



Programming Assignment #1 Part 1
CAS CS 460: Introduction to Database Systems

Due Date and Time: 09/24, 23:59 on gradescope.

In this assignment you will implement parts of a basic database management system called SimpleDB. For this assignment, you will focus on implementing the core modules required to access stored data on disk; in future, you will add support for various query processing operators, as well as transactions, locking, and concurrent queries.

SimpleDB is written in Java. We have provided you with a set of mostly unimplemented classes and interfaces. You will need to write the code for these classes. We will grade your code by running a set of system tests written using JUnit. We have provided a number of unit tests, which we will not use for grading but that you may find useful in verifying that your code works.

The remainder of this document describes the basic architecture of SimpleDB, gives some suggestions about how to start coding, and discusses how to hand in the assignment.

We recommend that you start as early as possible on this assignment.

1. Getting Started

These instructions are written for Unix-based platform (e.g., Linux, MacOS, etc.) Because the code is written in Java, it should work under Windows as well; however, some directions in this document may not apply. (Java 1.8.* is required)

Download the file from [HERE](#)

SimpleDB uses the Ant build tool to compile the code and run tests. Ant is similar to make, but the build file is written in XML and is somewhat better suited to Java code. Most modern Linux distributions include Ant.

To help you during development, we have provided some unit testing cases. Again, you should not rely on them exclusively to verify the correctness of your project, they are by no means comprehensive.

To run the unit tests, go to terminal and type the following commands:

```
$ cd CS460-pa1  
$ # run all unit tests  
$ ant test  
$ # run a specific unit test  
$ ant runtest -Dtest=TupleTest
```

You should see output similar to:

```
# build output...  
test:  
 [junit] Running simpledb.CatalogTest  
 [junit] Testsuite: simpledb.CatalogTest  
 [junit] Tests run: 5, Failures: 4, Errors: 1, Time elapsed: 0.037 sec [junit] Tests run: 5, Failures: 4,  
 Errors: 1, Time elapsed: 0.037 sec  
# ... stack traces and error reports ...
```

The output above indicates problems occurred during compilation; this is because the code we have given you doesn't yet work. As you complete parts of the assignment, you will work towards passing additional unit tests. If you wish to write new unit tests as you code, they should be added to the test/simpledb directory.

1.1. Running System Tests

We have also provided tests that will eventually be used for grading. These tests are structured as JUnit tests that live in the test/simpledb/systemtest directory. To run all of the system tests, use the systemtest build target:

If the tests pass, you will see something like the following:

```
$ ant systemtest
# ... build output ...
[junit] Testsuite: simpledb systemtest ScanTest

[junit] Tests run: 3, Failures: 0, Errors: 0, Time elapsed: 7.278 sec
[junit] Tests run: 3, Failures: 0, Errors: 0, Time elapsed: 7.278 sec
[junit] Testcase: testSmall took 0.937 sec
[junit] Testcase: testLarge took 5.276 sec
[junit] Testcase: testRandom took 1.049 sec
```

```
BUILD SUCCESSFUL
Total time: 52 seconds
```

1.2. Implementation hints

Before beginning to write code, we **encourage** you to read through this entire document to get a feel for the high-level design of SimpleDB. You will need to fill in any piece of code that is not implemented. It will be obvious where we think you should write code. You may add private methods and/or helper classes. You may change APIs, but make sure our grading tests still run and make sure to mention, explain your decisions in your writeup.

In addition to the methods that you need to fill out for this assignment, the class interfaces contain numerous methods that you need not implement until subsequent assignments. These will either be indicated per class:

```
///Not necessary for lab1.
public class Insert implements DbIterator {
```

or per method:

```
public boolean deleteTuple(Tuple t) throws DbException {
    //some code goes here
    //not necessary for lab1
    return false;
}
```

The code that you submit should compile without having to modify these methods

1.3. Transactions, locking, and recovery

As you look through the interfaces, we have provided you, you will see a number of references to locking, transactions, and recovery. You do not need to support these features here, but you should keep these parameters in the interfaces of your code

because you will be implementing transactions and locking in a future assignment. The test code we have provided you with generates a fake transaction ID that is passed into the operators of the query it runs; you should pass this transaction ID into other operators and the buffer pool.

2. SimpleDB Architecture and Implementation Guide

SimpleDB consists of:

- Classes that represent fields, tuples, and tuple schemas;
- Classes that apply predicates and conditions to tuples;
- One or more access methods (e.g., heap files) that store relations on disk and provide a way to iterate through tuples of those relations;
- A collection of operator classes (e.g., select, join, insert, delete, etc.) that process tuples;
- A buffer pool that caches active tuples and pages in memory and handles concurrency control and transactions (neither of which you need to worry about for now); and,
- A catalog that stores information about available tables and their schemas.

SimpleDB does not include many things that you may think of as being a part of a "database." In particular, SimpleDB does not have:

- (In this assignment), a SQL front end or parser that allows you to type queries directly into SimpleDB.
- Views.
- Data types except integers and fixed length strings.
- (In this assignment) Query optimizer.
- Indices.

In the rest of this Section, we describe each of the main components of SimpleDB that you will need to implement in this assignment. You should use the exercises in this discussion to guide your implementation. This document is by no means a complete specification for SimpleDB; you will need to make decisions about how to design and implement various parts of the system. Note that for this assignment you do not need to implement any operators (e.g., select, join, project) except sequential scan. You will add support for additional operators in future assignments.

2.1. The Database Class

The Database class provides access to a collection of static objects that are the global state of the database. In particular, this includes methods to access the catalog (the list of all the tables in the database), the buffer pool (the collection of database file pages that are currently resident in memory), and the log file. You will not need to worry about the log file in this assignment. We have implemented the Database class for you. You should take a look at this file as you will need to access these objects.

2.2. Fields and Tuples

Tuples in SimpleDB are quite basic. They consist of a collection of Field objects. The Tuple.Field is an interface that different data types (e.g., integer, string) implement. Tuple objects are created by the underlying access methods (e.g., heap files, or B-trees), as described in the next section. Tuples also have a type (or schema), called a tuple descriptor, represented by a TupleDesc object. This object consists of a collection of Type objects, one per field in the tuple, each of which describes the type of the corresponding field.

Exercise 1. Implement the skeleton methods in:

[src/simpledb/TupleDesc.java](#)



src/simpledb/Tuple.java

At this point, your code should pass the unit tests **TupleTest** and **TupleDescTest**.

2.3. Catalog

The catalog (class `Catalog` in `SimpleDB`) consists of a list of the tables and schemas of the tables that are currently in the database. You will need to support the ability to add a new table, as well as getting information about a particular table. Associated with each table is a `TupleDesc` object that allows operators to determine the types and number of fields in a table.

The global catalog is a single instance of `Catalog` that is allocated for the entire `SimpleDB` process. The global catalog can be retrieved via the method `Database.getCatalog()`, and the same goes for the global buffer pool (using `Database.getBufferPool()`).

Exercise 2. Implement the skeleton methods in:

src/simpledb/Catalog.java

At this point, your code should pass the unit tests in **CatalogTest**.

2.4. BufferPool

The buffer pool (class `BufferPool` in `SimpleDB`) is responsible for caching pages in memory that have been recently read from disk. All operators read and write pages from various files on disk through the buffer pool. It consists of a fixed number of pages, defined by the `numPages` parameter to the `BufferPool` constructor. In later assignments, you will implement an eviction policy. For this one, you only need to implement the constructor and the `BufferPool.getPage()` method used by the `SeqScan` operator. The `BufferPool` should store up to `numPages` pages. For this assignment, if more than `numPages` requests are made for different pages, then instead of implementing an eviction policy, you may throw a `DbException`. In future you will be required to implement an eviction policy.

The `Database` class provides a static method, `Database.getBufferPool()`, that returns a reference to the single `BufferPool` instance for the entire `SimpleDB` process.

Exercise 3. Implement the `getPage()` method in:

src/simpledb/BufferPool.java

We have not provided unit tests for `BufferPool`. The functionality you implemented will be tested in the implementation of `HeapFile` below. You should use the `DbFile.readPage` method to access pages of a `DbFile`.

2.5. HeapFile access method

Access methods provide a way to read or write data from disk that is arranged in a specific way. Common access methods include heap files (unsorted files of tuples) and B-trees; for this assignment, you will only implement a heap file access method, and we have written some of the code for you.

A `HeapFile` object is arranged into a set of pages, each of which consists of a fixed number of bytes for storing tuples, (defined by the constant `BufferPool.PAGE_SIZE`), including a header. In `SimpleDB`, there is one `HeapFile` object for each table in the database. Each page in a `HeapFile` is arranged as a set of slots, each of which can hold one tuple (tuples for a given table in `SimpleDB` are all of the same size). In addition to these slots, each page has a header that consists of a bitmap with one bit per tuple slot. If the bit corresponding to a particular tuple is 1, it indicates that the tuple is valid; if it is 0, the tuple is invalid (e.g.,



has been deleted or was never initialized.) Pages of HeapFile objects are of type HeapPage which implements the Page interface. Pages are stored in the buffer pool but are read and written by the HeapFile class.

SimpleDB stores heap files on disk in more or less the same format they are stored in memory. Each file consists of page data arranged consecutively on disk. Each page consists of one or more bytes representing the header, followed by the page size bytes of actual page content. Each tuple requires tuple size * 8 bits for its content and 1 bit for the header. Thus, the number of tuples that can fit in a single page is:

$$\text{tuples per page} = \text{floor}((\text{page size} * 8) / (\text{tuple size} * 8 + 1))$$

Where tuple size is the size of a tuple in the page in bytes. The idea here is that each tuple requires one additional bit of storage in the header. We compute the number of bits in a page (by multiplying page size by 8), and divide this quantity by the number of bits in a tuple (including this extra header bit) to get the number of tuples per page. The floor operation rounds down to the nearest integer number of tuples (we don't want to store partial tuples on a page!). Once we know the number of tuples per page, the number of bytes required to store the header is simply:

$$\text{headerBytes} = \text{ceiling}(\text{tupsPerPage}/8)$$

The ceiling operation rounds up to the nearest integer number of bytes (we never store less than a full byte of header information.)

The low (least significant) bits of each byte represents the status of the slots that are earlier in the file. Hence, the lowest bit of the first byte represents whether or not the first slot in the page is in use.

Also, note that the high-order bits of the last byte may not correspond to a slot that is actually in the file, since the number of slots may not be a multiple of 8. Also note that all Java virtual machines are big-endian.

Exercise 4. Implement the skeleton methods in:

src/simpledb/HeapPageId.java

src/simpledb/RecordID.java

src/simpledb/HeapPage.java

Although you will not use them directly now, we ask you to implement `getNumEmptySlots()` and `isSlotUsed()` in `HeapPage`. These require pushing around bits in the page header. You may find it helpful to look at the other methods that have been provided in `HeapPage` or in `src/simpledb/HeapFileEncoder.java` to understand the layout of pages.

You will also need to implement an Iterator over the tuples in the page, which may involve an auxiliary class or data structure.

At this point, your code should pass the unit tests in **HeapPageIdTest**, **RecordIdTest**, and **HeapPageReadTest**.

After you have implemented `HeapPage`, you will write methods for `HeapFile` to calculate the number of pages in a file and to read a page from the file. You will then be able to fetch tuples from a file stored on disk.

Exercise 5. Implement the skeleton methods in:

src/simpledb/HeapFile.java

To read a page from disk, you will first need to calculate the correct offset in the file. Hint: you will need random access to the file in order to read and write pages at arbitrary offsets. You should not call `BufferPool` methods when reading a page from disk.

You will also need to implement the `HeapFile.iterator()` method, which should iterate through through the tuples of each page in the `HeapFile`. The iterator must use the `BufferPool.getPage()` method to access pages in the `HeapFile`. This method loads the page into the buffer pool and will eventually be used (in a later assignment) to implement locking-based concurrency control and recovery. Do not load the entire table into memory on the `open()` call -- this will cause an out of memory error for very large tables.

At this point your code should pass the unit tests in **HeapFileReadTest**.

2.6. Operators

Operators are responsible for the actual execution of the query plan. They implement the operations of the relational algebra. In `SimpleDB`, operators are iterator based; each operator implements the `DbIterator` interface.

Operators are connected together into a plan by passing lower-level operators into the constructors of higher-level operators, i.e., by 'chaining them together.' Special access method operators at the leaves of the plan are responsible for reading data from the disk (and hence do not have any operators below them).

At the top of the plan, the program interacting with `SimpleDB` simply calls `getNext` on the root operator; this operator then calls `getNext` on its children, and so on, until these leaf operators are called. They fetch tuples from disk and pass them up the tree (as return arguments to `getNext`); tuples propagate up the plan in this way until they are output at the root or combined or rejected by another operator in the plan.

For this time, you will only need to implement one `SimpleDB` operator.

Exercise 6. Implement the skeleton methods in:

[src/simpledb/SeqScan.java](#)

This operator sequentially scans all the tuples from the pages of the table specified by the `tableid` in the constructor. This operator should access tuples through the `DbFile.iterator()` method.

At this point, you should be able to complete the **ScanTest system test**.

Good work! You will fill in other operators in subsequent assignments.

3. Logistics

You must submit your code (see below) as well as a short (2 pages, maximum) writeup describing your approach. This writeup should:

- Describe any design decisions you made. These may be minimal for pa1.
- Discuss and justify any changes you made to the API.
- Describe any missing or incomplete elements of your code.
- Describe how long you spent on the assignment, and whether there was anything you found particularly difficult or confusing.



3.1. Collaboration

This assignment should be manageable for a single person, but if you prefer to work with a partner, this is also OK (which is the default setup in our class). Larger groups than 2 are not allowed. Please indicate clearly who you worked with, if anyone, on your individual writeup.

3.2. Submitting your assignment

Please only submit source code files on gradescope! Do not submit .class files or anything else. If you worked in a **group** make a **GROUP** submission. Also **do not forget** to mention the names of the **people** that collaborated in the write-up!

3.3 Grading

85% of your grade will be based on whether or not your code passes the system test when we run over it. These tests will be a superset of the tests we have provided. Before handing in your code, you should make sure your code produces no errors from both **ant test** and **ant systemtest**.

Before testing, we will replace your build.xml and the entire contents of the test directory with our version of these files This means you cannot change the format of dat files You should also be careful changing our APIs. You should test that your code compiles the unmodified tests.

It will look roughly like this:

```
$ tar xvf cs660-pa1.tar.gz $  
cd ./CS460-pa1  
[replace build.xml and test] $ ant test  
$ ant systemtest  
[additional tests]
```

If any of these commands fail, we'll be unhappy, and, therefore, so will your grade. An additional 15% of your grade will be based on the quality of your writeup and our subjective evaluation of your code.

3.3 Academic Honesty

Assignments must be completed by **you** or **your group**. Representing the work of another person as your own is expressly forbidden. This includes "borrowing", "stealing", copying programs/solutions or parts of them from others. We **will** use an automated plagiarism checker. Cheating will not be tolerated under any circumstances.

See the [CAS Academic Conduct Code](#), in particular regarding plagiarism and cheating on exams. A student suspected to violate this code will be reported to the Academic Conduct Committee, and if found culpable, the student will receive a grade of "F" for the course.

We hope you enjoy hacking on this assignment!