

(If you're a little stuck, you might refer back to the `Bag` definition in Chapter 1, which is not identical but is quite similar.)

Talk to me if you're not sure about any of them. There are three or four really important ones, plus a few that would be more optional. Make sure to mark which ones would be `const`.

Once you're pretty confident about your list, write a file `Set.h` that encodes this information in the form of valid C++ method headers. We would like to make our `Sets` able to hold any type of element, as with the classes defined in the book; to make that happen, you just need to precede the class header with

```
template <class Thing>
```

and then use `Thing` as the name of the type the `Set` would hold, whenever you add a value or search for a value or anything like that. (Feel free to use a different name than `Thing`.)

Because our `Set` class is meant to define an interface, we want to mark its methods as “pure virtual” (see p. 45): the implications of this we'll discuss in class, but the mechanics simply involve marking it `virtual` and setting the body to zero. That is, if you had written a method

```
int getSomeValue() const;
```

you would mark it pure virtual by writing

```
virtual int getSomeValue() const = 0;
```

Write a simple test file called `test_VectorSet.u` that, for now, just `#includes` your `Set.h` file and has an empty test suite. Compile that file to confirm that your header has no errors.

Test cases

Now that we have a public interface, we can start planning our test cases. On the next page, first describe a few useful examples (which will become the test fixture). Then, write some sequences of method calls, using those examples, that collectively verify that a `Set` would correctly contain its elements, and does not count or distinguish duplicates.

Starting an implementation

Eventually, we'll write `Set` implementations that run efficiently and mimic the standard implementations, but before we worry about efficiency we have to aim for correctness. Our first implementation will be `VectorSet`, and will use the `vector` built in to C++ to store the data.¹ Its main inefficiency will be that when the user requests to add an element, it will have to check to see that it's not already in the set before adding it.

Edit a file `VectorSet.h` to start working on the class definition. The `VectorSet` will declare itself to be a subclass of `Set` by using the following class header:

```
template <class Thing>
class VectorSet : public Set<Thing>
```

(again, feel free to use a word other than `Thing`). Inside the class, you'll start by making a private instance variable that is a `vector` to hold the data; and then for every pure virtual method in the `Set` definition, you'll

¹Using the terminology of Section CI 4.1–4.2, a `VectorSet` object “has-a” `vector`, but “is-a” `Set`. I also alluded to this in the “big picture” lecture on Monday.

write a stub method in the `VectorSet` definition (for now). Note: because it is a templated class, *all* the code for `VectorSet` will go in the `.h` file.

Testing it

Now that you have the bare bones of an implementation, go ahead and type the test cases you wrote out earlier into the file `test_VectorSet.u` you created earlier. For reasons I'll explain tomorrow, I want you to use pointers here; in your fixture, you'll have lines that look like this:

```
Set<int>* example1 = new VectorSet<int>();
```

You'll probably have more than one, and some of them could be sets of string or whatever, and you should use names more descriptive than “`example1`”. If you have something you want to do to some of them as part of setting up your test fixture, remember that you can create a `setup` block with arbitrary code. And since you're using `new` to create the examples, you should also write a `teardown` block that `deletes` them with statements like

```
delete example1;
```

Once you have your test file typed in, compile it and run it to confirm that everything compiles. If you run your test now, most if not all of the tests will still be failing—they're still just stubs!

Actually writing it

Now go back and start filling in the stub methods. At this point you can compile and test fairly frequently. The more frequently you do so, the easier it will be to find bugs that you inadvertently introduce.

Several of the methods will be quite short, and can simply call an existing method of `vector`! Don't write more than you have to.

Another implementation

Once you've finished `VectorSet`, write a different class called `LazyVectorSet`. From a user perspective, the results it gives should be exactly the same (but

may take more or less time) as a `VectorSet`. The difference is that when the user requests to add an element, it always just adds it (using `push_back`) to the internal `vector`, even if this creates duplicates—making this a cheap operation—but then it has to do a bit more work when it removes something.

Testing that one

The tests for `LazyVectorSet` should be identical to the ones for the other set, right? Copy your existing test file to one called `test_LazyVectorSet.u` and replace all occurrences of `VectorSet` (which should only be at the top of the fixture) with `LazyVectorSet`, and compile the test suite and run it. Debug your `LazyVectorSet` and keep testing it until it passes as well.

Handing in

Hand your code in by 4pm Wednesday, as `lab9` .

RUBRIC

- 1 Present in lab
- 1 Appropriate readme
- Set**
- 1 Method headers
- ¹/₂ pure virtual
- ¹/₂ compiles ♣
- VectorSet**
- 1 Class definition as subclass ♣
- 1 Test suite tests correct behaviour (fail ok) ♣
- 1 Either `add` or `contains` is defined and correct
- 1 `add`, `contains`, and `remove` are correct ♣
- LazyVectorSet**
- 1 Class definition, subclass, `add` is correct ♣
- 1 `remove` is correct ♣

♣ indicates point is only available if the code compiles, with at least a stub for the relevant method(s).

Extra

Produce a table of times and a group of graphs à la Lab 8 to show the efficiency differences between `VectorSet` and `LazyVectorSet`.