

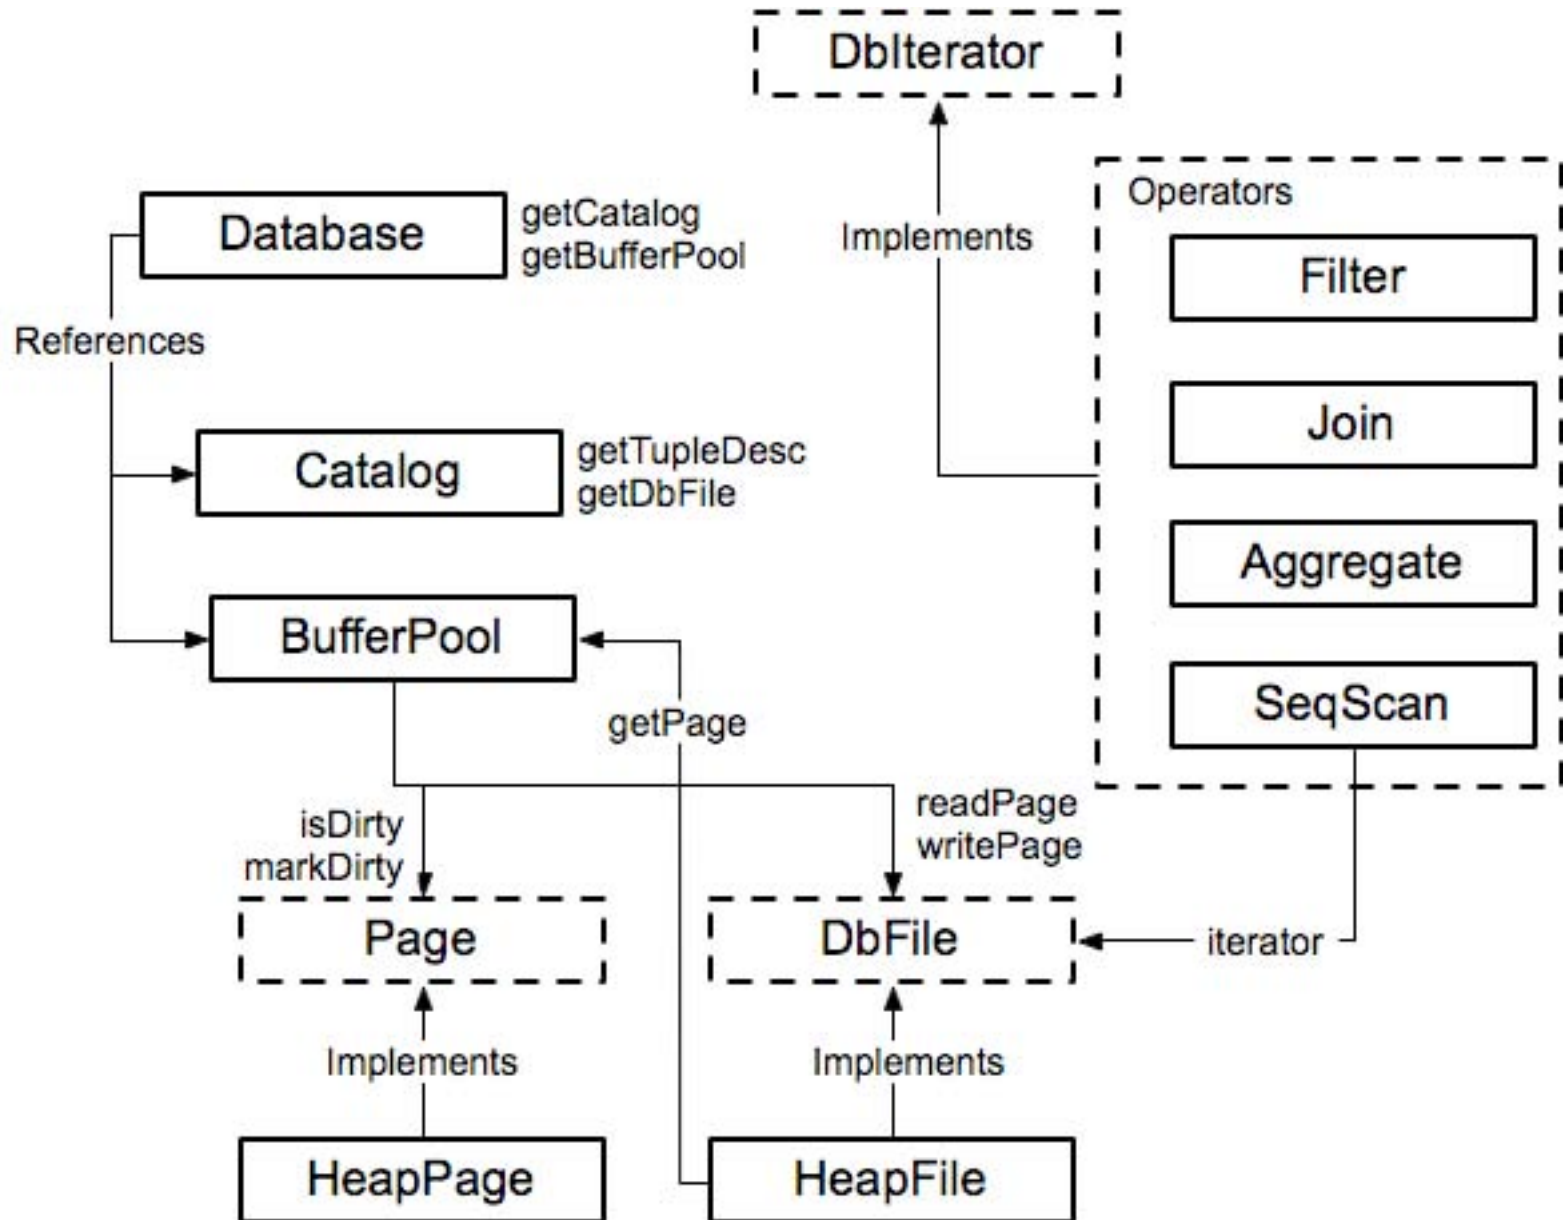
SimpleDB Overview

9/18/2008

What is SimpleDB?

- A basic database system
- What it has
 - Heapfiles
 - Basic Operators (Scan, Filter, JOIN, Aggregate)
 - Buffer Pool
 - Transactions
 - SQL Front-end
- Things it doesn't have
 - Query optimizer
 - Fancy relational operators (UNION, etc)
 - Recovery
 - Indices

Module Diagram



Catalog

- Catalog stores a list of available tables, TupleDesc
 - void addTable(DbFile d, TupleDesc d)
 - DbFile getTable(int tableid)
 - TupleDesc getTupleDesc(int tableid)
- Not persisted to disk

DbIterator.java

- Iterator class implemented by all operators
 - open()
 - close()
 - getTupleDesc()
 - hasNext()
 - next()
 - rewind()
- Iterator model: chain iterators together

```
// construct a 3-column table schema
```

```
Type types[] = new Type[]{ Type.INT_TYPE, Type.INT_TYPE, Type.INT_TYPE };
```

```
String names[] = new String[]{ "field0", "field1", "field2" };
```

```
TupleDesc descriptor = new TupleDesc(types, names);
```

```
// create the table, associate it with some_data_file.dat
```

```
// and tell the catalog about the schema of this table.
```

```
HeapFile table1 = new HeapFile(new File("some_data_file.dat"), descriptor);
```

```
Database.getCatalog().addTable(table1);
```

```
// construct the query: we use a simple SeqScan, which spoonfeeds
```

```
// tuples via its iterator.
```

```
TransactionId tid = new TransactionId();
```

```
SeqScan f = new SeqScan(tid, table1.id());
```

```
// and run it
```

```
f.open();
```

```
while (f.hasNext()) {
```

```
    Tuple tup = f.next();
```

```
    System.out.println(tup);
```

```
}
```

```
f.close();
```

```
Database.getBufferPool().transactionComplete();
```

HeapFile.java

- An array of HeapPages on disk
- Javadoc is your friend!
- Implement everything except addTuple and removeTuple

HeapPage.java

- Format
 - Header is a bitmap
 - Page contents are an array of fixed-length Tuples
- Full page size = `BufferPool.PAGE_SIZE`
- Number of bits in Header = number of Tuples
- Header size + size of tuples = `BufferPool.PAGE_SIZE`

HeapFileEncoder.java

- Because you haven't implemented insertTuple, you have no way to create data files
- HeapFileEncoder converts CSV files to HeapFiles
- Usage:
 - `java -jar dist/simpledb.jar convert csv-file.txt numFields`
- Produces a file `csv-file.dat`, that can be passed to HeapFile constructor.

BufferPool.java

- Manages cache of pages
 - Evicts pages when cache is full [not lab 1]
- All page accesses should use getPage
 - Even from inside DbFile!

You will eventually implement

- locking for transactions
- Flushing of pages for recovery

Compiling, Testing, and Running

- Compilation done through the ant tool
 - Works a lot like make
- Two kinds of tests:
 - Unit tests
 - System Tests
- Demo on debugging using unit tests.

MIT OpenCourseWare
<http://ocw.mit.edu>

6.830 / 6.814 Database Systems
Fall 2010

For information about citing these materials or our Terms of Use, visit: <http://ocw.mit.edu/terms>.