

Laboratory Exercises, C++ Programming

General information:

- The course has four compulsory laboratory exercises.
- You shall work in groups of two students. Sign up for the labs at sam.cs.lth.se/Labs.
- The labs require homework. Before each lab session, you must have done as much as you can of the assignments (A1, A2, ...) in the lab, written and tested the programs, and so on. Use the scheduled lab sessions for getting answers to your questions.
- Extra labs are organized only for students who cannot attend a lab because of illness. Notify Sven Gestegård Robertz (Sven.Robertz@cs.lth.se) if you fall ill, *before* the lab.

The labs are about:

1. Basic C++ programming, compiling, linking.
2. Introduction to the standard library.
3. Debugging
4. Strings and streams.

Practical information:

- You will use many half-written “program skeletons” during the lab. You must download the necessary files from the course homepage before you start working on the lab assignments.
- The lab files are available in gzipped tar format. The archives can be unpacked from the command line like this:

```
tar xzf lab1.tar.gz
```

That creates a directory *lab1* in the current directory.

Good sources of information about C++:

- <http://www.cppreference.com>
- <http://www.cplusplus.com>

Programmering i C++, godkända laborationsuppgifter

Skriv ditt namn och din namnteckning nedan:

Namn:

Namnteckning:

Godkänd laboration	Datum	Laborationsledarens namnteckning
1		
2		
3		
4		

1 Basic C++ Programming, Compiling, Linking

Objective: to introduce C++ programming in a Unix environment.

Read:

- Book: basic C++, variables and types including expressions, statements, functions, simple classes, `ifstream`, `ofstream`.
- GCC manual: <http://gcc.gnu.org/onlinedocs/>
- GNU make: <http://www.gnu.org/software/make/manual/>

1 Introduction

Different C++ compilers are available in a Unix environment, for example `g++` from GNU (see <http://gcc.gnu.org/>) and `clang++` from the Clang project (see <http://clang.llvm.org/>). The GNU Compiler Collection, GCC, includes compilers for many languages, the Clang collection only for “C-style” languages. `g++` and `clang++` are mostly compatible and used in the same way (same compiler options, etc.). In the remainder of the lab we mention only `g++`, but everything holds for `clang++` as well.

Actually, `g++` is not a compiler but a “driver” that invokes other programs:

Preprocessor (`cpp`): takes a C++ source file and handles preprocessor directives (`#include` files, `#define` macros, conditional compilation with `#ifdef`).

Compiler: the actual compiler that translates the input file into assembly language.

Assembler (`as`): translates the assembly code into machine code, which is stored in object files.

Linker (`ld`): collects object files into an executable file.

A C++ source code file is recognized by its extension. The two commonly used extensions are `.cc` (recommended by GNU) and `.cpp`. The source files contain definitions. To enable separate compilation, declarations are collected in header files with the extension `.h`. To distinguish C++ headers from C headers other extensions are sometimes used, such as `.hpp` or `.hh`. We will use the extension `.h`.

A C++ program normally consists of many classes that are defined in separate files. It must be possible to compile the files separately. The program source code is usually organized like this (a main program that uses a class `Editor`):

- Define the editor class in a file `editor.h`:

```
#ifndef EDITOR_H // include guard
#define EDITOR_H

#include <string> // include necessary headers here

class Editor {
public:
    /* Creates a text editor containing the text t */
    Editor(const std::string& t) : text(t) {}
    size_type find_left_par(size_type pos) const;

    // ... functions to edit the text (insert and delete characters)
private:
    std::string text;
};
#endif
```

- Define the class member functions in a file *editor.cc*

```
#include "editor.h"
// include other necessary headers

size_type Editor::find_left_par(size_type pos) const { ... }
...
```

- Define the main function in a file *test_editor.cc*:

```
#include "editor.h"
// include other necessary headers

...
int main() {
    Editor ed( ... );
    test_equals( ed.find_left_par(15), 11);
    ...
}
```

The include guard is necessary to prevent multiple definitions of names. Do *not* write function definitions in a header file (except inline functions and function templates).

Compiling and running a program: When compiling from the command line, the g++ command line looks like this:

```
g++ [options] [-o outfile] infile1 [infile2 ...]
```

The *.cc* files are compiled separately. The resulting object files (*.o* files) are linked into an executable file *test_editor*, which is then executed:

```
g++ -std=c++11 -c editor.cc
g++ -std=c++11 -c test_editor.cc
g++ -o test_editor test_editor.o editor.o
./test_editor
```

The *-c* option directs the driver to stop before the linking phase and produce an object file, named as the source file but with the extension *.o* instead of *.cc*.

In order to reduce compilation times, separate compilation as in the above example is usually desired, but gcc also supports giving multiple source files in the same command. Then, that means to both compile and link. The three calls to g++ above can therefore also be written

```
g++ -std=c++11 -o test_editor test_editor.cc editor.cc
```

A1. Write a “Hello, world!” program in a file *hello.cc*, compile and test it.

2 Options and messages

There are more options to the `g++` command than were mentioned in section 1. Your source files must compile correctly using the following command line:

```
g++ -c -O2 -Wall -Wextra -pedantic-errors -Wold-style-cast -std=c++11 file.cc
```

Short explanations (you can read more about these and other options in the `gcc` manual):

<code>-c</code>	just produce object code, do not link
<code>-O2</code>	optimize the object code (perform nearly all supported optimizations)
<code>-Wall</code>	print most warnings
<code>-Wextra</code>	print extra warnings
<code>-pedantic-errors</code>	treat “serious” warnings as errors
<code>-Wold-style-cast</code>	warn for old-style casts, e.g., <code>(int)</code> instead of <code>static_cast<int></code>
<code>-std=c++11</code>	follow the new C++ standard (use <code>-std=c++0x</code> on early versions of <code>g++</code>)
<code>-stdlib=libc++</code>	Clang only — use Clang’s own standard library instead of GNU’s <code>libstdc++</code>

Do not disregard warning messages. Even though the compiler chooses to “only” issue warnings, your program is most likely erroneous or at least questionable. It is strongly recommended that you add the option `-Werror`, which treats warnings as errors.

Some of the warning messages are produced by the optimizer and will therefore not be output if the `-O2` flag is not used. But you must be aware that optimization takes time, and on a slow machine you may wish to remove this flag during development to save compilation time. Some platforms define higher optimization levels, `-O3`, `-O4`, etc. You should not use these optimization levels unless you know very well what their implications are.

It is important that you become used to reading and understanding the GCC error messages. The messages are sometimes long and may be difficult to understand, especially when the errors involve the standard library template classes (or any other complex template classes).

3 Introduction to make

You have to type a lot in order to compile and link C++ programs — the command lines are long, and it is easy to forget an option or two. You also have to remember to recompile all files that depend on a file that you have modified.

There are tools that make it easier to compile and link, “build”, programs. These may be integrated development environments (Eclipse, Visual Studio, ...) or separate command line tools. In Unix, *make* is the most important tool. Make works like this:

- it reads a “makefile” when it is invoked. Usually, the makefile is named *Makefile*¹.
- The makefile contains a description of dependencies between files (which files that must be recompiled/relinked if a file is updated).
- The makefile also contains a description of how to perform the compilation/linking.

As an example, we take the program from section 1. The files *editor.cc* and *test_editor.cc* must be compiled and then linked. Instead of typing the command lines, you just enter the command `make`. Make reads the makefile and executes the necessary commands.

¹ By default `make` uses *Makefile* in the current directory. The command line option `-f` can be used to specify another filename, e.g., `make -f AnotherMakefile`

A minimal makefile, without all the compiler options, looks like this:

```
# The following rule means: "if test_editor is older than test_editor.o
# or editor.o, then link test_editor".
test_editor: test_editor.o editor.o
    g++ -o test_editor test_editor.o editor.o

# Rules to create the object files.
test_editor.o: test_editor.cc editor.h
    g++ -std=c++11 -c test_editor.cc

editor_o: editor.cc editor.h
    g++ -std=c++11 -c editor.cc
```

A rule specifies how a file (the *target*), which is to be generated, depends on other files (the *prerequisites*). The line following the rule contains a shell command, a *recipe*, that generates the target. The recipe is executed if any of the prerequisites are older than the target. It must be preceded by a TAB character, *not* eight spaces.

A2. The file *Makefile* in the *lab1* directory contains the makefile described above. The files *editor.h*, *editor.cc*, and *test_editor.cc* are in the same directory. Experiment:

Run `make`. Run `make` again. Delete the executable program and run `make` again. Change one or more of the source files (it is sufficient to touch them) and see what happens. Run `make test_editor.o`. Run `make notarget`. Read the manual² and try other options.

3.1 Invoking make

Make has many command line options (see the manual for details). A few quite useful ones are

-f filename specify the name of the makefile to use

-B unconditionally make all targets

-C dir change to directory *dir* and run `make` there

-n "dry run": just print the commands that would be executed but do not execute them

4 More Advanced Makefiles

4.1 Implicit Rules

Make has *implicit rules* for many common tasks, for example producing *.o*-files from *.cc*-files. The recipe for this task is:

```
$(CXX) $(CPPFLAGS) $(CXXFLAGS) -c -o $@ $<
```

CXX, *CPPFLAGS*, and *CXXFLAGS* are variables that the user can define. The expression $$(VARIABLE)$ evaluates a variable, returning its value. *CXX* is the name of the C++ compiler, *CPPFLAGS* are the options to the preprocessor, *CXXFLAGS* are the options to the compiler. $@$$ expands to the name of the target, $<$ expands to the first of the prerequisites.

There is also an implicit rule for linking, where the recipe (after some variable expansions) looks like this:

```
$(CC) $(LDFLAGS) $^ $(LOADLIBES) $(LDLIBS) -o $@
```

² See section 9 (*How to run make*) of the manual. The options are summarized in section 9.7.

LDLFLAGS are options to the linker, such as `-Ldirectory`. `LOADLIBES` and `LDLIBS`³ are variables intended to contain libraries, such as `-l1ab1`. `$$` expands to all prerequisites. So this is a good rule, except for one thing: it uses `$(CC)` to link, and `CC` is by default the C compiler `gcc`, not `g++`. But if you change the definition of `CC`, the implicit rule works also for C++:

```
# Define the linker
CC = $(CXX)
```

4.2 Phony Targets

`make` by default creates the first target that it finds in the makefile. By convention, the first target should be named *all*, and it should make all the targets. But suppose that a file *all* is created in the directory that contains the makefile. If that file is newer than the *test_editor* file, a `make` invocation will do nothing but say `make: Nothing to be done for 'all'.`, which is not the desired behavior. The solution is to specify the target *all* as a *phony target*, like this:

```
all: test_editor
.PHONY: all
```

Another common phony target is *clean*. Its purpose is to remove intermediate files, such as object files, and it has no prerequisites. It typically looks like this:

```
.PHONY: clean
clean:
    rm -f *.o test_editor
```

4.3 Generating Prerequisites Automatically

While you're working with a project the prerequisites are often changed. New `#include` directives are added and others are removed. In order for `make` to have correct information about the dependencies, the makefile must be modified accordingly. This is a tedious task, and it is easy to forget a dependency.

The C++ preprocessor can be used to generate prerequisites automatically. The option `-MMD`⁴ makes the preprocessor look at all `#include` directives and produce a file with the extension *.d* which contains the corresponding prerequisite. Suppose the file *test_editor.cc* contains the following `#include` directive:

```
#include "editor.h"
```

The compiler produces a file *test_editor.d* with the following contents:

```
test_editor.o: test_editor.cc editor.h
```

The *.d* files are included in the makefile, so it functions the same way as if we had written the rules ourselves.

³ There doesn't seem to be any difference between `LOADLIBES` and `LDLIBS` — they always appear together and are concatenated. Use `LDLIBS`.

⁴ The option `-MMD` generates prerequisites as a side effect of compilation. If you only want the preprocessing but no actual compilation, `-MM` can be used.

4.4 Putting It All Together

The makefile below can be used as a template for makefiles in many (small) projects. To add a new target you must:

1. add the name of the executable to the definition of PROGS,
2. add a rule which specifies the object files that are necessary to produce the executable.

```
# Define the compiler and the linker. The linker must be defined since
# the implicit rule for linking uses CC as the linker. g++ can be
# changed to clang++.
CXX = g++
CC = $(CXX)

# Generate dependencies in *.d files
DEPFLAGS = -MT $@ -MMD -MP -MF $*.d

# Define preprocessor, compiler, and linker flags. Uncomment the # lines
# if you use clang++ and wish to use libc++ instead of GNU's libstdc++.
# -g is for debugging.
CPPFLAGS = -std=c++11 -I.
CXXFLAGS = -O2 -Wall -Wextra -pedantic-errors -Wold-style-cast
CXXFLAGS += -std=c++11
CXXFLAGS += -g
CXXFLAGS += $(DEPFLAGS)
LDFLAGS = -g
#CPPFLAGS += -stdlib=libc++
#CXXFLAGS += -stdlib=libc++
#LDFLAGS += -stdlib=libc++

# Targets
PROGS = test_editor print_argv

all: $(PROGS)

# Targets rely on implicit rules for compiling and linking
test_editor: test_editor.o editor.o
print_argv: print_argv.o

# Phony targets
.PHONY: all clean

# Standard clean
clean:
    rm -f *.o $(PROGS)

# Include the *.d files
SRC = $(wildcard *.cc)
include $(SRC:.cc=.d)
```

- A3. The makefile with automatic dependencies is in the file *MakefileWithDeps*. Rename this file to *Makefile*, and experiment. The compiler will warn about unused parameters. These warnings will disappear when you implement the member functions. Look at the generated *.d* files. Use the makefile to build your “Hello world!” program.

5 Writing small programs

- A4. Editors for program text usually help with matching of parentheses: when you type a right parenthesis, `)`, `]` or `}` the editor highlights the corresponding left parenthesis.

Example with matching parentheses marked (the dots `...` represent a sequence of characters except parentheses):

```

... (... (... [... ] ...) ...) ... { ... } ...
           |---|           |---|
           |-----|
           |-----|

```

Implement the function `find_left_par` in the class `Editor`, test with `editortest.cc`.

- A5. Implement two functions for encoding and decoding:

```

/* For any character c, encode(c) is a character different from c */
unsigned char encode(unsigned char c);

/* For any character c, decode(encode(c)) == c */
unsigned char decode(unsigned char c);

```

Use a simple method for coding and decoding. Test your encoding and decoding routines with `test_coding.cc`.

Then write a program, `encode`, that reads a text file⁵, encodes it, and writes the encoded text to another file. The program can ask for a filename as in the following execution

```

./encode
enter filename.
myfile

```

and write the encoded contents to `myfile.enc`.

Alternatively, you can give the file name on the command line, like this:

```

./encode myfile

```

Command-line parameters are passed to `main` in an array of C-strings. The prototype to use is `int main(int argc, const char** argv)`. See `print_argv.cc` for an example of how to use command line arguments.

Write another program, `decode`, that reads an encoded file, decodes it, and writes the decoded text to another file `FILENAME.dec`. Add rules to the makefile for building the programs.

Test your programs and check that a file that is first encoded and then decoded is identical to the original. Use the Unix `diff` command.

Note: the programs will work also for files that are UTF-8 encoded. In UTF-8 characters outside the “ASCII range” are encoded in two bytes, and the `encode` and `decode` functions will be called twice for each such character.

6 Finding Errors

With the GNU debugger, `gdb`, you can control a running program (step through the program, set breakpoints, inspect variable values, etc.). Debug information is inserted into the executable program when you compile and link with the option `-g`. Preferably you should also turn off optimization (no `-O2` option). From `g++` version 4.8 there is a new option `-Og`, which turns on all optimizations that do not conflict with debugging.

⁵ Note that you cannot use `while (infile >> ch)` to read all characters in `infile`, since `>>` skips whitespace (why?). Use `infile.get(ch)` instead. Output with `outfile << ch` should be ok, but `outfile.put(ch)` looks more symmetric.

A program is executed under control of gdb like this:

```
gdb ./program
```

Some useful commands:

<code>help [command]</code>	Get help about gdb commands.
<code>run [args...]</code>	Run the program (with arguments).
<code>start [args...]</code>	Set a temporary breakpoint at <code>main</code> and run the program
<code>continue</code>	Continue execution.
<code>next</code>	Step to the next line <i>over</i> function calls.
<code>step</code>	Step to the next line <i>into</i> function calls.
<code>finish</code>	Continue until just after the current function returns.
<code>where</code>	Print the call stack.
<code>list [nbr]</code>	List 10 lines around the current line or around line <code>nbr</code> (the following lines if repeated).
<code>break func</code>	Set a breakpoint on the first line of a function <code>func</code> .
<code>break nbr</code>	Set a breakpoint at line <code>nbr</code> in the current file.
<code>print expr</code>	Print the value of the expression <code>expr</code> .
<code>watch var</code>	Set a watchpoint, i.e., watch all modifications of a variable. Can be very slow but can be the best solution to find some bugs.

`gdb` also has a mode (called “text user interface”) which splits the window into one window for the source code and one for the command prompt. This mode is toggled with the key sequence `ctrl-x a` (i.e., first type `ctrl-x` and then type an `a`). In the code window, the arrow keys scroll the source code, in the command window the arrow keys move in the command history. The key sequence `ctrl-x o` jumps to the other windows.

A6. Run the test programs under control of `gdb`, try the commands.

6.1 Memory-Related Errors

This section is for information only, and the `valgrind` tool is not available under Windows. If you use a Mac or a Linux box, you are strongly encouraged to use `valgrind`.

In Java, many errors are caught by the compiler (use of uninitialized variables) or by the runtime system (addressing outside array bounds, dereferencing null pointers, etc.). In C++, errors of this kind are not caught, instead they result in erroneous results or faults during program execution. Furthermore, you get no information about where in the program the error occurred. Since deallocation of dynamic memory in C++ is manual, you also have a whole new class of errors (dangling pointers, double delete, memory leaks).

`Valgrind` is a tool (available under Linux and Mac OS X) that helps you find memory-related errors at the precise locations at which they occur. It does this by emulating an x86 processor and supplying each data bit with information about the usage of the bit. This results in slower program execution, but this is more than compensated for by the reduced time spent in searching for bugs.

`Valgrind` is easy to use. Compile and link as usual, then execute like this:

```
valgrind ./program
```

When an error occurs, you get an error message and a stack trace (and a lot of other information). At the end of execution, `valgrind` prints a “leak summary” which indicates the amount of dynamic memory that hasn’t been properly freed.

7 Object Code Libraries

A lot of software is shipped in the form of libraries, e.g. class packages. In order to use a library, a developer does not need the source code, only the object files and the headers. Object file libraries may contain thousands of files and cannot reasonably be shipped as separate files. Instead, the files are collected into library files that are directly usable by the linker.

7.1 Static Libraries

The simplest kind of library is a *static library*. The linker treats the object files in a static library in the same way as other object files, i.e., all code is linked into the executable files. In Unix, a static library is an *archive file*, `lib<name>.a`. In addition to the object files, an archive contains an index of the symbols that are defined in the object files.

A collection of object files `f1.o`, `f2.o`, `f3.o`, ..., are collected into a library `libfoo.a` using the `ar` command:

```
ar crv libfoo.a f1.o f2.o f3.o ...
```

(Some Unix versions require that you also create the symbol table with `ranlib libfoo.a`.) In order to link a program `main.o` with the object files `obj1.o`, `obj2.o` and with the object files in the library `libfoo.a`, you use the following command line:

```
g++ -o main main.o obj1.o obj2.o -L. -lfoo
```

The linker searches for libraries in certain system directories. The `-L.` option makes the linker search also in the current directory.⁶ The library name (N.B.! the *name*, *without* the prefix `lib` and the suffix `.a`) is given after `-l`.

For debugging, it can sometimes be interesting to look at the symbols defined in an object file or library. For this, the utility `nm` can be used. The symbols in an object file `foo.o` is listed with the command

```
nm foo.o
```

By default, `nm` lists all symbols in a file. To restrict it to just defined or undefined symbols, the options `-defined-only` and `-undefined-only` can be used. If you run it, you see that `c++` function names are *mangled* to avoid name clashes for overloaded functions and member functions. With GNU `nm`, the option `-demangle` makes the names more readable. If that option is not available, the utility `c++filt` can be used to demangle symbol names, as in the command line

```
nm -undefined-only main.o | c++filt
```

A7. (optional) Collect the object files `editor.o` and `coding.o` in a library `liblab1.a`. Change the makefile so the programs (`test_editor`, `encode`, `decode`) are linked with the library. The `-L` option belongs in `LDFLAGS`, the `-l` option in `LDLIBS`.

Note that this does not tell make how to create `liblab1.a`. For that, add a rule

```
liblab1.a: coding.o editor.o
    ar crv liblab1.a coding.o editor.o
```

Note that you cannot easily write rules to make the programs depend on the lib, so you must first make `liblab1.a` and then `make`. See section 8 for how CMake handles this.

Please note that putting code into a library is usually a way to separate common, reusable, and stable, parts (the libraries) from more specific, and often more actively developed parts (the main program). For the remainder of the labs in this course, it is probably overkill to make parts of the code into libraries even if it is generic.

⁶ You may have several `-L` and `-l` options on a command line. Example, where the current directory and the directory `/usr/local/mylib` are searched for the libraries `libfoo1.a` and `libfoo2.a`:

```
g++ -o main main.o obj1.o obj2.o -L. -L/usr/local/mylib -lfoo1 -lfoo2
```

7.2 Shared Libraries

Since most programs use large amounts of code from libraries, executable files can grow very large. Instead of linking library code into each executable that needs it the code can be loaded at runtime. The object files should then be in *shared libraries*. When linking programs with shared libraries, the files from the library are not actually linked into the executable. Instead a “pointer” is established from the program to the library.

In Unix shared library files are named *lib<name>.so[.x.y.z]* (*.so* for shared objects, *.x.y.z* is an optional version number). The linker uses the environment variable `LD_LIBRARY_PATH` as the search path for shared libraries. In Microsoft Windows shared libraries are known as DLL files (dynamically loadable libraries).

- A8.** (Advanced, optional) Create a shared library with the object files *editor.o* and *coding.o*. Link the executables using the shared library. Make sure they run correctly. Compare the sizes of the dynamically linked executables to the statically linked (there will not be a big difference, since the library files are small).

Use the command `ldd` (list dynamic dependencies) to inspect the linkage of your programs. Shared libraries are created by the linker, not the `ar` archiver. Use the `gcc` and `ld` manpages (and, if needed, other manpages) to explain the following sequence of operations:

```
g++ -fPIC -std=c++11 -c *.cc
g++ -shared -Wl,-soname,liblab1.so.1 -o liblab1.so.1.0 list.o coding.o
ln -s liblab1.so.1.0 liblab1.so.1
ln -s liblab1.so.1 liblab1.so
```

You then link with `-L. -llab1` as before. The linker merely checks that all referenced symbols are in the shared library. Before you execute the program, you must define `LD_LIBRARY_PATH` so it includes the current directory. You do this with the following command (on the command line):

```
export LD_LIBRARY_PATH=.:$LD_LIBRARY_PATH
```

8 CMake, a system for generating build scripts

Make is a standard tool for building programs, but it is quite low-level and for larger projects, a more high-level build system is commonly used. One example is CMake, which is used to specify how to build a project in an operating system and in a compiler-independent manner. CMake then generates the required Makefiles (or project files for one of the supported IDEs).

The directory *cmake-example* contains a small project, consisting of an example library, a configuration file, and a main program. This gives a brief overview of how CMake works.

With *cmake*, you usually build the project in a directory separate from the source, typically named *build*. This has the advantages that you can easily make several separate builds (e.g. testing and production) simply by doing them in separate build directories. It also means that the generated files are kept separate from the source code, so that removing them is done by simply removing the entire build directory. The steps to create the build files and then build the project are, assuming you are standing in the project root directory, in this case *cmake-example*:

```
mkdir build && cd build
cmake ..
make
```

- A9. Study the *cmake-example* project and then write *CMakeLists.txt* files for *lab1* and use that to build the programs. The files that went into *liblab1.a* in the previous section (7.1) should be moved to a subdirectory. (Copy your files to a separate directory for this task, in order to keep your makefile-only solution.) Verify that both the library and the main program is rebuilt if you make changes to the library source.

2 Introduction to the Standard Library

Objective: to solve a moderately large problem using C++. Some parts of the standard library that haven't yet been introduced in the course will be used. They are introduced in section 3.

1 Spelling Correction

Most word processors can check the spelling of a document and suggest corrections to misspelled words. Often, a dictionary is used — words that aren't in the dictionary are considered to be misspelled. The suggested corrections are the words in the dictionary that are “similar” to the misspelled word.

Your task is to write a class `Dictionary` which can be used as in the following example:

```
void check_word(const string& word, const Dictionary& dict)
{
    if (dict.contains(word)) {
        cout << "Correct." << endl;
    } else {
        vector<string> suggestions = dict.get_suggestions(word);
        if (suggestions.empty()) {
            cout << "Wrong, no suggestions." << endl;
        } else {
            cout << "Wrong. Suggestions:" << endl;
            for (const auto& w : suggestions) {
                cout << "    " << w << endl;
            }
        }
    }
}

int main() {
    Dictionary dict;
    string word;
    while (cin >> word) {
        transform(word.begin(), word.end(), word.begin(), ::tolower);
        check_word(word, dict);
    }
}
```

Examples:

```
expertise
Correct.
seperate
Wrong. Suggestions:
    separate
    desperate
    federate
    generate
    imperate
```

Notes:

- The function `contains` (section 2.2) must be efficient (fast).
- In `get_suggestions` you can spend time on finding good suggestions for corrections.
- It can be advantageous to “preprocess” the file which contains the dictionary (section 2.1).
- It is not certain that the data structures which you shall use are optimal (or even necessary), but you shall solve the assignments as they are given. You are encouraged to improve the program, but do that as a separate project.

The following shall be done in `get_suggestions`:

1. Search the dictionary and find candidates for corrections (section 2.3). To begin with, the words in the dictionary which have approximately the same number of letters (plus/minus one letter) as the misspelled word should be considered. Of these candidates, the words which contain at least half of the “trigrams” of the misspelled word should be kept. A trigram is three adjacent letters — for example, the word `summer` contains the trigrams `sum` `umm` `mme` `mer`.
2. Sort the candidate list so the “best” candidates are first in the list (section 2.4). The sort key is the cost to change the misspelled word to one of the candidate words.
3. Keep the first 5 candidates in the list (section 2.5).

Expressed in program code:

```
vector<string> Dictionary::get_suggestions(const string& word) const {
    vector<string> suggestions;
    add_trigram_suggestions(suggestions, word);
    rank_suggestions(suggestions, word);
    trim_suggestions(suggestions);
    return suggestions;
}
```

2 Assignments

2.1 Preprocess the Dictionary

- A1. The file `/usr/share/dict/words` contains a large number of words (one word per line). The file is UTF-8 encoded; ignore this and treat all characters as 8-bit. Write a program which reads the file and creates a new file `words.txt` in the current directory. Each line in the file shall contain a word, the number of trigrams in the word, and the trigrams.⁷ The trigrams shall be sorted in alphabetical order; upper case letters shall be changed to lower case. Example:

```
...
hand 2 and han
handbag 5 and bag dba han ndb
handbook 6 and boo dbo han ndb ook
...
```

Copy the *Makefile* from the *lab1* directory, modify it to build the program, build, test.

2.2 Determine If a Word is Correct

- A2. Implement the constructor and the function `contains` in `Dictionary`. The preprocessed list of words is in the file `words.txt`. The words shall be stored in an `unordered_set<string>`. Wait with the trigrams until assignment A4.

Modify the makefile (the main program shown in section 1 is in *spell.cc*), build, test.

2.3 Use Trigrams to Find Candidates

- A3. The words together with their trigrams must be stored in the dictionary. Each word shall be stored in an object of the following class:

⁷ Note that there are short words with zero trigrams.


```

class Word {
public:
    /* Creates a word w with the sorted trigrams t */
    Word(const std::string& w, const std::vector<std::string>& t);

    /* Returns the word */
    std::string get_word() const;

    /* Returns how many of the trigrams in t that are present
       in this word's trigram vector */
    unsigned int get_matches(const std::vector<std::string>& t) const;
};

```

Implement this class. The trigram vector given to the constructor is sorted in alphabetical order (see assignment A1). The function `get_matches` counts how many of the trigrams in the parameter vector that are present in the word's trigram vector.⁸ You may assume that the trigrams in the parameter vector also are sorted in alphabetical order. Use this fact to write an efficient implementation of `get_matches`.

- A4. The class `Dictionary` shall have a member variable that contains all words with their trigrams. It must be possible to quickly find words which have approximately the same length as the misspelled word. Therefore, the words shall be stored in the following array:

```

vector<Word> words[25]; // words[i] = the words with i letters,
                       // ignore words longer than 25 letters

```

Modify the `Dictionary` constructor so the `Word` objects are created and stored in `words`, implement the function `add_trigram_suggestions`. Use a constant instead of the number 25.

2.4 Sort the Candidate List

After `add_trigram_suggestions` the suggestion list can contain a large number of candidate words. Some of the candidates are "better" than others. The list shall be sorted so the best candidates appear first. The sorting condition shall be the "edit distance" (also called "Levenshtein distance") from the misspelled word to the candidate word.

The cost $d(i, j)$ to change the i first characters in a word p to the j first characters in another word q can be computed with the following formula (i and j start from 1):

$$\begin{aligned}
 d(i, 0) &= i \\
 d(0, j) &= j \\
 d(i, j) &= \text{minimum of } \begin{cases} \text{if } p_i = q_j \text{ then } d(i-1, j-1) \text{ else } d(i-1, j-1) + 1, \\ d(i-1, j) + 1, \\ d(i, j-1) + 1. \end{cases}
 \end{aligned}$$

The minimum computation considers the cost for replacing a character, inserting a character and deleting a character. The cost to change p to q , that is the edit distance, is $d(p.length, q.length)$.

- A5. Implement the function `rank_suggestions`. Do *not* write a recursive function, it would be very inefficient. Instead, let d be a matrix (with $d(i, 0) = i$ and $d(0, j) = j$) and compute the elements in row order (dynamic programming). Declare d with the type `int [26] [26]` to avoid problems with a dynamically allocated matrix.

⁸ You don't have to consider multiple occurrences of the same trigram.

2.5 Keep the Best Candidates

A6. Implement the function `trim_suggestions`.

3 More Information About the Assignments

- In the main program in *spell.cc*, the call to `transform` applies the function `tolower` to all the characters in a string (between `begin()` and `end()`), and stores the function result in the same place. `tolower` converts a character from upper case to lower case. The scope operator `::` is necessary to get the right version of the overloaded `tolower` function.
- To sort a vector `v`, call `std::sort(v.begin(), v.end())` (include `<algorithm>`).
- The standard library class `unordered_set` is in header `<unordered_set>`. An element is inserted in a set with the function `insert(element)`. The function `count(element)` returns the number of occurrences of an element (0 or 1).
- Here's one way to sort the suggested candidates in edit distance order (another way is to use a `map`):
 - Define a vector with elements that are pairs with the first component an `int`, the second a `string`: `vector<pair<int, string>`.
 - Compute the edit distance for each candidate word, insert the distance and the word in the vector: `push_back(make_pair(dist, word))`.
 - Sort the vector (pairs have an operator`<` that first compares the first component, so the elements will be sorted according to increasing edit distance).
 - For a pair `p`, `p.first` is the first component, `p.second` the second component.
- Read more about computing edit distance on the web. You may also consider using the Damerau–Levenshtein distance.
- A vector can be resized to size `n` with `resize(n)`.

3 Debugging

Objective: to practice debugging a C++ program

1 A (broken) system for access control

A company has created a system for controlling access to its facilities. To open a door, an employee inserts a card in a card-reader. The system looks up the card number in a table. If the card number is found, the name of the employee is presented on a display and the door is unlocked. Otherwise the door remains locked.

The class `UserTable` is used to check who a certain card number belongs to. The access control system includes the following code:

```
UserTable t;
...
int cardNbr = sensor.getCardNbr();

User u = t.find(cardNbr);           // Find the owner of the card

if (u == UserTable::user_not_found) {
    display.showText("** access denied **");
} else {
    display.showText("Welcome, " + u.getName());
    door.unlock();
}
```

The interesting line is marked with a comment. Given a card number (an integer) the system looks up a user. The variables `sensor`, `display`, and `door` refer to objects outside the scope of this task. Can you understand the general operation of the above code?

The class `UserTable` is unfortunately broken. Your task is to find and correct the mistakes. Present your solution according to the table at the end of this section, together with a working test program.

2 Tasks

1. Study the source code in the files `User.h`, `User.cc`, `UserTable.h` och `UserTable.cc`. There is also a text file, `users.txt`, containing all users and their card numbers. Look at the file and see how it is organized. In the class `UserTable` the user database is read from the text file and stored in an array.

2. We will now debug the class `UserTable`. To test the class, it is impractical to try a lot of cards in the card reader. A better option is to create a simple test-program that creates a `UserTable` object and calls the function `find` in the same way the card-reader does. Then you can simply make up and test different card numbers yourself.

Create such a test program (including a `main` function) that

- creates a `UserTable` object,
- calls the function `find(int)` to find the user with card number 21330,
- calls the function `find(int)` to find the user with card number 21331 and
- calls the function `find(std::string)` to find the user named "Jens Holmgren".

What is the result of the program? What did you expect? (Compare with the file `users.txt`.)

3. The class `UserTable` has a member function for printing the contents of the table. Extend your test program with a call to this function and see what is printed.

Study the constructor for the class `UserTable`. Here, the user database is read from a file and stored in an array. Either the reading from the file is not working, or the read lines are not stored correctly in the array.

Add some lines to your test program so that a new (made up) user with card number 1234 is added, and then print the contents of the table. Run your test program. How many users are now in the table? How many did you expect?

4. Find and fix the mistake that you identified in the previous task. To get further clues you can ask yourself the following questions:

- The function `add` inserts a `User` object at the right position in the array. What positions are the objects put in? That is, what value does the variable `pos` get? (You can either use debug printouts, or run it in the debugger and use breakpoints. Remember to build without optimization when using the debugger.)
- Does the function `getNbrUsers` return the correct result? (Test.)

When you have corrected the mistake, check if searching for a card number works. Verify that your test program calls `find(int)` with a card number, and that the result is the same card number and the correct user name. Compare with the corresponding line in `users.txt`. Always compile with warnings enabled (the option `-Wall` to `g++`).

5. It has been reported that `find(std::string name)` does not work. The function is implemented with *linear search*: starting at the first position, go through the array until a person with the name `name` is found. Study the implementation of the function to find the bug.

When the function is corrected, the search for Jens Holmgren should work.

6. Verify that searching works. There is function `testFindNumber()` that can be used to test the member function `UserTable::find(int)`. Let your test program call that test function and verify that all users (card numbers) now can be found.

7. Check the test function. It has been reported that `testFindNumber()` performs the tests correctly, but that the tested `UserTable` object is left in an unusable state afterwards. Investigate what happens when `testFindNumber()` is called.

Note. The names in the table have been generated from the 100 most common swedish family names and male and female names⁹ (excluding names containing åäö). They do not refer to actual persons.

⁹ Source: SCB, <http://www.scb.se/namnstatistik/>.

4 Strings and Streams. Testing.

Objective: to practice using the standard library string and stream classes.

Read:

- Book: strings, streams, function templates, exceptions.

1 Unit testing

When writing code, testing is important and a common methodology in modern *agile* software development methodologies is *test driven design* or “test first”. One major benefit of writing tests before writing the code is that this helps with understanding the problem and structuring the code, starting from the desired result and working ones way towards the solution. That is an application of the principle of *programming by intention* (or “wishful thinking”) and helps designing functions with suitable parameters and return types.

In the strict formulation, you are only allowed to write new code if there is a unit test that doesn’t pass. So to add new functionality you first write a unit test and make the test call the desired new function. In doing so, you specify both what the arguments to the function will be, and what type and value the function should return for the given arguments. This also includes defining any new types that you need to express the desired functionality. At this stage, your test will not compile, so now you add an empty function (with `return 0;`, `return false;` or what is suitable). Run the tests and make sure that they fail — if not, either the functionality is already supported or the test case is wrong. Then, implement the function to make the test case pass.

Another good testing principle is to write *unit tests*, that tests “the smallest testable unit” (e.g., functions, classes) in addition to large-scale tests that verifies the function of the entire system. Unit tests are valuable during development as they make it easy to check if additions to the system have affected the behaviour of the old (apparently unrelated) functionality. If the unit tests ran successfully before a modification to the code, they should also run successfully after it.

It is also often a good idea to write test programs that do not require manual user input or manually checking the output. For two small examples of such test programs, study the examples given in lab1: `test_coding` and `test_editor`.

A0. Read through the assignments of this lab (A1 – A5), and write test programs for each assignment. For instance, for the first assignment (A1), create a file (or string) where you have manually removed the HTML tags from the given HTML file and write a program that calls your tag removal function and compares its result with the manually created file (string). Start with smaller unit tests, e.g. for testing the removal of HTML tags and the replacement of special characters. For those, write small test cases that test just one thing (“unit of functionality”).

For the second assignment (A2), you can translate the example with the numbers 0–35 into code, testing that your functions returns both the correct strings (“CCPPC...”) and the corresponding prime number sequence).

For the third assignment (A3), there is a main program (that requires user input) given. You can base your test program on that. A good option here is to give the streams to use as parameters instead of hard-coding `std::cin` and `std::cout`. Then, one can use `std::stringstream` to automate both the input and checking the results. (See section 3.2 of this lab for info on stringstream.)

For the last assignment (A5), remember to also test that your function throws exceptions correctly.

2 Class string

2.1 Introduction

In C, a string is a null-terminated array of characters. This representation is the cause of many errors: overwriting array bounds, trying to access arrays through uninitialized or incorrect pointers, and leaving dangling pointers after an array has been deallocated. The `<cstring>` library contains operations on C-style strings, such as copying and comparing strings.

C++ strings hide the physical representation of the sequence of characters. The exact implementation of the `string` class is not defined by the C++ standard.

The `string` identifier is not actually a class, but a type alias for a specialized template:

```
using string = std::basic_string<char>;
```

This means that `string` is a string containing characters of type `char`. There are other string specializations for strings containing “wide characters”. We will ignore all “internationalization” issues and assume that all characters fit in one byte.

`string::size_type` is a type used for indexing in a string. `string::npos` (“no position”) is a value indicating a position beyond the end of the string; it is returned by functions that search for characters when the characters aren’t found.

2.2 Operations on Strings

The following class specification shows most of the operations on strings:

```
class string {
public:
    /** construction **/
    string(); // creates an empty string
    string(const string& s); // creates a copy, also has move constructor
    string(const char* cs); // creates a string with the characters from cs
    string(size_type n, char ch); // creates a string with n copies of ch

    /** information **/
    size_type size(); // number of characters

    /** character access **/
    const char& operator[](size_type pos) const;
    char& operator[](size_type pos);

    /** substrings */
    string substr(size_type start, size_type n = npos); // the substring starting
                                                    // at position start containing n characters

    /** inserting, replacing, and removing **/
    string& insert(size_type pos, const string& s); // inserts s at position pos
    string& append(const string& s); // appends s at the end
    string& replace(size_type start, size_type n, const string& s); // replaces n
                                                    // characters starting at pos with s
    void erase(size_type start = 0, size_type n = npos); // removes n
                                                    // characters starting at pos

    /** assignment and concatenation **/
    string& operator=(const string& s); // also move assignment
    string& operator=(const char* cs);
    string& operator=(char ch);
    string& operator+=(const string& s); // also const char* and char

    /** access to C-style string representation **/
    const char* c_str();
    /** finding things (see below) **/
}
```


- Note that there is no constructor `string(char)`. Use `string(1, char)` instead.
- The subscript functions (operator[]) do not check for a valid index. There are similar `at()` functions that do check, and that throw `out_of_range` if the index is not valid.
- The `substr()` member function takes a starting position as its first argument and the number of characters as the second argument. This is different from the `substring()` method in `java.lang.String`, where the second argument is the end position of the substring.
- There are overloads, for C-style strings or characters, of most of the functions.
- Strings have iterators like the standard library collections (e.g., `std::vector`).
- There is a bewildering variety of member functions for finding strings, C-style strings or characters. They all return `npos` if the search fails. The functions have the following signature (the `string` parameter may also be a C-style string or a character):

```
size_type FIND_VARIANT(const string& s, size_type pos = 0) const;
```

`s` is the string to search for, `pos` is the starting position. (The default value for `pos` is `npos`, not 0, in the functions that search backwards).

The “find variants” are `find` (find a string, forwards), `rfind` (find a string, backwards), `find_first_of` and `find_last_of` (find one of the characters in a string, forwards or backwards), `find_first_not_of` and `find_last_not_of` (find a character that is not one of the characters in a string, forwards or backwards). For example:

```
void f() {
    string s = "accdcde";
    auto i1 = s.find("cd");           // i1 = 2 (s[2]=='c' && s[3]=='d')
    auto i2 = s.rfind("cd");         // i2 = 4 (s[4]=='c' && s[5]=='d')
    auto i3 = s.find_first_of("cd"); // i3 = 1 (s[1]=='c')
    auto i4 = s.find_last_of("cd");  // i4 = 5 (s[5]=='d')
    auto i5 = s.find_first_not_of("cd"); // i5 = 0 (s[0]!='c' && s[0]!='d')
    auto i6 = s.find_last_not_of("cd"); // i6 = 6 (s[6]!='c' && s[6]!='d')
}
```

There are global overloaded operator functions for concatenation (operator+) and for comparison (operator==, operator<, etc.). They all have the expected meaning. Note that you cannot use + to concatenate a string with a number, only with another string, C-style string or character (this is unlike Java). In the new standard, there are functions that convert strings to numbers and vice versa: `stod("123.45") => double`, `to_string(123) => "123"`.

- A1.** Write a class that reads a file and removes HTML tags and translates HTML-encoded special characters. The class should be used like this:

```
int main() {
    TagRemover tr(std::cin); // read from cin
    tr.print(std::cout);    // print on cout
}
```

- All tags should be removed from the output. A tag starts with a < and ends with a >.
- You can assume that there are no nested tags.
- Tags may start on one line and end on another line.
- Line separators should be kept in the output.
- You don't have to handle all special characters, only `<`, `>`, ` `, `&`; (corresponding to < > space &).
- Make sure that you use the standard library. Manual iteration and copying character by character is not a good solution.
- Assignments like this should be a good fit for regular expressions. Study and use the C++ regex library if you're interested¹⁰.

Copy the makefile from one of the previous labs, modify it, build and test.

¹⁰ However, be warned that regular expressions grow quite complicated very quickly, and a more straight forward solution using `find` and `replace` or `erase`, as outlined above, is often preferable for both correctness and readability.

- A2. The Sieve of Eratosthenes is an ancient method for finding all prime numbers less than some fixed number M . It starts by enumerating all numbers in the interval $[0, M]$ and assuming they are all primes. The first two numbers, 0 and 1 are marked, as they are not primes. The algorithm then starts with the number 2, marks all subsequent multiples of 2 as composites, finds the next prime, marks all multiples, ... When the initial sequence is exhausted, the numbers not marked as composites are the primes in $[0, M]$.

In this assignment you shall use a string for the enumeration. Initialize a string of appropriate length to PPPPP...PPP. The characters at positions that are not prime numbers should be changed to C.

Example with the numbers 0–35:


	1	2	3	
	0	1	2	3
Initial:	012345678901234567890123456789012345			
Find 2, mark 4,6,8,...:	CCPPPPPPPPPPPPPPPPPPPPPPPPPPPPPPPPPP			
Find 3, mark 6,9,12,...:	CCPPPCPCPCPCPCPCPCPCPCPCPCPCPCPCPCPC			
Find 5, mark 10,15,20,...:	CCPPPCPCPCPCPCPCPCPCPCPCPCPCPCPCPCPC			
Find 7, mark 14,21,28,35:	CCPPPCPCPCPCPCPCPCPCPCPCPCPCPCPCPCPC			
Find 11, mark 22,33:	CCPPPCPCPCPCPCPCPCPCPCPCPCPCPCPCPCPC			
...				

- Write a program that prints the prime numbers between 1 and 200 and also the largest prime that is less than 100,000.
- Do not expose the internal string representation in your interface. For instance, if you want to return a sequence of primes, use `std::vector<int>` and not a string.

3 The iostream Library

3.1 Input/Output of User-Defined Objects

In addition to the stream classes for input or output there are `iostream`'s that allow both reading and writing. The stream classes are organized in the following (simplified) hierarchy:



The classes `ios_base` and `ios` contain, among other things, information about the stream state. There are, for example, functions `bool good()` (the state is ok) and `bool eof()` (end-of-file has been reached). There is also a conversion operator `operator bool()` that returns true if the state is good, and a `bool operator!()` that returns true if the state is not good. We have used these operators with input files, writing for example `while (infile >> ch)` and `if (!infile)`.

To do formatted stream input and output of objects of user-defined classes, `operator>>` and `operator<<` must be overloaded.

- A3.** The files `date.h`, `date.cc`, and `date_test.cc` describe a simple date class. Implement the class and add operators for input and output of dates (`operator>>` and `operator<<`). Dates should be output in the form 2015-01-10. The input operator should accept dates in the same format. (You may consider dates written like 2015-1-10 and 2015 -001 - 10 as legal, if you wish.)

The input operator should set the stream state appropriately, for example `is.setstate (ios_base::failbit)` when a format error is encountered. Write your code so that the right hand operand of `operator>>` is not changed if the conversion fails.

3.2 String Streams

The string stream classes (`istringstream` and `ostringstream`) function as their “file” counterparts (`ifstream` and `ofstream`). The only difference is that characters are read from/written to a string instead of a file. In the following assignments you will use string streams to convert objects to and from a string representation (in the new standard, this can be performed with functions like `to_string` and `stod`, but only for numbers).

- A4.** In Java, the class `Object` defines a method `toString()` that is supposed to produce a “readable representation” of an object. This method can be overridden in subclasses.

Write a template function `toString` for the same purpose. Also write a test program. Example:

```
double d = 1.234;
Date today;
std::string sd = toString(d);
std::string st = toString(today);
```

You may assume that the argument object can be output with `<<`.

- A5.** Type casting in C++ can be performed with, for example, the `static_cast` operator. Casting from a string to a numeric value is not supported, since this involves executing code that converts a sequence of characters to a number.

Write a function template `string_cast` that can be used to cast a string to an object of another type. Examples of usage:

```
try {
    int i = string_cast<int>("123");
    double d = string_cast<double>("12.34");
    Date date = string_cast<Date>("2015-01-10");
} catch (std::invalid_argument& e) {
    cout << "Error: " << e.what() << endl;
}
```

You may assume that the argument object can be read from a stream with `>>`. The function should throw `std::invalid_argument` (defined in header `<stdexcept>`) if the string could not be converted.