Lab 7
Due 9 April

Handout 9
CSCI 136: Spring, 2007
2 April

The Two Towers

1 Short Answer

Write an iterator that iterates over the characters of a `String`. Repeated calls to the iterator's `next()` method will return the first character of the string, then the second character, and so on.

More specifically, you are to complete the design of the class `CharacterIterator` and implement the listed constructor and methods:

```
public class CharacterIterator extends AbstractIterator<Character> {

  public CharacterIterator(String str) { ... }
  public Character next() { ... }
  public boolean hasNext() { ... }
  public void reset() { ... }
  public Character get() { ... }
}
```

The `Character` class is the wrapper class for `char` values. You use it much like we use `Integer` for `int` values. The Java compiler will automatically convert `char` values to `Character` objects when necessary. **Hand in your solution at the beginning of lab on Wednesday.**

2 Lab Program

(This is a version of the Lab from Chapter 8, modified to use generic classes.)

Goal. To solve a difficult problem using Iterators.

Discussion. Suppose that we are given n uniquely sized cubic blocks and that each block has a face area between 1 and n. If we build two towers by stacking these blocks, how close can we get their heights? The following two towers built by stacking 15 blocks, for example, differ in height by only 129 millions of an inch (each unit is one-tenth of an inch):

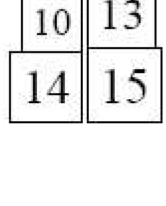

Still, this stacking is only the *second-best* solution! To find the best stacking, we could consider all the possible configurations. We do know one thing: the total height of the two towers is computed by summing the heights of all the blocks:

$$h = \sum_{i=1}^{n} \sqrt{i}$$

If we consider all the subsets of the n blocks, we can think of the subset as the set of blocks that make up, say, the left tower. We need only keep track of that subset that comes closest to $h/2$ without exceeding it.

In this lab, we will represent a set of n distinct objects by a `Vector`, and we will construct an `Iterator` that returns each of the 2^n subsets.

Procedure The trick to understanding how to generate a subset of n values from a `Vector` is to first consider how to generate a subset of indices of elements from 0 to $n-1$. Once this simpler problem is solved, we can use the indices to help us build a Vector (or subset) of values identified by the indices.

There are exactly 2^n subsets of values 0 to $n-1$. We can see this by imagining that a coin is tossed n times—once for each value—and the value is added to the subset if the coin flip shows a head. Since there are $2 \times 2 \times \cdots \times 2 = 2^n$ different sequences of coin tosses, there are 2^n different sets.

We can also think of the coin tosses as determining the values for n different digits in a binary number. The 2^n different sequences generate binary numbers in the range 0 through $2^n - 1$. Given this, we can see a line of attack: count from 0 to $2^n - 1$ and use the binary digits (bits) of the number to determine which of the original values of the Vector are to be included in a subset. Computer scientists work with binary numbers frequently, so there are a number of useful things to remember:

- An int type is represented by 32 bits. A long is represented by 64 bits. For maximum flexibility, it would be useful to use long integers to represent sets of up to 64 elements.
- The arithmetic shift operator << can be used to quickly compute powers of 2. The value 2^i can be computed by shifting a unit bit (1) i places to the left. In Java we write this `1L << i`. This works only for nonnegative, integral powers. The constant `1L` is the value one stored as a 64-bit long value. Using this constant ensures that we are using a 64-bit shift operation resulting in a `long` value instead of a 32-bit operation resulting in an `int` value.
- The "bitwise and" of two numbers can be used to determine the value of a single bit in a number's binary representation. To retrieve bit `i` of a long integer `m` we need only compute `m & (1L << i)`.

Armed with this information, the process of generating subsets is fairly straightforward. One line of attack is the following:

1. Construct a new extension to the `AbstractIterator` class. (By extending the `AbstractIterator` we support both the `Iterator` and `Enumeration` interfaces.) This new class should have a constructor that takes a `Vector` as its sole argument. Subsets of this `Vector` will be returned as the `Iterator` progresses.

   Name this class `SubsetIterator`, and be sure to import `structure5.*` and `java.util.Iterator` at the top of your file. Your `SubsetIterator` should be completely generic. It should know nothing about the values it is iterating over. Thus, the declaration will be

   ```
   public class SubsetIterator<E> extends AbstractIterator<Vector<E>>
   ```

2. Internally, a `long` value is used to represent the current subset. This value increases from 0 (the empty set) to $2^n - 1$ (the entire set of values) as the `Iterator` progresses. Write a `reset` method that resets the subset counter to 0.
3. Write a `hasNext` method that returns `true` if the current value is a reasonable representation of a subset.

4. Write a `get` method that returns a new `Vector` of values that are part of the current subset. If bit `i` of the current counter is 1, element `i` of the original `Vector` is included in the resulting subset `Vector`.

5. Write a `next` method. Remember it returns the current subset before incrementing the counter.

6. For an `Iterator` you would normally have to write a `remove` method. If you extend the `AbstractIterator` class, this method is provided and will do nothing (this is reasonable).

You can now test your new `SubsetIterator` by having it print all the subsets of a `Vector` of values. For example, write a `main` method for your `SubsetIterator` that creates a `Vector` with the `Integers` from 0 through 7, creates a `SubsetIterator` with this `Vector`, and then prints out all subsets returned. Make sure you end up with all 256 different subsets printed.

To solve the two-towers problem, write a main method in a new class `TwoTowers` that inserts the values $\sqrt{1}, \sqrt{2}, ..., \sqrt{n}$ into a `Vector<Double>` object. (To compute the square root of `n`, you can use the `Math.sqrt(n)` method.) A `SubsetIterator` is then used to construct 2^n subsets of these values. The values of each subset are summed, and the sum that comes closest to, but does not exceed, the value $h/2$ is remembered. After all the subsets have been considered, print the best solution.

Thought Questions. Consider the following questions as you complete the lab:

1. What is the best solution to the 15-block problem?

2. How long does it take your program to find the answer to the 20-block problem? You may time programs with the Unix `time` command, as in the following:

   ```
   time java -Xint TwoTowers
   ```

   (The "-Xint" flag turns off some optimizations in the JVM and will give you more reliable results.)

   Based on the time taken to solve the 20-block problem, about how long do you expect it would take to solve the 21-block problem? What is the actual time? How about the 25-block problem? Do these agree with your expectations, given the time complexity of the problem? What about the 40- and 50-block problems? (These will take a *very* long time. Just estimate based on the run times of the smaller problems).

3. This method of exhaustively checking the subsets of blocks will not work for very large problems. Consider, for example, the problem with 50 blocks: there are 2^{50} different subsets. One approach is to repeatedly pick and evaluate random subsets of blocks (stop the computation after 1 second of elapsed time, printing the best subset found). How would you implement `randomSubset`, a new `SubsetIterator` method that returns a random subset?

2.1 Deliverables

Create and submit a tar file `lab7.tar` that includes the following before midnight on the due date:

1. Your `SubsetIterator.java` and `TwoTowers.java` files.

2. A `README` file that includes your answers to the three thought questions.

To create a `tar` file, use the "`tar`" command to archive the full contents of a directory into a single file. For example, the command `tar -cf lab7.tar lab7dir` creates a file called `lab7.tar` containing the full contents of the directory `lab7dir`. You can then run `turnin -c 136 lab7.tar` to submit the tar file.

As in all labs, you will be graded on design, documentation, style, and correctness. Be sure to document your program with appropriate comments, including a general description at the top of the file, a description of each method with pre- and post-conditions where appropriate. Also use comments and descriptive variable names to clarify sections of the code which may not be clear to someone trying to understand it.