ExtentValueHandle): ExtentValueHandle
```

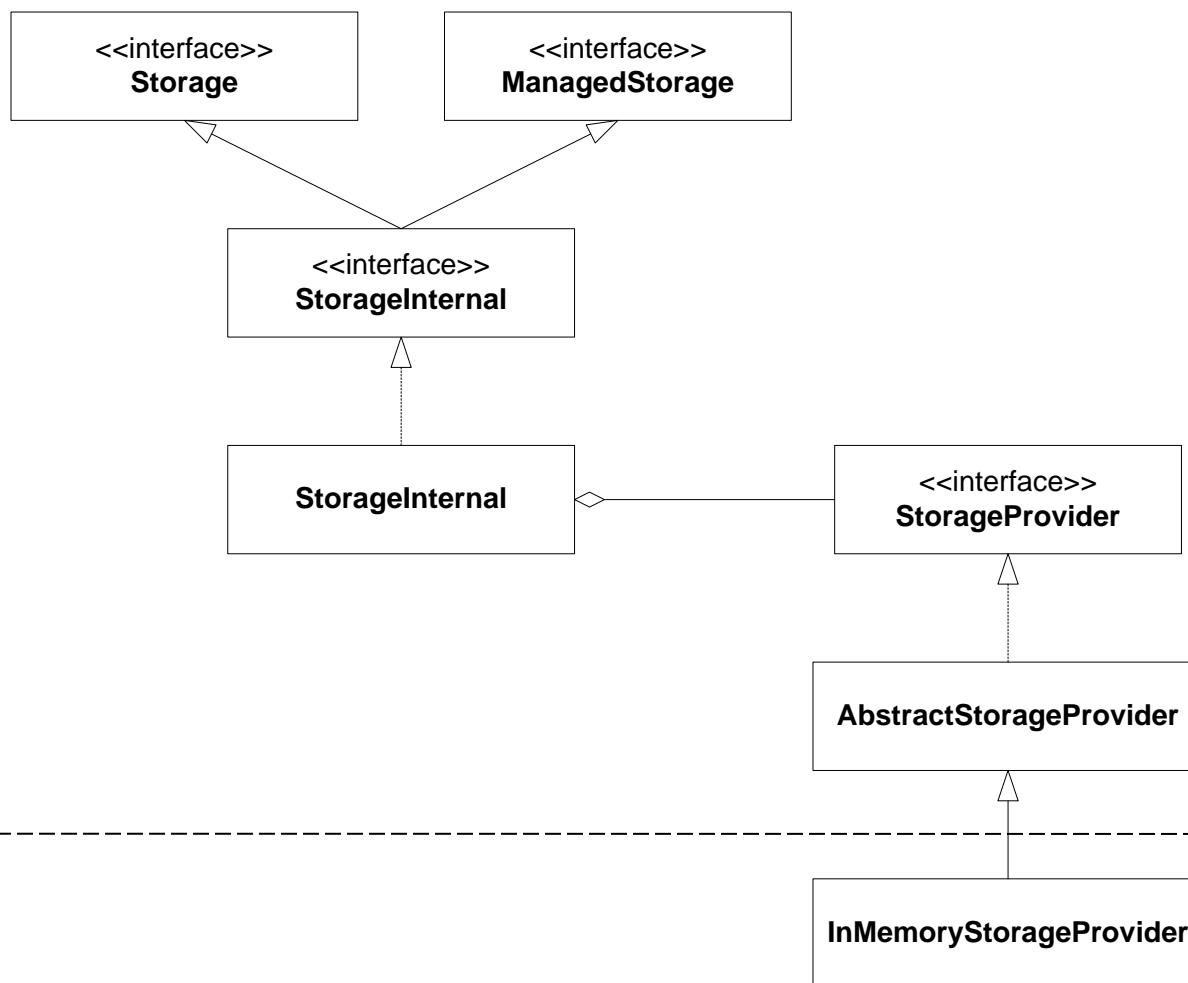
Storage Layer Concepts



Storage Layer

- Consists of three main parts
- Application Programming Interface
 - used by the model layer
 - encapsulates high-level functionality and concepts
- Internal
 - internal functionality
 - common and managed functionality
- Service Provider Interface
 - interface for storage providers that provide low-level functionality
 - db4o, Berkeley DB, in-memory

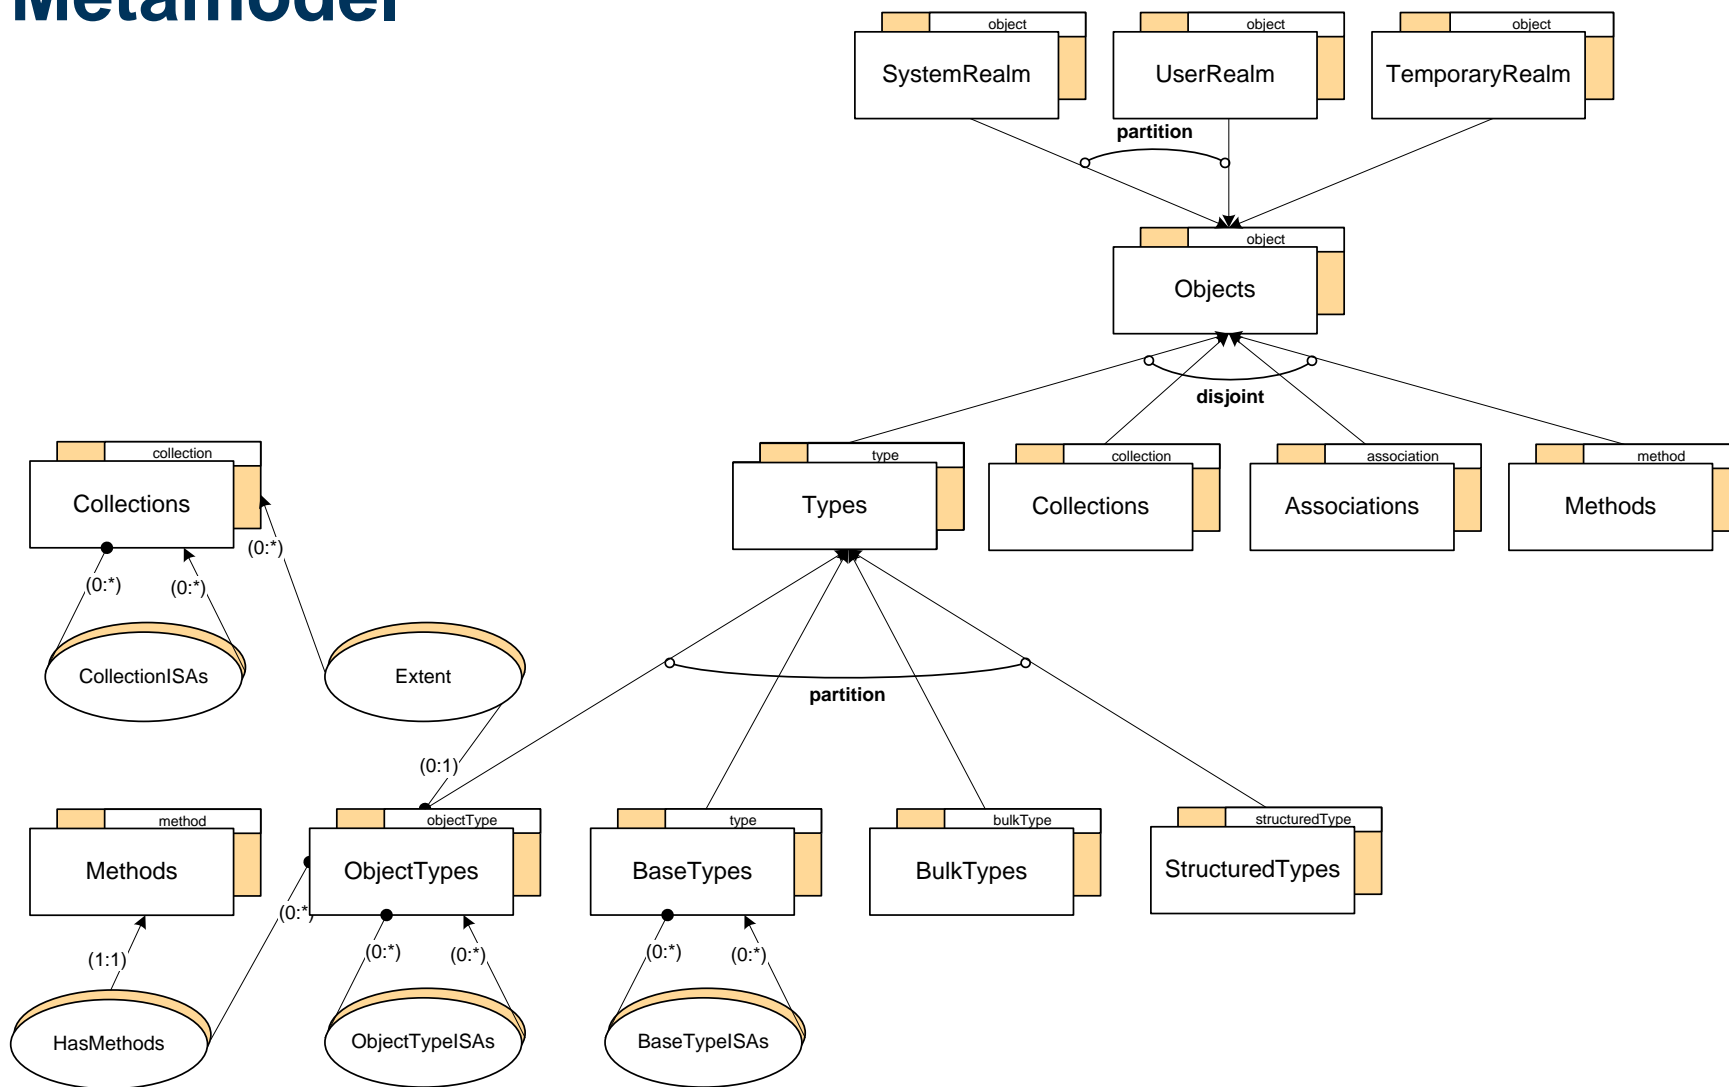
Storage Layer Design



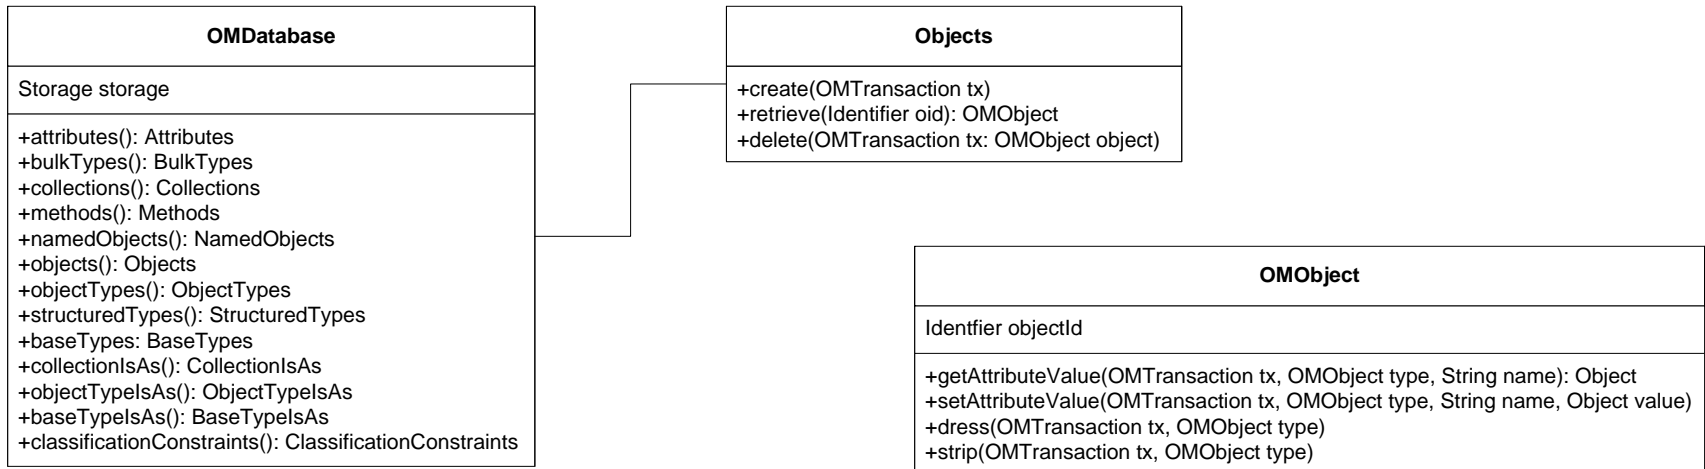
OMS Avon Model Layer

- One generic Java abstraction to represent OM objects
 - single point of extensibility
 - flexibility for database evolution
 - central control for transactions and recovery
- Classes for managing generic objects
 - create, retrieve, update and delete operations
- Utility classes to access and interpret generic objects
 - cache metadata and access data from the storage layer
- Entire OM metamodel is bootstrapped
 - metamodel is expressed using OM (metacircularity)
 - different flavours of the metamodel
 - metamodel extensibility

Metamodel



Application Programming Interface



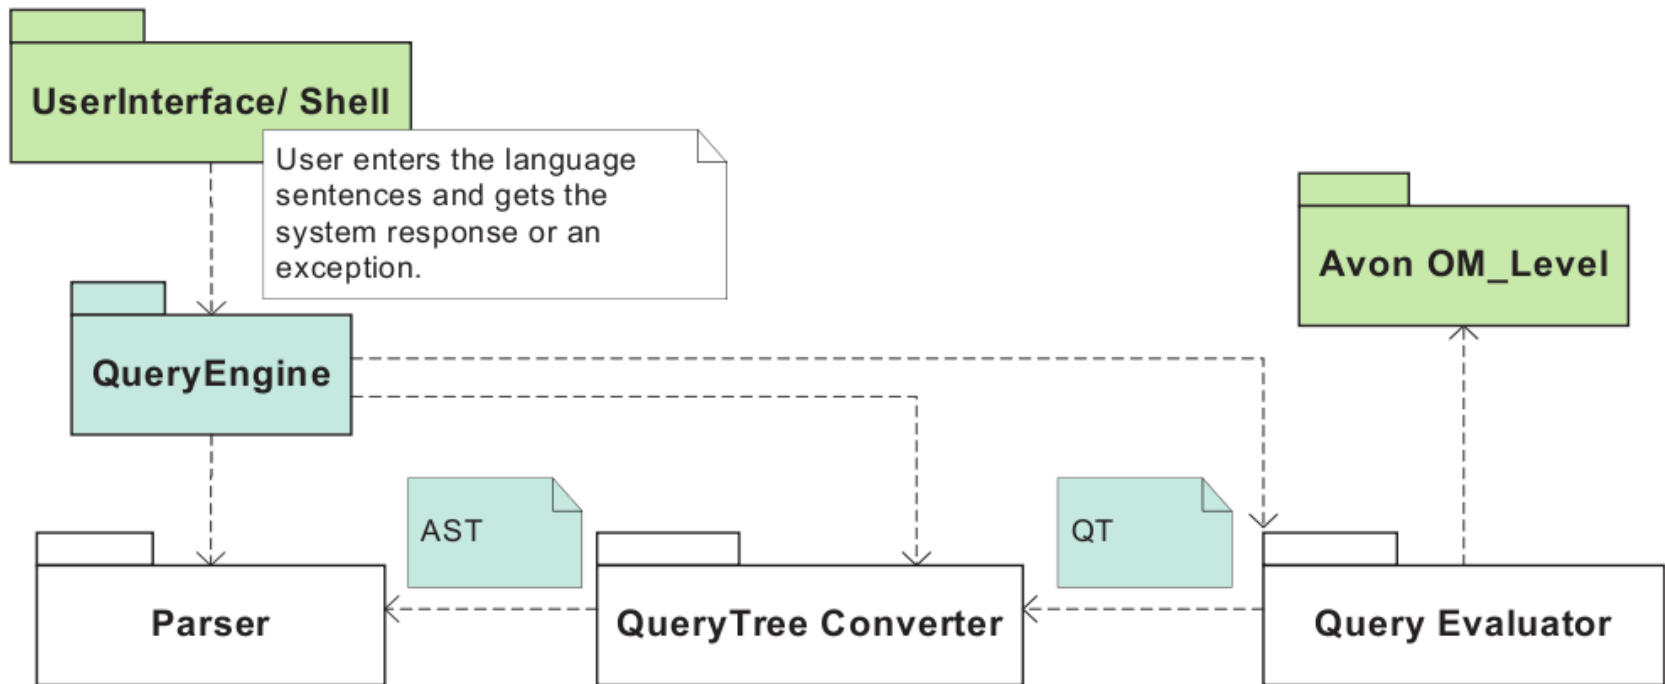
```

OMDatabase database = OMDatabaseManager.openDatabase("contacts.oms");
OMTransaction tx = database.beginTransaction();
OMOBJECT object = database.objects().create(tx);
OMOBJECT tPerson = database.namedObjects().retrieve(tx, "person");
object.dress(tx, tPerson);
OMOBJECT tContact = database.namedObjects().retrieve(tx, "contact");
object.setAttributeValue(tx, tContact, "name", "Moirra C. Norrie");
tx.commit();
  
```

Database Modules

- OMS Avon supports database modules that can extend or adapt the system for special application domains
- A database module consists of
 - metamodel extension to define new concepts
 - functionality that manages the new concepts
 - query language extension to interact with new concepts
- Existing database modules
 - main system
 - event system
 - peer-to-peer data sharing
 - personal information management
 - Web content management
 - ...

OML Query Engine

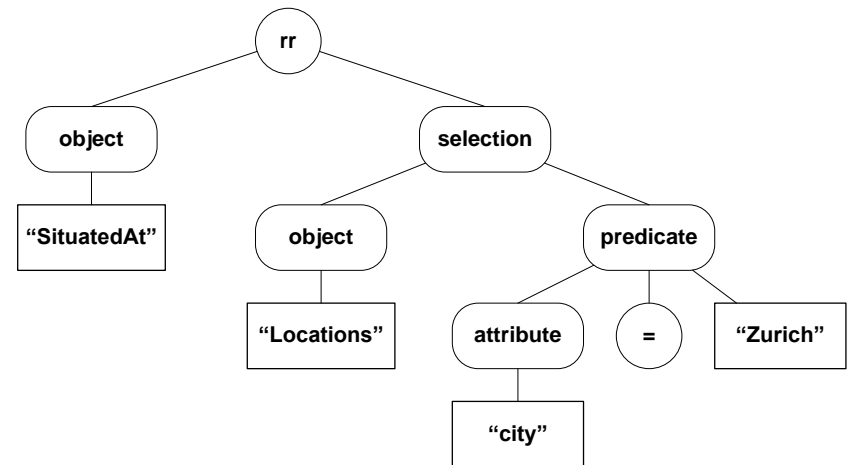


OML Query Evaluation

- **Parser**
 - used JavaCC to generate parser and lexer
 - returns an abstract syntax tree (AST)
- **Query Tree Converter**
 - uses the visitor design pattern to process AST in post-order
 - transforms the AST into a query tree (QT)
- **Query Tree Evaluator**
 - uses the visitor design pattern to process QT in post-order
 - returns only the last result from the OML script
 - stores intermediated results in the node structure

Query Tree

```
SituatedAt  
  rr(  
    all $l in Locations  
    having  
      ($l.city = "Zurich")  
  )
```



- QT nodes are atomic construct
- Used to build different database operations
 - selection, domain, range, iteration, object access

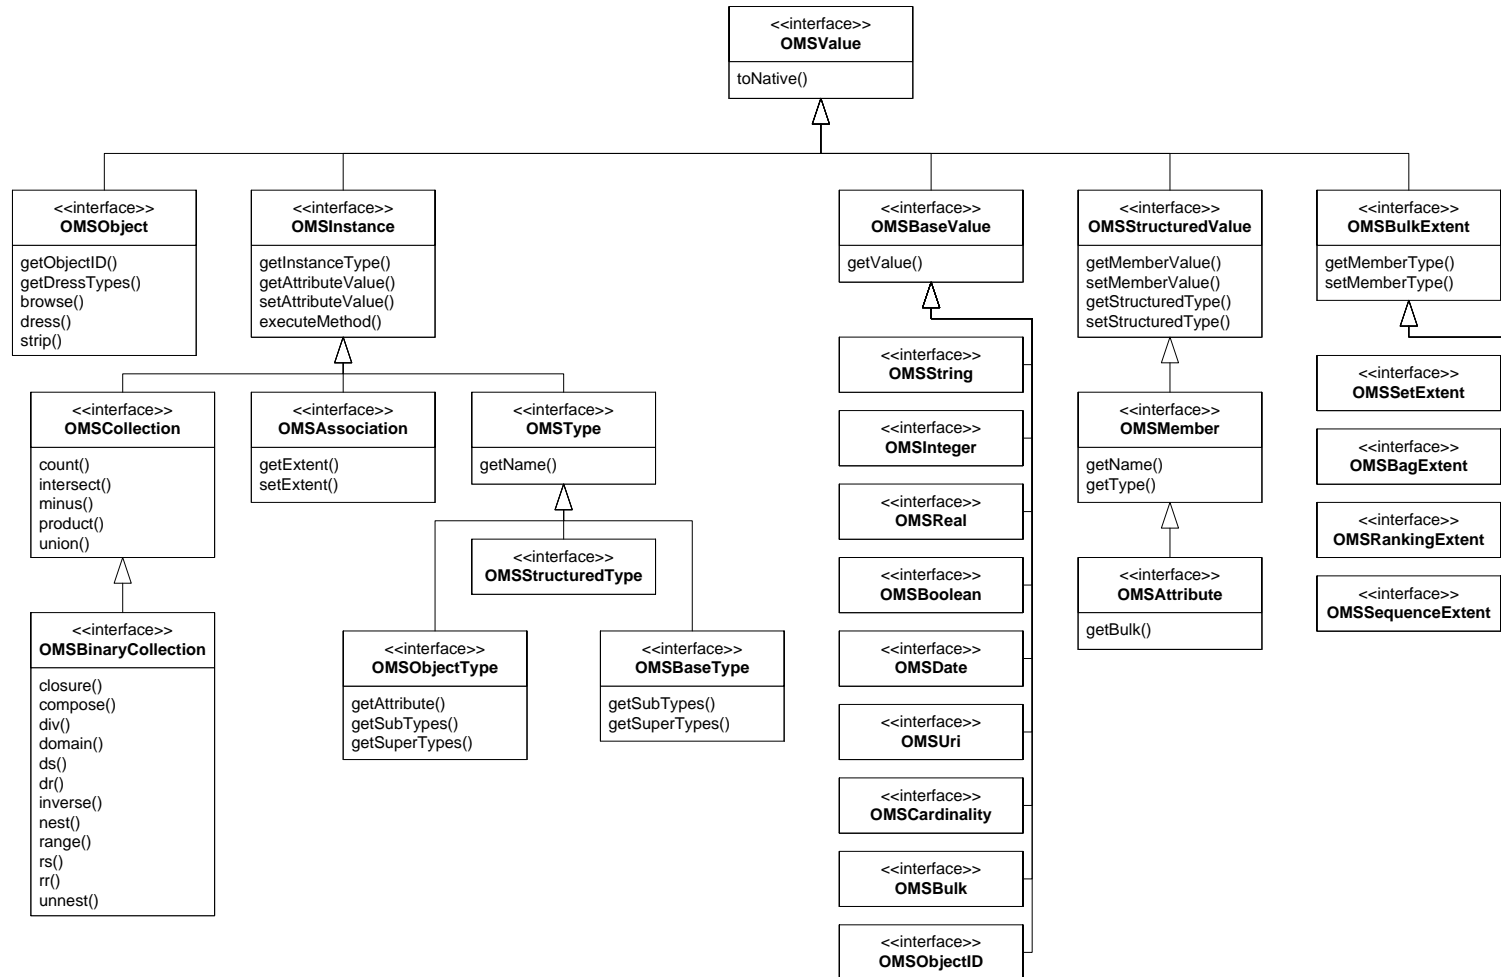
OMS Avon Interface Layer

- OMS Avon provides alternative interfaces for application development
- OMSjp
 - uniform access to heterogeneous OMS databases
 - programmatic interface based on Java
 - equivalent to JDBC in the OMS world
 - maps Java types to OM types
 - provides Java abstractions for OM system concepts
- Object Model Language (OML)
- Graphical User Interface
 - OMSjp Eclipse plug-in with graphical schema editor
 - OMSjp Browser

OMSDriver and OMSDatabase

- OMSDriver
 - provides database management functionality
 - abstraction of underlying implementation
 - one driver per supported platform
 - driver manager loads drivers based on configuration file
 - driver is configured via URL
 - **omsjp:platform://user/password@host:port/database**
- OMSDatabase
 - provides database functionality
 - create, update and delete object
 - query interface
 - schema management
 - import and export of databases

OMSjp Value Framework

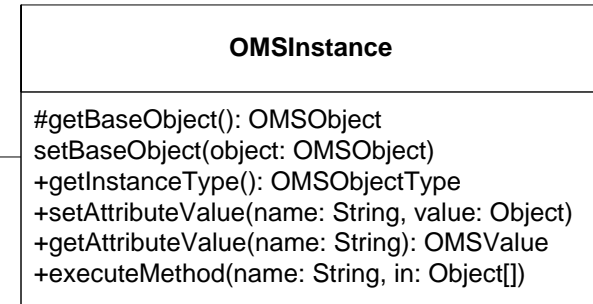
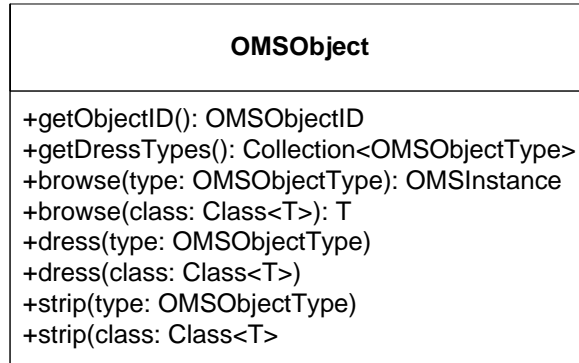


Impedance Mismatch

- Mapping the OM object-model to the Java object-model creates a new impedance mismatch
 - multiple instantiation
 - multiple inheritance
- A single object is represented by multiple classes
 - a metamodel class of type **OMSObject**
 - several model classes of type **OMSInstance**
- Application-specific instance classes can be used instead of generic instance classes
 - if data model does not multiple inheritance, Java inheritance can be used to implement application-specific model classes
 - if data model does use multiple inheritance, application-specific model classes cannot use Java inheritance

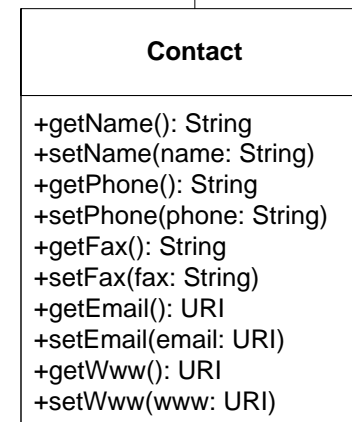
OMSObject and OMSInstance

Metamodel



...

Model



Using Application-Specific Instances

```
public class Contact extends AbstractInstance {  
    public void setName(final String name) throws OMSEException {  
        this.setAttributeValue("name", name);  
    }  
    public String getName() throws OMSEException {  
        OMSSString s = (OMSSString) this.getAttributeValue("name");  
        return s.getString();  
    }  
    ...  
}
```

- Façade design pattern
- Access
 - method browse() of class OMSSObject
 - automatically if context defines a type, e.g. in a query
- Mapping file defines instance type registrations
 - contact = ch.ethz.globis.demo.contacts.Contact

OMSjp in Action

```
// connect driver and open database
OMSDriver driver =
    OMSDriverManager.getDriver("omsjp:avon:local://localhost");
OMSDatabase db = driver.openDatabase("contacts.oms");

// retrieve and access an object
Contact contact = db.getObject(Contact.class, "name", "Moirra C. Norrie");
URI www = contact.getWww();

// dress an object and set new attribute value
Person person = contact.getBaseObject().dress(Person.class);
person.setTitle("Prof");

// retrieve a collection and perform a query
OMSBinaryCollection worksFor =
    (OMSBinaryCollection) db.getCollection("WorksFor");
Organisation organisation =
    (Organisation) worksFor.dr(contact).range().first();

// close database and disconnect driver
driver.closeDatabase();
driver.disconnect();
```

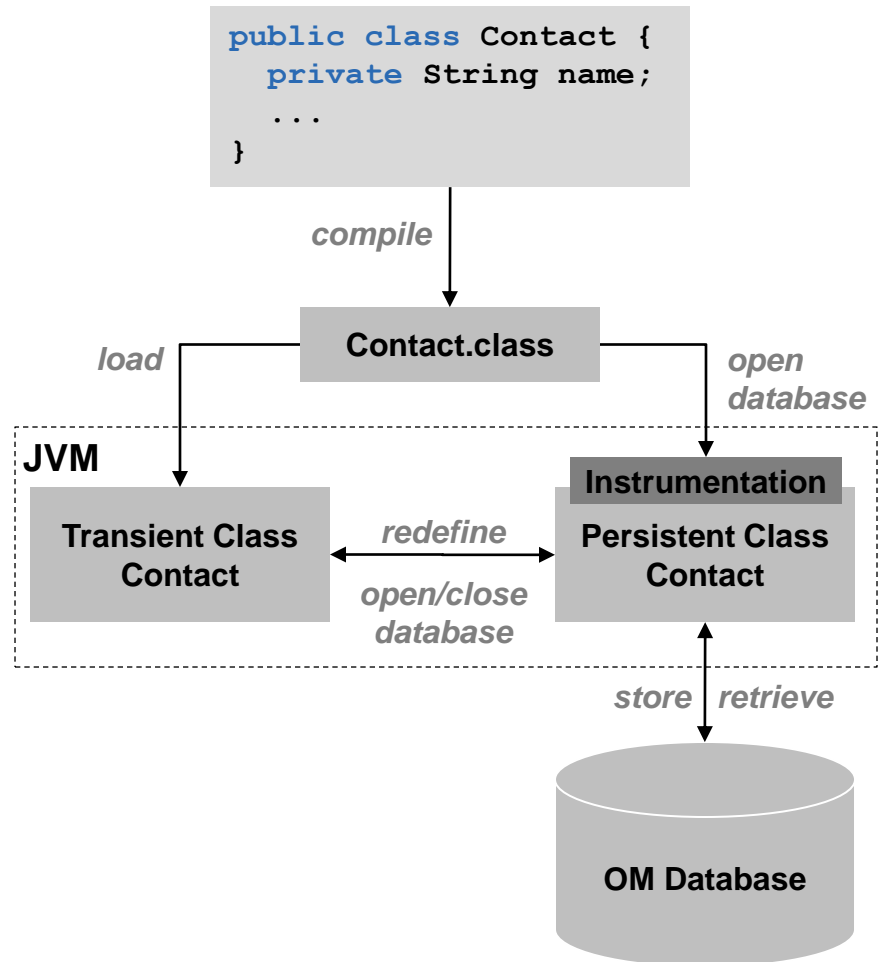
Using OML in OMSjp

```
// query for a single object
OMSValue result = db.evaluateQuery(
    "first(all $p in Persons having ($p.name = 'Moira Norrie'))";
Person moira = (Person) result;
moira.setName("Moira C. Norrie");

// query for a collection of object
OMSValue result = db.evaluateQuery("dom(SituatedAt rr(all $l in Locations
    having ($l.city = 'Zurich'))");
OMSCollection collection = (OMSCollection) result;
OMSBulkValue< ? > extent = collection.getExtent();
for (Object o: extent) {
    Contact contact = (Contact) o;
    System.out.println(contact.getName());
}
```


Dynamic Bytecode Instrumentation

- Avoid impedance mismatch to provide transparent persistence
- Code injected at run-time
 - persistence agent is registered in JVM when database is opened
 - agent transforms all classes that are loaded by the JVM
 - when database is closed, transient agent is registered
- Instrumentation supports
 - OMSjp façade pattern
 - type generation and hierarchies
 - multi-valued attributes
 - associations



Instance Instrumentation

```
public class Person implements OMSTInstance {
    private String name;
    private OMSTObject baseObject;
    public Person(OMSTObject object) {
        this.baseObject = object;
    }
    public void setName(String name) {
        this.name = name;
        this.setAttributeValue("name", this.name);
    }
    public String getName() {
        this.name = (String) this.getAttributeValue("name");
        return this.name;
    }
    public void setAttributeValue(String name, Object object) {
        baseObject.setAttributeValue(this.getClass(), name, object);
    }
    Object getAttributeValue(String name) {
        return baseObject.getAttributeValue(this.getClass(), name);
    }
}
```

Class Hierarchy Instrumentation


Java Object Model

OM Data Model

Static


```
public class Contact {
    private String name;
    ...
}
```

```
public class Person extends Contact {
    private String title;
    ...
}
```



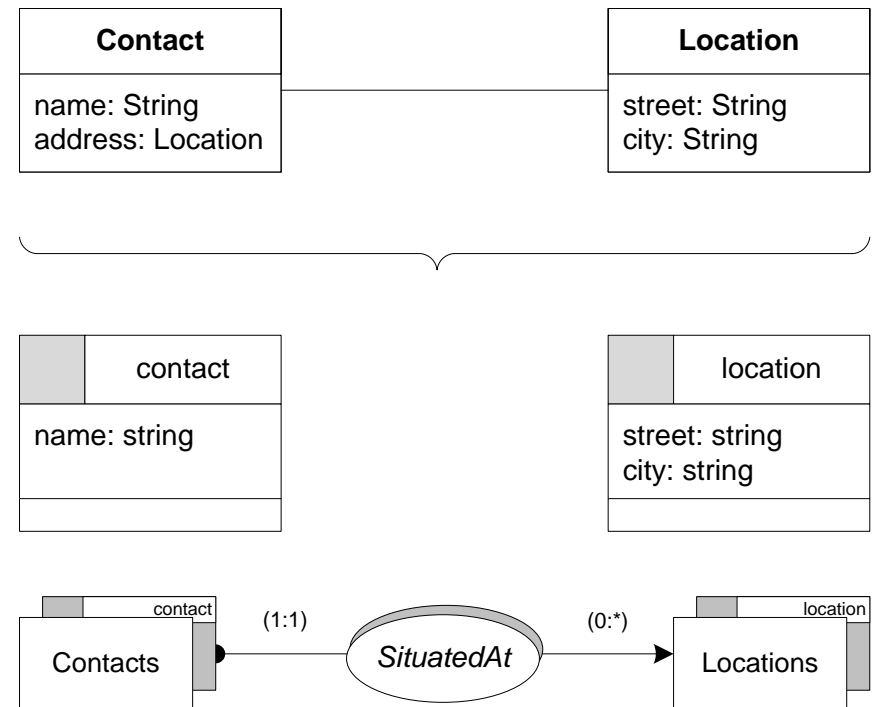
Dynamic

```
Person christoph = new Person();
christoph.setName("Christop Lins");
christoph.setTitle("MSc");
```



Type and Association Instrumentation

- Type creation
 - existence of OM type is checked using the fully qualified Java type name
 - Java class analysed using reflection
 - Java types mapped to OM types
 - every type gets an extent collection
- Object references
 - OM types are created for properties that reference non-basic and non-collection types
 - persistence by reachability
 - agent can be configured to create associations instead of references to raise semantic expressiveness



Collection Instrumentation

- Multi-valued properties are represented using OM extents
 - `Set<T>` is mapped to `OMSSet<T>`
 - `List<T>` is mapped to `OMSSequence<T>`
 - other collection types are currently not supported by the agent
- OM extents provide support for
 - collection behaviour
 - collection algebra
 - management of large collections

```
public class Person implements OMSInstance {  
  
    private final Set<String> topics;  
    private final List<String> tasks;  
  
    public Person() {  
        this.topic = new OMSSet<String>();  
        this.tasks = new OMSSequence<String>();  
    }  
  
    public void addTopic(String topic) {  
        topic.add(topic);  
    }  
  
    ...  
}
```

Conventions

■ Classes

- must not be abstract
- must have a public default constructor

■ Properties

- must be accessible through getter and setter accessor methods that follow the Java Beans convention
 - `public Type getProperty()`
 - `public void setProperty(Type)`
- all other method names must not contain the **get** and **set** prefixes
- keyword **transient** can be used for non-persistent properties
- Arrays are not supported as class members
- List and Set are supported, all other Java Collections are forbidden.
- direct initialisation of class members as part of their declaration is not supported

Seamless Access to OM Functionality

```
// connect driver and open database
OMSDriver driver =
    OMSDriverManager.getDriver("omsjp:avon:local://localhost");
OMSDatabase db = driver.openDatabase("contacts.oms");

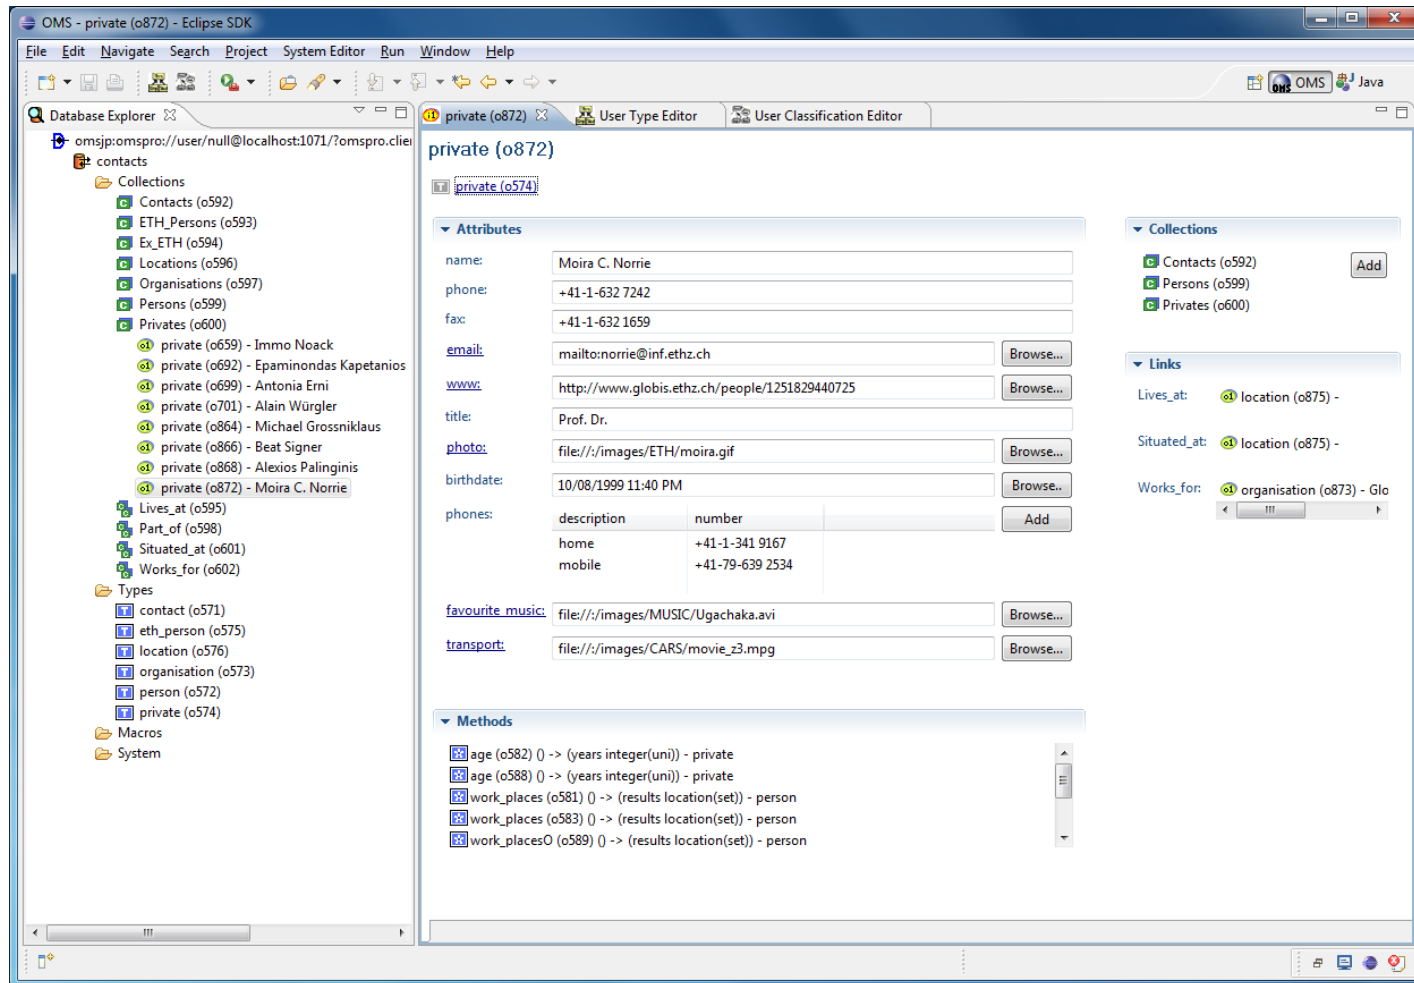
// create a persistent person object
Person person = new Person();
person.setName("Christoph Lins");

// evolve person to a private contact
OMSObject object = ((OMSInstance) person).getBaseObject();
Private private = object.dress(Private.class);
private.setMusic("http://www.youtube.com/watch?v=kWKWYuZvmhc");

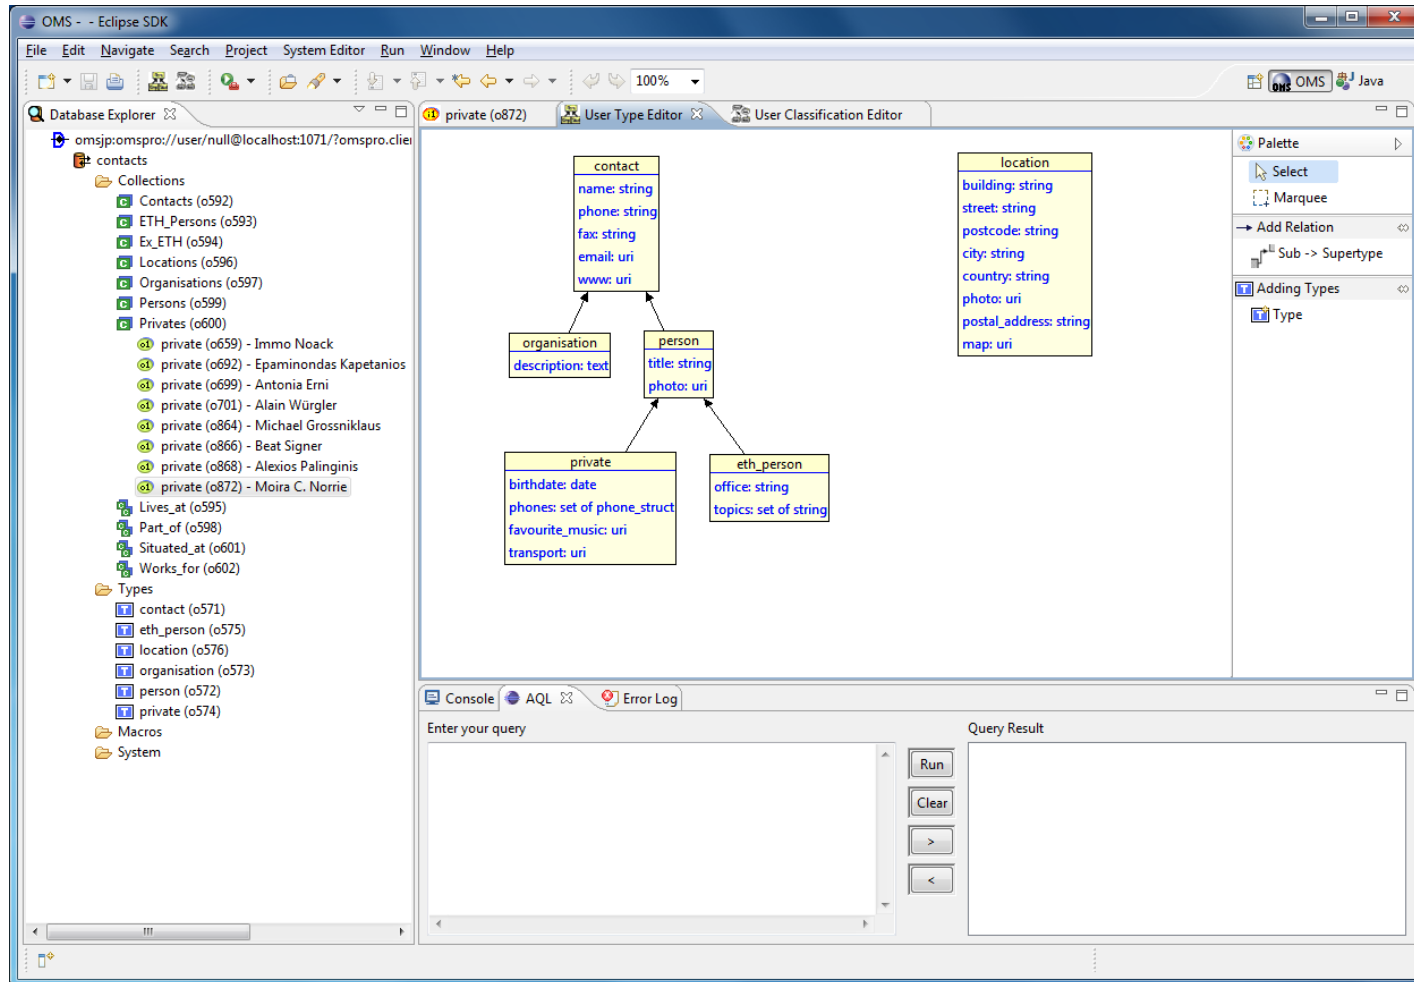
// create a transient person object
driver.closeDatabase();
person = new Person();
person.setName("Barney Stinson");

// disconnect driver
driver.disconnect();
```

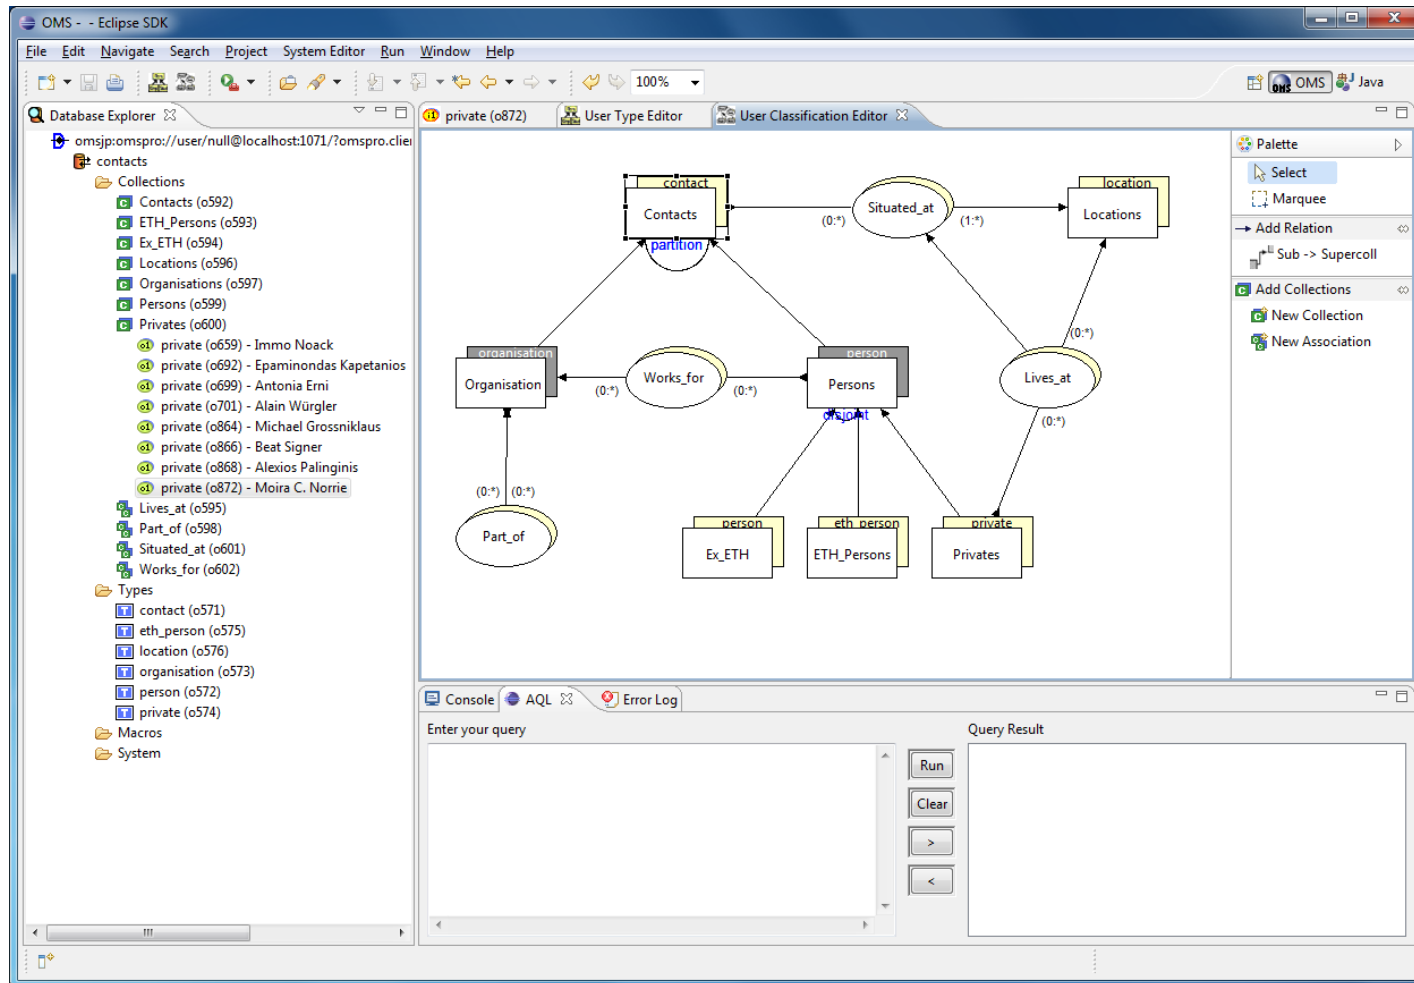
OMSjp Eclipse Plug-in: Object Browser



OMSjp Eclipse Plug-in: Type Editor



OMSjp Eclipse Plug-in: Classification Editor



References

- M. C. Norrie, M. Grossniklaus, C. Decurtins, A. de Spindler, A. Vancea and S. Leone: **Semantic Data Management for db4o**, In: *Proceedings of ICOODB, 2008*
- M. Grossniklaus: **OMSjp – A Uniform Java Interface to Heterogenous OMS Platforms**, *White Paper, 2005*
- B. Aeppli: **Eclipse-Based Front-End for OMSjp**, *Diploma Thesis, 2005*
- C. Schmid: **Design and Development of an Eclipse Schema Editor for OMSjp**, *Diploma Thesis, 2005*
- C. Lins: **Exploiting Dynamic Bytecode Instrumentation to Support Transparent Persistence**, *Semester Thesis, 2009*

Next Week

Course Review

- Exam Information
- Summary
- OODBMS Architectures

