

Introduction to C++ Complementary Library

CCL

Luxury for Developers

2nd Edition

Covers CCL version 0.2 Beta

Written by Daniel T. McGinnis

Published on December 26th, 2023

Copyright © 2022, 2023 Daniel T. McGinnis

Table of Contents

Chapter 1: Getting Started

- Library Organization

- Stability Promises

- Consuming CCL

 - Using CCL without CMake (in Visual Studio, Xcode, Code::Blocks, etc.)

 - Using CCL in a CMake Project

 - Building the Examples

- Configuration Constants

 - CMP_CONFIG_LTO_REQUESTED

 - CMP_CONFIG_LTO_ENABLED

 - CMP_CONFIG_HEADER_ONLY

 - CMP_CONFIG_DEFAULT_BUFFER_CAPACITY

 - CMP_CONFIG_DEFAULT_STDOUT_BUFFER_CAPACITY

 - CMP_CONFIG_DEFAULT_STDERR_BUFFER_CAPACITY

 - CMP_CONFIG_DEFAULT_STDIN_BUFFER_CAPACITY

 - CMP_CONFIG_DESKTOP_GUI_PACKAGE_EXCLUDED

 - CMP_CONFIG_IO_PACKAGE_EXCLUDED

Chapter 2: Unicode

- Reviewing ASCII

- Say Hello to Unicode

- Unicode Terminology

- Unicode in the Standard Library

- Unicode in CCL

 - Printing Unicode Strings to Standard Output

 - Reading Unicode Strings from Standard Input

 - Iterating Over Unicode Strings

 - Converting Unicode Strings

 - Reading Unicode Strings from Files

 - Writing Unicode Strings to Files

- Endianness

- Handling Endianness in CCL

- Source File Encoding

Chapter 3: I/O

- Introduction

 - Transfer Resources

 - Transfer Streams

 - How They Fit Together

- Files

- Transfer Streams

 - Data Input Streams

 - Data Output Streams

 - Text Input Streams

 - Text Output Streams

- Container Transfer Resources

- Wrapping Up

Chapter 4: Desktop GUI

- Supported Platforms

- An Example GUI Application

 - Overview

 - #include <cmp/desktop_gui/main.hpp>

 - #include <cmp/desktop_gui/desktop_gui_application.hpp>

- Reacting to Input

- Wrapping Up

Chapter 1: Getting Started

Welcome to CCL. CCL is a native C++ library for developing high-performance cross-platform applications. This chapter will get you started with CCL and discuss its essential terminology and practices. While this book doesn't explore every nook and cranny of CCL, reading this chapter alone will give you a firm grasp of the fundamentals of CCL. This chapter gives an overview of the library, and later chapters explore how to apply CCL to specific kinds of problems, like Unicode and file I/O.

Chapter Table of Contents

- Library Organization

- Stability Promises

- Consuming CCL

 - Using CCL without CMake (in Visual Studio, Xcode, Code::Blocks, etc.)

 - Using CCL in a CMake Project

 - Building the Examples

- Configuration Constants

 - CMP_CONFIG_LTO_REQUESTED

 - CMP_CONFIG_LTO_ENABLED

 - CMP_CONFIG_HEADER_ONLY

 - CMP_CONFIG_DEFAULT_BUFFER_CAPACITY

 - CMP_CONFIG_DEFAULT_STDOUT_BUFFER_CAPACITY

 - CMP_CONFIG_DEFAULT_STDERR_BUFFER_CAPACITY

 - CMP_CONFIG_DEFAULT_STDIN_BUFFER_CAPACITY

 - CMP_CONFIG_DESKTOP_GUI_PACKAGE_EXCLUDED

 - CMP_CONFIG_IO_PACKAGE_EXCLUDED

Library Organization

CCL is a **library**. It consists of several **packages**. Different software projects use different terminology to distinguish between their various parts. Boost, for example, is a set of C++ libraries, and each library is useful for a particular kind of task. For example, Boost.Container or Boost.Asio. In this case, the "parts" of Boost are called "libraries". Qt is a framework and its parts are called "modules". For example, Qt Core, Qt Network, etc. CCL is a library, and its parts are called "packages", and each package is useful for a particular kind of task. These packages are listed below.

- CCL Core
- CCL Desktop GUI
- CCL IO
- CCL Unicode

Each package is discussed in detail in the following chapters. For now, know that CCL Core is the only package that every other package depends on, CCL Desktop GUI provides basic window and input capabilities on desktop (like spinning up a GUI window and reacting to key presses), CCL IO handles tasks related to I/O (like reading and writing files), and CCL Unicode handles tasks related to Unicode (like iterating Unicode strings or converting between different encoding forms).

It is important to note that some packages can be implemented purely in portable C++, whereas others require the operating system's native API to get the job done. The supported operating systems are Windows, macOS, GNU/Linux and FreeBSD. The term "supported operating systems" here means operating systems that CCL is specifically designed for, compiled on, and tested on before every release. OS-dependent packages can be disabled with compile-time flags if your particular operating system is not supported or if the functionality in the package is not needed. Currently, the only OS-dependent packages are CCL Desktop GUI and CCL IO.

Stability Promises

CCL does not guarantee backwards compatibility. This means that if you have an existing application that uses CCL and you want to upgrade to a newer version of CCL when it comes out, then you may have to change your source code but you will have to at least recompile your application. This makes upgrading to newer versions of CCL less practical but it makes it easier to improve CCL as time goes forward. Source-breaking changes will always be highlighted in the release notes, along with suggestions on how to do the upgrade.

Consuming CCL

When you use CCL, you can use all packages or exclude OS-dependent packages, but in either case, CCL is used as one library. You cannot use CCL packages individually like you can with Boost libraries.

CCL can be used header-only or compiled as a static or dynamic library. The advantage of compiling CCL and linking it with your code is that you cut down on compile times because you don't recompile CCL code every time you compile your own code. The advantage of using CCL header-only is that you get the guarantee that you literally only compile and ship the code that is needed by your application and nothing else, so if you use one class from CCL, your binary will only include the code for that one class and whatever else that one class depends on, but you won't necessarily pay for the cost of the entire library.

The following sections explain how to consume CCL based on what approach you take to build your code. We look at how to integrate CCL into a CMake-based project and how to use CCL in popular IDEs like Visual Studio or Xcode. Then we look at how to build and run the examples.

Using CCL without CMake (in Visual Studio, Xcode, Code::Blocks, etc.)

CMake is a great build tool, but there are a lot of software projects out there that are tied to a particular IDE, and there are several reasons why this is common, so this section goes over how to consume CCL in IDEs without using CMake.

After downloading CCL, copy the CCL directory to your location of choosing and then go to your IDE's settings dialog and add the following directories to your include paths:

- <CCL root>/api/include
- <CCL root>/api/src

The reason you have to add the **src** directory as an include path is because every CCL header knows whether it's being used header-only, and if it is then it includes the corresponding source file. This setup allows the interface (which is in the header files) to be physically separate from the implementation (which is in the source files) and still allow users of the library to choose whether to compile it or use it header-only.

Once your include paths have been set up, you have to choose whether to use CCL header-only or compile it as a static or dynamic library.

To use it header-only, make sure that the preprocessor constant **CMP_CONFIG_HEADER_ONLY** is defined as **true** before including any CCL header files in your code. This can be done more easily by using your IDE's settings dialog to create a global preprocessor constant. Then just include CCL headers in your code and build your project the way you normally would in the IDE you're using.

Note that when the **CMP_CONFIG_HEADER_ONLY** constant is **true**, you are unable to use CCL Desktop GUI. This is because on macOS, CCL doesn't use C++, instead it uses Objective-C++ in order to access macOS' native system API. To prevent that language choice from leaking out of CCL, and to enable users of CCL to stick to C++ on all platforms, CCL disables header-only support in CCL Desktop GUI, because that way all Objective-C++ code is hidden in the source files whereas the header files are strictly C++. This means that if you wish to use CCL header-only then you must go without CCL Desktop GUI, and if you wish to use CCL Desktop GUI then you must use CCL as a static or dynamic library and link your application with it.

To use CCL as a static or dynamic library, don't define **CMP_CONFIG_HEADER_ONLY** at all, or just define it as **false**, and then actually compile CCL with CMake and link the compiled binary with your project using your IDE's settings dialog. Then just include CCL headers in your code and build your project normally through your IDE.

To compile CCL, simply build it like any other CMake library:

```
cd ccl
mkdir build
cd build
cmake ..
cmake --build . --config Release
ctest
```

This will produce a dynamic library. If you wish to use CCL as a static library then add the **-DCMP_BUILD_STATIC=true** flag to the CMake configuration command. Like this:

```
cmake .. -DCMP_BUILD_STATIC=true
```

There are several other configuration constants besides **CMP_CONFIG_HEADER_ONLY** that you can define to customize CCL, which we will discuss later in this chapter. These are simply preprocessor constants that begin with **CMP_CONFIG_**. When you use CCL as a header-only library, and wish to specify a value for a particular configuration constant, all you have to do is define it as a global preprocessor constant with the value you want, which you can do through your IDE's settings dialog. However, when you use CCL as a static or dynamic library, make sure that whatever configuration constants you define in your IDE's settings dialog are spelled the same in the CMake configuration command. For example, if you define **CMP_CONFIG_IO_PACKAGE_EXCLUDED** to be **true** in your IDE's settings dialog, and wish to use CCL as a compiled library that you link with your project, then make sure to include **-DCMP_CONFIG_IO_PACKAGE_EXCLUDED=true** in the configuration step of CMake. Ultimately that would look like this:

```
cmake .. -DCMP_CONFIG_IO_PACKAGE_EXCLUDED=true
```

This of course also works with static builds, which for completeness is illustrated below:

```
cmake .. -DCMP_BUILD_STATIC=true -DCMP_CONFIG_IO_PACKAGE_EXCLUDED=true
```

Of course, feel free to specify a generator with the **-G** flag, if you wish. CCL is known to work with the latest versions of Visual C++, GCC and Clang.

The unit tests are executed before every version of CCL is released, but it is a good idea to run them after you build the library just to make sure that everything is working as expected. This is done with the **ctest** command. After that, CCL is built and ready to be linked with your project.

Using CCL in a CMake Project

This sections goes over how to consume CCL within a CMake project.

CMake is very flexible in that it lets you define the directory structure of your project. Let's look at an example application that comes with CCL. The root directory of CCL has a subdirectory called **examples** where you can find the **unicode_stdio_example** directory, which contains a simple CMake-based application that uses CCL to print Unicode text to standard output and read Unicode text from standard input. The CMakeLists.txt file for that application is shown below:

```
# Copyright (C) 2022 Daniel T. McGinnis
# SPDX-License-Identifier: BSL-1.0

cmake_minimum_required(VERSION 3.17)
project(cc1 CXX)

if (NOT DEFINED CMP_DEV)
    add_subdirectory(../../api cmp_build)
endif ()

add_executable(
    unicode_stdio_example
    src/main.cpp
)
set_target_properties(
    unicode_stdio_example PROPERTIES
    LINKER_LANGUAGE CXX
    CXX_STANDARD 20
    CXX_STANDARD_REQUIRED TRUE
    CXX_EXTENSIONS OFF
)
target_link_libraries(
    unicode_stdio_example
    cmp
)
```

Let's go over this code piece by piece. The instructions here will assume you are starting a brand new CMake project, and that you are taking inspiration from the code above to write your own CMakeLists.txt file. Of course, if you *copy* the code above, be sure to not carry over the copyright notice as that would imply that I am the owner of your CMakeLists.txt file, which of course is not correct. Also, the second comment at the top specifies that this file is licensed under the Boost Software License 1.0, feel free to not carry over that comment as well.

```
cmake_minimum_required(VERSION 3.17)
```

This line of code simply states that the minimum version of CMake required to build this application is version 3.17. Requiring version 3.17 is not a law, your own CMake project may require a different minimum version.

```
project(cc1 CXX)
```

This line of code says that **cc1** is a C++ project. Change **cc1** to the name of your own CMake project.

```
if (NOT DEFINED CMP_DEV)
    add_subdirectory(../../api cmp_build)
endif ()
```

This code tests whether the **CMP_DEV** variable is defined, which I define privately during development to provide for a better experience for me as I develop CCL. You don't need this if statement, the important part for you is the **add_subdirectory** command. Use this **add_subdirectory** line to pull in the **cmp** library target defined in the CMakeLists.txt file of the **api** directory, which your application's executable target will later link with. Change this line of code to match where you put CCL. For example, if you put CCL in the same directory as your CMakeLists.txt file, then this line should read **add_subdirectory(ccl/api cmp_build)**.

```
add_executable(  
    unicode_stdio_example  
    src/main.cpp  
)
```

This piece of code adds an executable target called **unicode_stdio_example**. Change this to the name you desire to use for your own application's executable target. Additionally, we are listing all the source files of this target, which is just the **main.cpp** file in the **src** directory. Again, change this to your heart's content. Add all the other source files your project uses to this list, if any. I do recommend that you put your source files in the **src** directory but that is optional as well.

```
set_target_properties(  
    unicode_stdio_example PROPERTIES  
    LINKER_LANGUAGE CXX  
    CXX_STANDARD 20  
    CXX_STANDARD_REQUIRED TRUE  
    CXX_EXTENSIONS OFF  
)
```

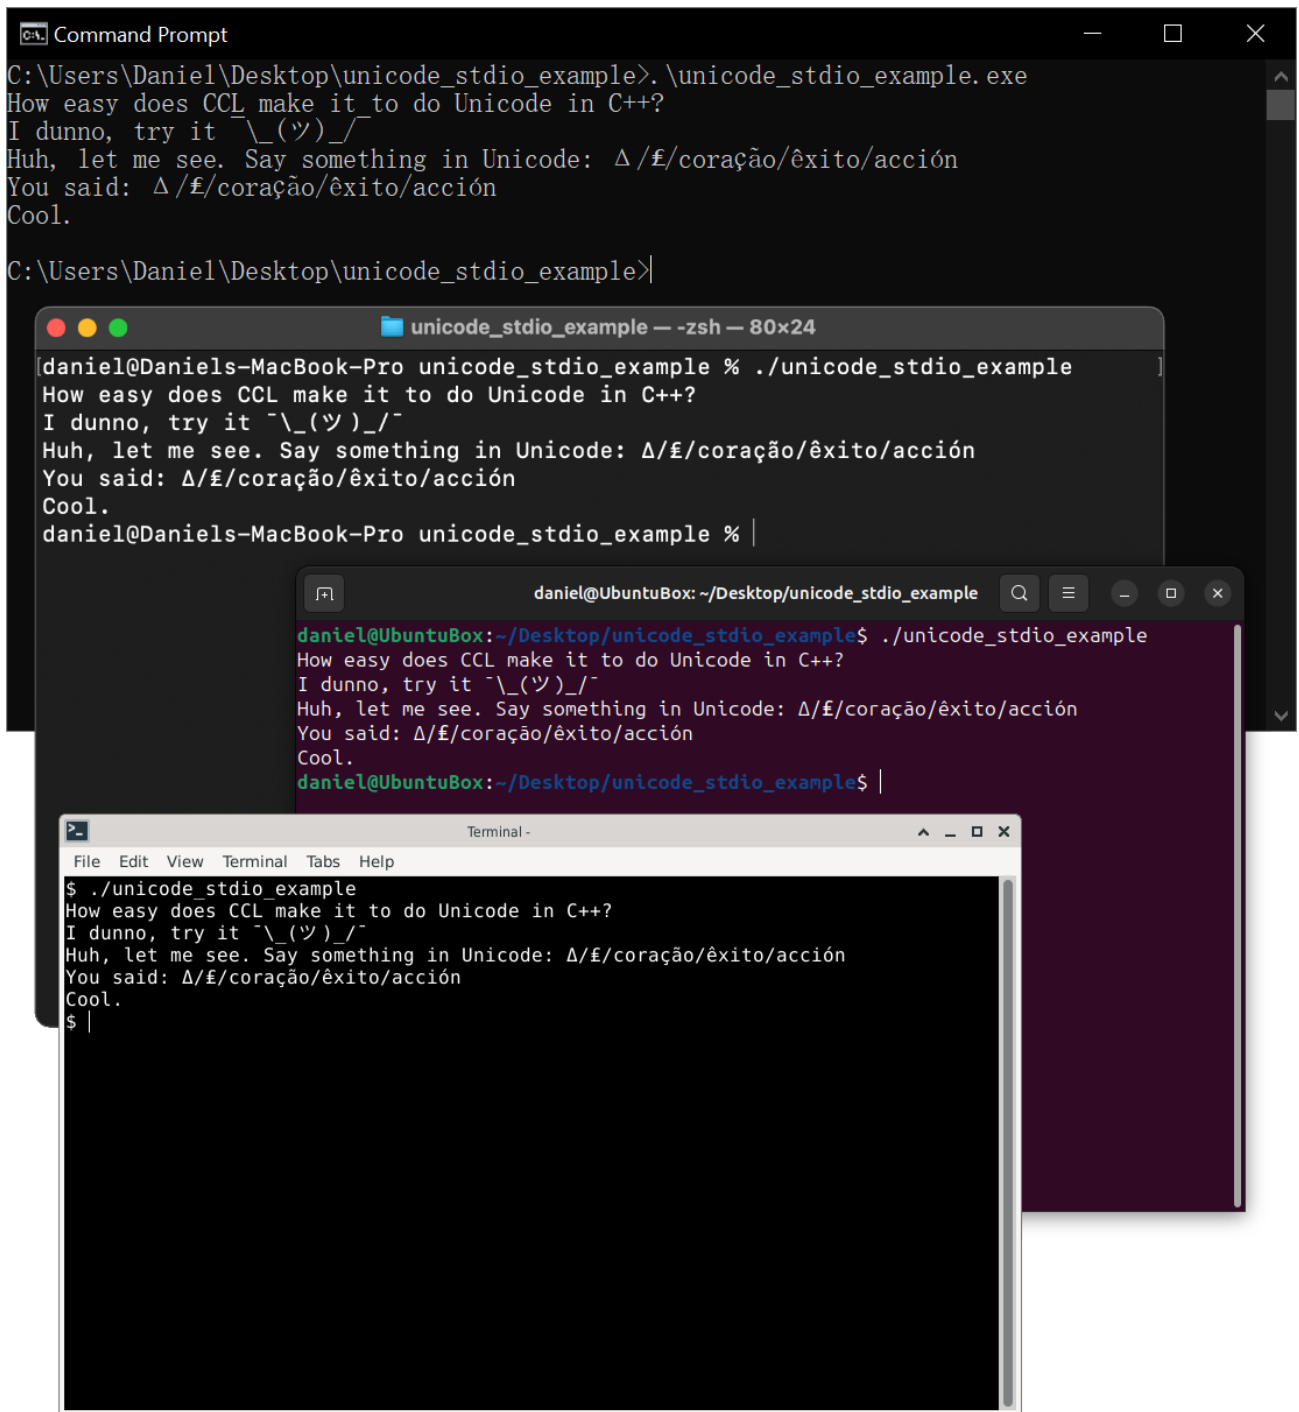
This piece of code confines the **unicode_stdio_example** target to standard C++20. Change **unicode_stdio_example** to the name of your actual executable target, specified in the **add_executable** command shown previously. Feel free to change this piece of code if you want to use compiler extensions, you don't necessarily have to conform to the C++ standard if you don't want to, but it is important to understand that CCL **requires** that you use C++20 or later. C++17 and earlier standards are not supported.

```
target_link_libraries(  
    unicode_stdio_example  
    cmp  
)
```

This piece of code links the **unicode_stdio_example** executable target with the **cmp** library target. Change **unicode_stdio_example** to your own target's name. This will make CCL available to your application. Now just include CCL headers in your code, build it with CMake, and you're set. Here's the contents of the **main.cpp** file for this example:

```
// Copyright (C) 2022 Daniel T. McGinnis  
// SPDX-License-Identifier: BSL-1.0  
  
#include <iostream>  
#include <cmp/io/uio.hpp>  
  
int  
main ()  
{  
    std::cout << "How easy does CCL make it to do Unicode in C++?\n";  
    cmp::uout << u8"I dunno, try it `\\_(ツ)_/`\n";  
    cmp::uout << u"Huh, let me see. Say something in Unicode: ";  
    std::u8string input;  
    cmp::uin >> input;  
    cmp::uout << U"You said: " << input << '\n';  
    cmp::uout << "Cool.\n";  
  
    return 0;  
} // function -----
```

Here is what the program above outputs on Windows, macOS, Ubuntu and FreeBSD:



Building the Examples

The examples don't get built by default, they have to be built separately, and individually. You basically just follow the same steps for building CCL, except you point the CMake configuration command to the desired example and you skip the **ctest** command. To build **unicode_stdio_example**, just run:

```
cd ccl
mkdir build
cd build
cmake ../examples/unicode_stdio_example
cmake --build . --config Release
```

You'll get an executable file, just make sure that the CCL library file is in the same directory as the executable file. For instance, if you build an example with Visual C++, then place **cmp.dll** in the same directory as **unicode_stdio_example.exe**. Then just run the executable file and check out the example.

Building **desktop_gui_example** does require first making sure your system is properly set up for developing GUI applications. On Windows and macOS, there are no steps you need to take to use CCL Desktop GUI. On Linux and FreeBSD, however, you do need to make sure that pkg-config and the GTK 4 development package are installed. If you're on Linux, run the following commands:

```
sudo apt-get install pkg-config
sudo apt-get install libgtk-4-dev
```

If you're on FreeBSD, run these commands:

```
su
pkg install pkgconf
pkg install gtk4-4.10.4
pkg install noto-2.0
```

The specific version of GTK 4 that is available when you configure your FreeBSD development machine may be different from what is shown above. Version 4.10.4 is the latest version of GTK that is available through pkg at the time of writing, but you can find the exact version that is available to you by running:

```
pkg search gtk
```

Just look for the package called **gtk4-4.x.y** where **x** and **y** may be newer than the **10** and **4** shown earlier.

The **noto** package is needed on FreeBSD to equip your terminal with the necessary Unicode glyphs. Without it, your Unicode text will show up garbled. This is a requirement for users of your applications as well, although it goes without saying that you need a glyph for a character if you wish to display it.

That's all the setup you need to use CCL Desktop GUI. To build **desktop_gui_example**, replace **unicode_stdio_example** in the instructions shown earlier with **desktop_gui_example**. This is what you end up with:

```
cd ccl
mkdir build
cd build
cmake ../examples/desktop_gui_example
cmake --build . --config Release
```

Again, make sure the CCL library file is in the same directory as the example's executable file before you run it. On Windows it's going to be **cmp.dll** or **libcmp.dll**, on Linux it'll be called **libcmp.so**, and so on.

Configuration Constants

CCL can be customized, or **configured**, with preprocessor constants called **configuration constants**. The full list, along with their default values, is shown in Table 1-1.

Name	Default Value
CMP_CONFIG_LTO_REQUESTED	true
CMP_CONFIG_LTO_ENABLED	false
CMP_CONFIG_HEADER_ONLY	false
CMP_CONFIG_DEFAULT_BUFFER_CAPACITY	1024
CMP_CONFIG_DEFAULT_STDOUT_BUFFER_CAPACITY	CMP_CONFIG_DEFAULT_BUFFER_CAPACITY
CMP_CONFIG_DEFAULT_STDERR_BUFFER_CAPACITY	CMP_CONFIG_DEFAULT_BUFFER_CAPACITY
CMP_CONFIG_DEFAULT_STDIN_BUFFER_CAPACITY	CMP_CONFIG_DEFAULT_BUFFER_CAPACITY
CMP_CONFIG_DESKTOP_GUI_PACKAGE_EXCLUDED	false
CMP_CONFIG_IO_PACKAGE_EXCLUDED	false

Table 1-1: CCL's configuration constants.

Let's take a closer look at each of these.

CMP_CONFIG_LTO_REQUESTED

This constant determines whether link-time optimization (also called interprocedural optimization, or IPO) is desired. If IPO is requested and supported then it will be enabled. Otherwise, it will be disabled.

CMP_CONFIG_LTO_ENABLED

This constant determines whether link-time optimization is enabled. You do not set this constant yourself. Instead, you set the **CMP_CONFIG_LTO_REQUESTED** constant, and then CMake will define **CMP_CONFIG_LTO_ENABLED** for you. If IPO is requested and supported, this constant will be set to **true**. Otherwise, it will be set to **false**. These LTO constants do nothing if you use CCL header-only.

CMP_CONFIG_HEADER_ONLY

This constant determines whether the library will be compiled or used header-only. The only package that can't be used header-only is CCL Desktop GUI, so if you use that package you can't use CCL header-only. Regardless of whether you compile CCL or use it header-only, make sure that CCL's **api/include** and **api/src** directories are both in your compiler's include paths. Then, if this constant is set to false, compile the library, include the headers in your code, and link with the compiled binary (statically or dynamically, both are supported). Or, if this constant is set to **true**, just include the headers in your code.

CMP_CONFIG_DEFAULT_BUFFER_CAPACITY

This constant determines the capacity of the I/O buffer of transfer resources in general.

CMP_CONFIG_DEFAULT_STDOUT_BUFFER_CAPACITY

This constant determines the capacity of the I/O buffer used when writing to standard output.

CMP_CONFIG_DEFAULT_STDERR_BUFFER_CAPACITY

This constant determines the capacity of the I/O buffer used when writing to standard error.

CMP_CONFIG_DEFAULT_STDIN_BUFFER_CAPACITY

This constant determines the capacity of the I/O buffer used when reading from standard input.

CMP_CONFIG_DESKTOP_GUI_PACKAGE_EXCLUDED

This constant determines whether the CCL Desktop GUI package should be excluded. A reason for doing this is if you are using a platform that this package doesn't support. Another reason is if you simply don't need the functionality provided by this package.

CMP_CONFIG_IO_PACKAGE_EXCLUDED

This constant determines whether the CCL IO package should be excluded. A reason for doing this is if you are using a platform that this package doesn't support. Another reason is if you simply don't need the functionality provided by this package.

Chapter 2: Unicode

Processing text in C++ can be difficult. C++ has extensive support for ASCII, allowing you to make programs that speak English, but if you want to make your software ready for use by people all around the world then you need support for Unicode. CCL provides some support for Unicode so you can stay in C++ and interoperate well with the world at large. However, to use CCL's Unicode facilities well you must first understand Unicode. This chapter introduces you to Unicode and how to use it effectively in C++ with CCL. CCL does not wrap existing Unicode libraries, it is written from scratch, giving you all the performance and flexibility enabled by C++20 in an elegant, modern API.

Chapter Table of Contents

[Reviewing ASCII](#)

[Say Hello to Unicode](#)

[Unicode Terminology](#)

[Unicode in the Standard Library](#)

[Unicode in CCL](#)

[Printing Unicode Strings to Standard Output](#)

[Reading Unicode Strings from Standard Input](#)

[Iterating Over Unicode Strings](#)

[Converting Unicode Strings](#)

[Reading Unicode Strings from Files](#)

[Writing Unicode Strings to Files](#)

[Endianness](#)

[Handling Endianness in CCL](#)

[Source File Encoding](#)

Reviewing ASCII

A **character set** is a system that makes it possible to represent characters in a computer. Character sets are usually implemented as a set of mappings between numbers and characters, where each number represents a unique character. The simplest of these is ASCII (American Standard Code for Information Interchange). The way ASCII works is it defines 128 different integers, where each one is associated with a unique character. So for example, the decimal number 65 represents an upper-case 'A'. ASCII encodes all letters of the English alphabet, all decimal digits, the basic punctuation marks, as well as some other common characters. Because a character encoded in ASCII requires only 7 bits of space, it easily fits in a single octet.

A common misunderstanding is that every byte is exactly 8 bits long. While 8 bits is certainly the most common size of a byte by far, some systems do have bytes of other sizes. The term **octet** refers to a number of exactly 8 bits. To simplify this discussion, this chapter will use the term "byte" from here on to refer to an 8-bit number. So we can say things like "each ASCII character takes up one byte".

The limitation of ASCII is that it doesn't represent all the characters used throughout the world, it only represents the characters that are common in English, hence the "American" in American Standard Code for Information Interchange. One solution to this problem would be to create a new character set where we simply list more number-to-character mappings, allowing more characters to be represented. This would increase the size of each character. You might say that it would be easy to just make each character use up all 8 bits so that you can represent 256 different characters and each character would still just take up one byte. That sounds interesting but 256 different numbers still doesn't let you encode all characters used in the world. You may suggest 16-bit characters. Certainly 65,536 different numbers is enough to encode a ton of characters, right? Well, it turns out that 16 bits is still not enough to encode all characters used in the world. Plus, it seems like a lot of work to specify all those characters! There is a solution, and the hard work has already been done. Enter Unicode.

Say Hello to Unicode

Unicode is a character set that assigns a unique number to pretty much every character used in the world. Every ASCII character is included in the Unicode character set, but Unicode adds thousands more characters to cater not only to speakers of English, but also to speakers of Spanish, French, German, Russian, Chinese, Japanese, Korean, and every other language in use today. Unicode also encodes historic characters that are no longer commonly used.

Unicode currently encodes more than a hundred thousand characters. You may assume then that every character in Unicode takes up more than 16 bits because 16 bits can only encode up to 65,536 different characters. That's not exactly how it works. Unicode has three different ways of expressing characters: UTF-8, UTF-16 and UTF-32. In UTF-32, every character is expressed as a single 32-bit number. In UTF-16, every character is expressed as one or two 16-bit numbers. In UTF-8, every character is expressed as one, two, three or four 8-bit numbers.

The advantage of UTF-32 is that every character really is a single value, allowing true random access (not as useful as it might seem, more on this later). The advantage of UTF-8 is that it is backwards-compatible with ASCII, that is, one byte containing an ASCII character is a valid Unicode character in UTF-8, without any processing needed. So going back to the example given earlier, ASCII associates the number 65 with the character 'A'. In UTF-32, that character would take up 32 bits but it would still just be the value 65. In UTF-8, the character would take up 8 bits and the value would also be 65, the most significant bit would just be zero. Example:

```
'A' = 65 (in both ASCII and Unicode)
65 in ASCII:                               1000001 (7 bits)
65 in UTF-8:                                01000001 (8 bits)
65 in UTF-16:                               00000000 01000001 (16 bits)
65 in UTF-32: 00000000 00000000 00000000 01000001 (32 bits)
```

Figure 2-1: How ASCII and Unicode represent the character 'A'.

Let's look at another example. The character '∞' (the symbol for infinity) maps to the number 8,734. Because this number doesn't fit in 7 bits, UTF-8 expresses it in more than one byte, but because it still fits comfortably

in 16 bits and 32 bits, UTF-16 represents it as a single 16-bit number and UTF-32 represents it as a single 32-bit number. This is what that looks like:

```
'∞' = 8,734 in Unicode
8,734 in UTF-8:           11100010 10001000 10011110
8,734 in UTF-16:         00100010 00011110
8,734 in UTF-32: 00000000 00000000 00100010 00011110
```

Figure 2-2: How Unicode represents the character '∞'.

Notice how in this case, the UTF-8 representation takes up more space (3 bytes) than the UTF-16 representation (2 bytes). Contrast this with how UTF-8 is more compact when representing the character 'A' (1 byte) than UTF-16 (2 bytes). UTF-8 isn't always more compact than UTF-16. UTF-8 is more compact in text files that have lots of ASCII characters, and UTF-16 is more compact in text files that have lots of characters from East-Asian languages.

Unicode Terminology

Unicode defines some terms that are important to understand in order to use Unicode effectively, which we will look at now. In Unicode, a **code point** is the number assigned to a character. For example, 65 is the number assigned to the character 'A', so 65 is a code point. 8,734 is the number assigned to the character '∞', so 8,734 is a code point. The code point 8,734 is expressed differently in UTF-8 than it is in UTF-16, but that doesn't change the fact that the number 8,734 is a single code point.

A **character** is a (possibly invisible) unit of text. A **grapheme cluster** is a visible unit of text as defined by a written language. The difference between a character and a grapheme cluster is that a grapheme cluster necessarily represents a graphic symbol that speakers of a language understand as indivisible, whereas a character may also represent other kinds of information necessary to represent text in a computer, for example line breaks, accents that are applied to letters, and format control codes. This distinction is important, so let's look at an example to clarify. 'ä' is a lower-case 'a' with a diaeresis applied. This is a single graphic symbol that represents an atomic unit of written text, just like the letter 'a' without a diaeresis on top. However, 'ä' can be encoded as one or two characters. One possible representation is the code point 228 ('ä'), another possible representation is the two code points 97 ('a') and 168 ('"). Unicode says that the diaeresis mark on its own is a character but not a grapheme cluster because people don't think of a single diaeresis mark as a full-blown unit of text, it is always placed on top of a letter like 'a'. Suddenly the term "character" doesn't sound as useful as it used to. So if you have a UTF-32 string object in C++, which is an array of characters, then a particular element in that string may be a single grapheme cluster as understood by speakers of a particular language, or it may represent a portion of a grapheme cluster. This is why the random access capability you get with UTF-32 isn't all that useful, since a single character in Unicode may not be what a user of a computer program would think of as an atomic unit of text. Since no Unicode encoding form provides random access to grapheme clusters, the variable length nature of UTF-8 is not a drawback because you still have to traverse a string object linearly if you want to manipulate grapheme clusters.

An **encoding form** is a method of expressing code points in memory and in storage. Unicode defines three encoding forms: UTF-8, UTF-16 and UTF-32. A **code unit** is the smallest value an encoding form works with, and it can either represent a code point or a part of a code point. The UTF-8 encoding form works with 8-bit code units, the UTF-16 encoding form works with 16-bit code units, and the UTF-32 encoding form works with 32-bit code units. UTF-8 can take up to four 8-bit code units to express one code point, UTF-16 can take up to two 16-bit code units to express one code point, and UTF-32 always uses one 32-bit code unit to express one code point.

Unicode in the Standard Library

The C++ standard library defines classes that represent Unicode strings. **std::u8string** represents a string encoded in UTF-8, **std::u16string** represents a string encoded in UTF-16, and **std::u32string** represents a string encoded in UTF-32. These classes store code units, that is, each element of a **std::u8string** is a **char8_t**, each element of a **std::u16string** is a **char16_t**, and each element of a **std::u32string** is a **char32_t**.

You can initialize strings of these types with string literals. The C++ language allows you to express string literals encoded in Unicode. You can initialize a Unicode string object with a Unicode string literal that contains fancy Unicode content, like this:

```
std::u8string utf8_string{u8"ㄟ(ツ)_/"};
std::u16string utf16_string{u"ㄟ(ツ)_/"};
std::u32string utf32_string{U"ㄟ(ツ)_/"};
```

That's about as far as the standard library goes. You can't read or write Unicode text files, you can't output Unicode strings to standard output or read them from standard input, there aren't any Unicode-aware algorithms, etc. This is where CCL comes in.

Unicode in CCL

CCL doesn't implement the entire Unicode standard but it does provide a lot of useful Unicode functionality. With CCL, you can print Unicode strings to standard output. You can read Unicode strings from standard input. You can iterate over the code points of Unicode strings of any encoding form. You can convert strings from any Unicode encoding form to any other Unicode encoding form. You can read and write Unicode text files. CCL also handles both little endian and big endian order. CCL's Unicode capabilities, while useful, are still quite limited if you consider the scope of the Unicode standard, but more features are on the way.

With that introduction out of the way, let's now go over how to do all these things with CCL. As we discuss CCL's Unicode capabilities, keep in mind that CCL does not define its own Unicode string types, it simply manipulates standard string types, aka **std::u8string**, **std::u16string** and **std::u32string** and corresponding string view types.

Printing Unicode Strings to Standard Output

You can write any Unicode-encoded string object to standard output and standard error with the **cmp::uout** and **cmp::uerr** objects. Here's what that looks like:

```
#include <cmp/io/uio.hpp>
...
std::u16string utf16_string{u"UTF-16 string here ㄟ(ツ)_/.\\n"};
cmp::uout << utf16_string << u8"UTF-8 literal here ㄟ(ツ)_/.\\n";
```

The only requirement is that your terminal or shell have a font that supports the Unicode characters you try to print out. If it does, you'll get gorgeous Unicode text. If it doesn't, you'll get weird-looking garbage characters.

Reading Unicode Strings from Standard Input

You can read Unicode-encoded strings from standard input with the **cmp::uin** object. Here's an example of how to read 3 Unicode strings from standard input:

```
#include <cmp/io/uio.hpp>
...
std::u8string first_string;
std::u16string second_string;
std::u32string third_string;

cmp::uin >> first_string >> second_string >> third_string;
// 3 Unicode strings were read from standard input.
```

Beware of the fact that, on Windows, reading Unicode text from standard input with **cmp::uin** will fail if standard input is redirected. This is due to a known bug in Windows, and until Microsoft fixes it, there is nothing CCL (or any other software) can do about it. On all other operating systems, **cmp::uin** works

normally, regardless of whether standard input has been redirected. If you need to be able to read from standard input on Windows regardless of whether it has been redirected then you cannot read Unicode text, you must fall back to non-Unicode mechanisms like `std::cin` or `scanf` or the `ReadFile` Windows API function.

Iterating Over Unicode Strings

You can iterate over a Unicode string object by code point, regardless of the encoding form of that string. You do this by constructing a `cmp::by_code_point` with the string you want to traverse, and then you just iterate over that object and you'll get a full code point at each iteration of the loop. Here's what that looks like:

```
#include <cmp/unicode/by_code_point.hpp>
...
std::u8string utf8_string{get_a_string_somewhat()};
for (char32_t current_code_point : cmp::by_code_point{utf8_string}) {
    process_a_code_point_somewhat(current_code_point);
}
// Now all code points stored in utf8_string have been processed.
```

Converting Unicode Strings

You can convert strings from one encoding form to another. To do that, call the functions `cmp::to_u8string`, `cmp::to_u16string` and `cmp::to_u32string` and pass in a standard Unicode string object (i.e. a `std::u8string`, a `std::u16string` or a `std::u32string`). Here's an example:

```
#include <cmp/unicode/algorithms.hpp>
...
std::u16string utf16_string{get_a_string_somewhat()};
std::u32string utf32_string{cmp::to_u32string(utf16_string)};
process_a_string_somewhat(utf32_string);
// utf16_string has been converted to UTF-32
// and the converted string has been processed.
```

Reading Unicode Strings from Files

You can read Unicode-encoded text files into Unicode-encoded string objects. While CCL doesn't define its own Unicode string types, it does define its own types for performing I/O, it doesn't integrate with the C++ IO streams in the standard library.

When reading Unicode text, you can mix and match all you want, so for example, you can read a line of text from a file encoded in UTF-8 and store that string as UTF-32 in a `std::u32string`. The chapter on I/O will go over how I/O is done in CCL but the basic gist of it is that to read from a file you must instantiate the `cmp::file` class with the path to the file you want to open, and then you read from that file through a stream. When you instantiate the stream, you specify the file object as well as the encoding that the file is in (that is, whether it is UTF-8, UTF-16 or UTF-32). When reading a Unicode text file, you have to know what encoding the file is in. That's just how Unicode works, you have to know the encoding beforehand. Here's some code to demonstrate how all of this works:

```
#include <cmp/io/file.hpp>
#include <cmp/io/text_input_stream.hpp>
...
cmp::file f{"utf8_file.txt", cmp::read_only, cmp::if_not_there::fail};

if (!f.is_open()) {
    std::cout << "utf8_file.txt could not be opened for reading.\n";
    return -1;
}
```

```

cmp::text_input_stream stream{f, cmp::utf8};
std::u32string utf32_line;
stream.read_line(utf32_line);
// Now do something with utf32_line.

```

Notice how **cmp::utf8** is specified when the stream is constructed. This tells the stream what encoding the file is in. The other options are **cmp::utf16** and **cmp::utf32** for reading UTF-16 and UTF-32 text respectively. It is important to note that the code above shows how to read UTF-8 text, which is not affected by endianness. To read UTF-16 text or UTF-32 text, you must check for a BOM (byte order mark), which is discussed later in this chapter.

Writing Unicode Strings to Files

You can also write Unicode-encoded string objects out to a file. You can mix and match all you want here too, so for example, you can write a string object encoded in UTF-16 out to a file as UTF-8. Again, you instantiate **cmp::file** with the path to the file you want to open and you write to that file through a stream. Here's what that looks like:

```

#include <cmp/io/file.hpp>
#include <cmp/io/text_output_stream.hpp>

...

cmp::file f{"utf8_file.txt", cmp::write_only, cmp::if_not_there::fail};

if (!f.is_open()) {
    std::cout << "utf8_file.txt could not be opened for writing.\n";
    return -1;
}

cmp::text_output_stream stream{f, cmp::utf8};
std::u16string utf16_string{get_a_string_somewhat()};
stream << utf16_string;
// Now utf16_string has been written out as UTF-8.

```

It is important to note that the code above shows how to write a string without first writing a BOM (byte order mark). Writing a BOM is not strictly required but it makes it easier for other programs to interpret the contents of a Unicode text file because it specifies what endianness the file's contents are in. Let's look at how that works now.

Endianness

One challenge to supporting Unicode is endianness. Endianness decides whether the bytes of a multi-byte value are ordered from most significant to least significant, or least significant to most significant, and it applies to every multi-byte value your CPU manipulates, for example **ints** and **floats**. There are two types of endianness: big endian and little endian. In big endian systems, the bytes are ordered from most significant first (in earlier addresses) to least significant last (in later addresses). In little endian systems, the bytes are ordered from least significant first (in earlier addresses) to most significant last (in later addresses).

Endianness is a property of the CPU architecture. For example, x86 CPUs arrange multi-byte values in little endian order, whereas SPARC V8 CPUs arrange multi-byte values in big endian order. The default endianness of TCP/IP is big endian, meaning when two computers communicate with each other, data sent across the network is formatted in big endian order, and if one of the computers uses a little endian CPU then that computer will have to reverse the bytes every time it sends and receives data so that it can correctly interpret and process the data.

Let's look at an example. Imagine a 16-bit number with the decimal value 512. In C++ this could be a **std::uint16_t**. The two bytes that make up that 16-bit number would show up like this in big endian order: [earliest address] 00000010 00000000 [latest address]. In little endian order, the bytes would show up like this: [earliest address] 00000000 00000010 [latest address]. In both cases, the mathematical value 512 is

represented, just in different ways. Notice how only the order of the **bytes** changes, not the order of the bits in each byte.

Because code units in UTF-16 and UTF-32 are comprised of more than one byte, it is important to know the order of those bytes. If you want to create a text file in UTF-16 or UTF-32, and you want any Unicode-compliant program to be able to read that file correctly, then you can either just default to big endian order or you can specify a **byte order mark** (BOM). The BOM is just the code point 0xFEFF, and it must be at the very beginning of the file. If a BOM is present, it must come before everything else in the file. When a program reads that code point it will know what endianness is being used in the text file.

UTF-16 BOMs are two bytes long, and UTF-32 BOMs are four bytes long. In UTF-16, if the first byte of the BOM is 0xFE and the second byte is 0xFF, then the BOM is signaling that the file's contents are in big endian order. If you order the bytes the other way around then the BOM signals little endian order. UTF-32 uses the same technique but two of the four bytes have the value zero because the code point 0xFEFF fits in just 2 bytes.

If no BOM is present then the Unicode standard says that, in the absence of a higher-level protocol, the default endianness is big endian. This means that Unicode does let you read a file according to little endian order if you know that the file is in little endian order, even when no BOM is present. This can happen if, for example, a video game always writes save files in little endian order, because that file is subject to a higher-level protocol (in this case the video game) and it isn't expected to be opened and read in a regular text editor. For general interchangeable text files, the rule is to use big endian order if no BOM is present.

Handling Endianness in CCL

CCL does provide a way to handle endianness. In CCL, when you want to read a file, you instantiate the `cmp::file` class specifying the path to the file you want to open. Then you instantiate a stream and pass in the file object you created earlier, as well as the encoding that the file is in. Then you read all data from the file through the stream object, but before reading anything, you call the `read_bom` function on the stream, which will read the BOM (if present) and adjust the endianness of the stream according to the BOM read from the file. Regardless of whether there really was a BOM in the file, the `read_bom` function will return the code point that was read, which you can compare with the constant `cmp::bom`. If the returned code point compares equal to `cmp::bom` then a BOM was read and the endianness of the stream was set to the endianness of the file, meaning all future read operations will produce valid and accurate data. However, if the returned code point does not compare equal to `cmp::bom` then a regular character was read, and is part of the text content of the file. You can do whatever you want with this character. If your intention was to read the entire file into a string object, then just append that character to a string object and then call the `append_all` function on the stream and give it the string object. That will make the specified string object contain the entire contents of the file. Here's some code to illustrate this:

```
#include <cmp/io/file.hpp>
#include <cmp/io/text_input_stream.hpp>

...

cmp::file f{"utf16_file.txt", cmp::read_only, cmp::if_not_there::fail};

if (!f.is_open()) {
    std::cout << "utf16_file.txt could not be opened for reading.\n";
    return -1;
}

cmp::text_input_stream stream{f, cmp::utf16};
std::u8string file_contents;
char32_t possible_bom{stream.read_bom()};
if (possible_bom != cmp::bom) {
    cmp::append_code_point(file_contents, possible_bom);
}
stream.append_all(file_contents);
// Now do something with file_contents. file_contents has the file's accurate
// contents regardless of whether there was a BOM. If no BOM was present then
// the endianness specified in the stream's constructor was used. The stream's
```

```
// constructor defaults to big endian if you don't specify the endianness, thus
// conforming to the Unicode standard with minimal effort.
```

To write a Unicode string to a file, instantiate the `cmp::file` class specifying the path to the file you want to open, and write the string to the file with a stream. To write a BOM to the file, just call the `write_bom` function on the stream. Here's some code to illustrate this:

```
#include <cmp/io/file.hpp>
#include <cmp/io/text_output_stream.hpp>

...

cmp::file f{"utf16_file.txt", cmp::write_only, cmp::if_not_there::fail};

if (!f.is_open()) {
    std::cout << "utf16_file.txt could not be opened for writing.\n";
    return -1;
}

cmp::text_output_stream stream{f, cmp::utf16};
std::u8string content{get_a_string_somewhat()};
stream.write_bom();
stream << content;
// Now a BOM and a string have been written to the file as UTF-16.
// Again, the stream defaults to big endian.
```

`std::endian` has three members: `std::endian::little` (which represents little endian), `std::endian::big` (which represents big endian), and `std::endian::native` (which represents the endianness of the platform you're compiling for). `std::endian::native` may be equivalent to `std::endian::little`, `std::endian::big` or neither. When `std::endian::native` is neither `std::endian::little` nor `std::endian::big`, that means the platform you're compiling for has **mixed** endianness. It is important to note that CCL assumes that a platform is either fully little endian or fully big endian, and will not work on platforms of mixed endianness. A compiler error will automatically be generated by a static assertion when mixed endianness is detected.

Source File Encoding

GCC and Clang assume that source files are encoded in UTF-8, whereas Visual C++ assumes UTF-16. However, you can tell Visual C++ that your source files are encoded in UTF-8 by giving it the `/utf-8` flag. By using `/utf-8` on Visual C++ and encoding your source files in UTF-8, you get the guarantee that the Unicode text you put in string literals, rendered beautifully in your code editor or IDE, will appear picture-perfect when your compiled program prints it out. To use UTF-8, use a Unicode-capable code editor or IDE, examples of which include Visual Studio, Visual Studio Code, CLion, Code::Blocks, and others. Most of these editors show the encoding of the file in the bottom-right corner of the window. If you see "UTF-8" down there then your file is encoded correctly and you can enjoy the advanced Unicode editing capabilities of your editor, and you can go nuts on typing out accents, math symbols, greek letters, and so on, in your string literals. If you link to CCL in a CMake project, then your own target will automatically get the `/utf-8` flag on Visual C++. You only need to provide this flag to Visual C++ if you use CCL header-only or if you link to CCL without CMake.

Chapter 3: I/O

CCL has its own infrastructure for doing I/O (input and output). In this chapter we look at how to read and write text and binary data. Currently, CCL only supports doing I/O on files and containers, as well as standard I/O, but a future release will add support for doing I/O on network sockets.

Chapter Table of Contents

- Introduction

 - Transfer Resources

 - Transfer Streams

 - How They Fit Together

- Files

 - Transfer Streams

 - Data Input Streams

 - Data Output Streams

 - Text Input Streams

 - Text Output Streams

- Container Transfer Resources

- Wrapping Up

Introduction

CCL separates the entities on which I/O takes place from the interpretation of the data that flows through them. In CCL, I/O takes place on transfer resources. A transfer resource knows how to read or write bytes, but has no concept of what those bytes mean. It is the job of transfer streams to interpret those bytes.

Transfer Resources

An **input resource** is a source of input, and an **output resource** is a destination for output. Input resources and output resources are collectively called **transfer resources**. An object that is both an input resource and an output resource is called an **I/O resource**. The `cmp::file` class represents a file and it is an I/O resource because it supports both reading and writing.

Creating your own transfer resource types is a more advanced topic and the details of doing this are not documented yet because CCL's internal API is still unstable. A future version of CCL will have a more mature internal API and a future version of this book will include a thorough explanation on how to use it to create your own transfer resource types.

Transfer Streams

An **input stream** is an object that reads data from an input resource, and an **output stream** is an object that writes data to an output resource. Input streams and output streams are collectively called **transfer streams**. An object that is both an input stream and an output stream is called an **I/O stream**. CCL provides two kinds of streams: those that carry binary data and those that carry text data. The full list of transfer stream templates in CCL is shown in Table 3-1.

Name	Purpose
<code>cmp::data_input_stream</code>	Reads binary data from an input resource.
<code>cmp::data_output_stream</code>	Writes binary data to an output resource.
<code>cmp::data_io_stream</code>	Reads and writes binary data from/to an I/O resource.
<code>cmp::text_input_stream</code>	Reads text from an input resource.
<code>cmp::text_output_stream</code>	Writes text to an output resource.
<code>cmp::text_io_stream</code>	Reads and writes text from/to an I/O resource.

Table 3-1: CCL's transfer stream templates.

These are templates and must be instantiated with the type of transfer resource that they are going to work with.

How They Fit Together

To perform I/O, you first instantiate the transfer resource, then you instantiate the transfer stream with the transfer resource, and lastly you use the transfer stream to do the I/O. As an example, let's look at how to read an integer from a file. To do that, you instantiate the `cmp::file` class with the path to the file you want to read from. When dealing with a file, you should then check if the file really is open, which may not be the case if, for example, the file you specified doesn't exist. Then, with the file object open, you instantiate the input stream with the file object. Finally, you read from the file through the stream object. Let's now look at how to do this in code:

```
#include <cmp/io/file.hpp>
#include <cmp/io/data_input_stream.hpp>

...

cmp::file f{"binary_file.dat", cmp::read_only, cmp::if_not_there::fail};

if (!f.is_open()) {
    cmp::uout << "binary_file.dat could not be opened for reading.\n";
    return -1;
}
```

```
cmp::data_input_stream stream{f, std::endian::native};
int number;
stream >> number;
```

```
cmp::uout << "The file contained the number " << number << '\n';
```

The transfer resources are decoupled from the transfer streams. In addition to `cmp::file`, CCL comes with container transfer resource types, which are used to do I/O on containers. The container I/O resource types are similar in spirit to `std::stringstream`, except they can read/write from/to any container that stores its elements contiguously, not just strings. Container transfer resources will be discussed later, but now let's take a closer look at the options we have with files, and subsequently we will look at transfer streams in detail.

Files

The `cmp::file` class has a constructor with the following parameters:

- `const std::filesystem::path& file_path`
The path to the file you want to open.
- `cmp::access_mode mode`
The access mode to use, indicating whether you want to open the file for reading, writing or both.
- `cmp::if_not_there if_file_not_there`
What to do if the requested file does not exist: whether to fail to open the file or to create it.
- `std::size_t buffer_capacity = cmp::io_buffer::default_buffer_capacity`
The capacity of the I/O buffer. Specify 0 for unbuffered operations.

For example, to open a file just for writing, and to create it if it doesn't already exist, write this:

```
cmp::file f{"path/to/file.txt", cmp::write_only, cmp::if_not_there::create};
```

To open a file for reading and writing, and to fail if it doesn't already exist, and to not buffer any I/O operations, write this:

```
cmp::file f{"path/to/file.dat", cmp::read_and_write, cmp::if_not_there::fail, 0};
```

Transfer Streams

Let's now turn our attention to transfer streams, starting with data input streams.

Data Input Streams

The `cmp::data_input_stream` class template has a constructor with the following parameters:

- `referenced_resource_type& resource`
The input resource to read from.
- `std::endian endianness = std::endian::big`
The endianness the input data is expected to be in.

For example, to get ready to read from a file object called `f`, whose contents are in little endian order, write this:

```
cmp::data_input_stream stream{f, std::endian::little};
```

Then, to read from the file, use the overloaded `>>` operator. Like this:

```
int integer;  
stream >> integer;
```

Data input streams support reading data in all fundamental types.

Data Output Streams

The `cmp::data_output_stream` class template has a constructor with the following parameters:

- `referenced_resource_type& resource`
The output resource to write to.
- `std::endian endianness = std::endian::big`
The endianness the output data is going to be in.

For example, to get ready to write to a file object called `f`, whose contents will be in little endian order, write this:

```
cmp::data_output_stream stream{f, std::endian::little};
```

Then, to write to the file, use the overloaded `<<` operator. Like this:

```
double number{get_some_number()};  
stream << number;
```

Data output streams support writing data in all fundamental types.

Text Input Streams

Text streams work the same way except that they must know the encoding to use.

The `cmp::text_input_stream` class template has a constructor with the following parameters:

- `referenced_resource_type& resource`
The input resource to read from.
- `encoding_form source_encoding_form`
The Unicode encoding form that the input text is expected to be in.
- `std::endian endianness = std::endian::big`
The endianness the input text is expected to be in.

For example, to get ready to read from a file object called `f`, whose contents is in UTF-8 and little endian order, write this:

```
cmp::text_input_stream stream{f, cmp::utf8, std::endian::little};
```

Then, to read from the file object, use the overloaded `>>` operator. Like this:

```
int integer;  
stream >> integer;
```

In this example, we are reading an integer, which means we are going to parse an integer from the text. If a valid integer cannot be parsed, we'll get an exception.

Text input streams support reading data in all fundamental types.

Text Output Streams

The `cmp::text_output_stream` class template has a constructor with the following parameters:

- referenced_resource_type& **resource**
The output resource to write to.
- encoding_form **target_encoding_form**
The Unicode encoding form that the output text is going to be in.
- std::endian **endianness** = std::endian::big
The endianness the output data is going to be in.

For example, to get ready to write to a file object called `f`, whose contents will be in UTF-8 and little endian order, write this:

```
cmp::text_output_stream stream{f, cmp::utf8, std::endian::little};
```

Then, to write to the file, use the overloaded `<<` operator. Like this:

```
double number{get_some_number()};  
stream << number;
```

In this example, we are writing a floating-point number, which means we are going to write out the text representation of that number.

Text output streams support writing data in all fundamental types.

Container Transfer Resources

Container transfer resources behave similarly to files, except you read from and write to a container, instead of a file. The full list of container transfer resource templates is shown in Table 3-2. An opaque container transfer resource owns the container that it operates on, whereas a transparent one does not.

Name	Purpose
<code>cmp::opaque_container_input_resource</code>	Owning input resource that reads from a container.
<code>cmp::opaque_container_output_resource</code>	Owning output resource that writes to a container.
<code>cmp::opaque_container_io_resource</code>	Owning I/O resource that reads and writes from/to a container.
<code>cmp::transparent_container_input_resource</code>	Non-owning input resource that reads from a container.
<code>cmp::transparent_container_output_resource</code>	Non-owning output resource that writes to a container.
<code>cmp::transparent_container_io_resource</code>	Non-owning I/O resource that reads and writes from/to a container.

Table 3-2: CCL's container transfer resource templates.

These are just templates and must be instantiated with a container type. Only containers that store its elements contiguously are supported. CCL provides type aliases for various specializations. Table 3-3 shows all specializations for opaque container transfer resource types, but transparent equivalents exist too and follow the same naming scheme.

Name	Purpose
<code>cmp::opaque_blob_input_resource</code>	Owning input resource that reads from a <code>std::vector<std::byte></code> .
<code>cmp::opaque_blob_output_resource</code>	Owning output resource that writes to a <code>std::vector<std::byte></code> .
<code>cmp::opaque_blob_io_resource</code>	Owning I/O resource that reads and writes from/to a <code>std::vector<std::byte></code> .
<code>cmp::opaque_string_input_resource</code>	Owning input resource that reads from a <code>std::string</code> .
<code>cmp::opaque_string_output_resource</code>	Owning output resource that writes to a <code>std::string</code> .
<code>cmp::opaque_string_io_resource</code>	Owning I/O resource that reads and writes from/to a <code>std::string</code> .
<code>cmp::opaque_u8string_input_resource</code>	Owning input resource that reads from a <code>std::u8string</code> .
<code>cmp::opaque_u8string_output_resource</code>	Owning output resource that writes to a <code>std::u8string</code> .
<code>cmp::opaque_u8string_io_resource</code>	Owning I/O resource that reads and writes from/to a <code>std::u8string</code> .
<code>cmp::opaque_u16string_input_resource</code>	Owning input resource that reads from a <code>std::u16string</code> .
<code>cmp::opaque_u16string_output_resource</code>	Owning output resource that writes to a <code>std::u16string</code> .
<code>cmp::opaque_u16string_io_resource</code>	Owning I/O resource that reads and writes from/to a <code>std::u16string</code> .
<code>cmp::opaque_u32string_input_resource</code>	Owning input resource that reads from a <code>std::u32string</code> .
<code>cmp::opaque_u32string_output_resource</code>	Owning output resource that writes to a <code>std::u32string</code> .
<code>cmp::opaque_u32string_io_resource</code>	Owning I/O resource that reads and writes from/to a <code>std::u32string</code> .
<code>cmp::opaque_wstring_input_resource</code>	Owning input resource that reads from a <code>std::wstring</code> .
<code>cmp::opaque_wstring_output_resource</code>	Owning output resource that writes to a <code>std::wstring</code> .
<code>cmp::opaque_wstring_io_resource</code>	Owning I/O resource that reads and writes from/to a <code>std::wstring</code> .

Table 3-3: CCL's opaque container transfer resource type aliases.

Transparent equivalents are named similarly, just replace `opaque_` with `transparent_`.

To illustrate, here is how you would read from a `std::u8string`:

```
const std::u8string& s{grab_some_string()};
cmp::transparent_u8string_input_resource resource{s};
cmp::text_input_stream stream{resource, cmp::utf8, std::endian::native};
std::u8string content;
stream.read_all(content);
cmp::uout << content << '\n';
```

This example is a little silly because we are just reading a `std::u8string` in its entirety with the call to `read_all`, and then printing it. The power of container transfer resources becomes clear when you need to read or write container data methodically, one piece at a time.

To write to a string, use the output counterparts: `cmp::transparent_u8string_output_resource` and `cmp::text_output_stream`. Like this:

```
std::u8string content;
cmp::transparent_u8string_output_resource resource{content};
cmp::text_output_stream stream{resource, cmp::utf8, std::endian::native};
```



```
stream << u"abc\nxyz";  
cmp::uout << content << '\n';
```

Here we create a container output resource, write a UTF-16 string literal to it through a stream, and then print the string we've built up. The stream knows that it's supposed to deliver UTF-8 content to its output resource so it automatically converts the UTF-16 text to UTF-8. This example isn't particularly fancy, it just serves to illustrate how these pieces fit together.

Wrapping Up

CCL is very capable when it comes to I/O. You can read and write both binary data and Unicode text. You can manipulate files, containers and strings. The general approach is to create your transfer resource first, then create your transfer stream with the transfer resource, and finally perform all of your I/O with the stream.

Chapter 4: Desktop GUI

CCL provides very simplistic GUI support, which will be extended significantly in future versions. The GUI capabilities available now are only the beginning, despite being very basic.

Chapter Table of Contents

Supported Platforms

An Example GUI Application

Overview

```
#include <cmp/desktop_gui/main.hpp>
```

```
#include <cmp/desktop_gui/desktop_gui_application.hpp>
```

Reacting to Input

Wrapping Up

Supported Platforms

Before we delve into code, let's just take a moment to appreciate what the different platforms that CCL Desktop GUI supports look like. CCL Desktop GUI runs on all desktop-flavored operating systems that CCL supports. Namely: Windows, macOS, GNU/Linux and FreeBSD. While Windows and macOS are desktop-exclusive, GNU/Linux and FreeBSD are not. CCL Desktop GUI doesn't run on non-desktop flavors of these operating systems, like Android for example.

An Example GUI Application

The code below was not taken from any example application that ships with CCL, but it serves to illustrate the structure of a GUI application made with CCL Desktop GUI.

```
#include <cmp/desktop_gui/main.hpp>
#include <cmp/desktop_gui/desktop_gui_application.hpp>
#include <cmp/desktop_gui/window.hpp>

int main (
    int argc,
    char** argv
) {
    cmp::desktop_gui_application app{argc, argv};
    cmp::window w{
        400,
        300,
        u8"Application Window",
        cmp::window_mode::standard
    };
    w.show();
    return app.run();
} // function -----
```

Overview

We start off by including CCL Desktop GUI headers. Then, in **main**, we create the application object, which will take care of initialization and, later on, the event loop. Then we create a 400 (width) by 300 (height) window with the title "Application Window" and call **show** on it, which displays it to the user. Then we enter the event loop by calling **run** on the application object, and when all windows are closed, the event loop will end and **main** will return 0 (zero).

#include <cmp/desktop_gui/main.hpp>

When developing GUI applications on Windows, there are two settings that are important to know about. One of them is the **subsystem** and the other is the **entry**. If you set the subsystem to **WINDOWS** then the application will not create a console window, which is often desirable when developing a GUI. If the entry is set to **mainCRTStartup** then the application will begin execution in the traditional **main** function, rather than the Windows-specific **WinMain** function. In CMake, you can specify **WIN32** right after the name of your executable target, which will set the subsystem to **WINDOWS**, which will prevent your application from creating a console window when it runs. However, it is still necessary to override the entry, and that is what the **main.hpp** header file does. This header file makes your compiler use the standards-conforming version of **main**, and it only needs to be included in one of your application's source files. Without it, your application would begin execution in the non-portable **WinMain** function. However, this header file is only necessary if you are on Windows. Furthermore, it has no effect on other platforms, it is always safe to include it, so if you're going for cross-platform development and you don't want to think about what operating system you're on, then the recommendation is that you include this header file unconditionally, as is done in the example code shown earlier.

CCL does not accomplish portability by defining **WinMain** and then calling the traditional **main** function. When using CCL, your applications actually use the standards-conforming **main** function natively (on all platforms) and that is where execution of your program really begins.

#include <cmp/desktop_gui/desktop_gui_application.hpp>

In CCL, there is the concept of an application object. The application object takes care of all the initialization required to make a certain kind of program work. The class `cmp::desktop_gui_application` represents an application object for desktop GUI applications. You create the application object in `main` and never again, as it is a singleton. In the case of `cmp::desktop_gui_application`, the application object will then allow you to spin up the event loop, which will make sure that any window you have open in your application will be notified of GUI events. Once all windows are closed, the event loop will stop, and in the case of the example code shown earlier, the program will return 0 and end.

Reacting to Input

Just creating and displaying a window isn't very useful, so let's make the application respond to input. To do that, we must subclass the `cmp::window` class, override the `handle_key_down_event` and `handle_key_up_event` functions, then instantiate that subclass and call `show` on it. Then, when the user pushes a key down or releases a key, our custom event-handling code will be executed. Here's the code:

```
#include <cmp/desktop_gui/main.hpp>
#include <cmp/desktop_gui/desktop_gui_application.hpp>
#include <cmp/desktop_gui/window.hpp>

class app_window
    : public cmp::window
{
public:
    app_window ()
        : cmp::window{
            400,
            300,
            u8"Application Window",
            cmp::window_mode::standard
        }
    {
    } // function -----

    void
    handle_key_down_event (
        cmp::key_event& ev
    )
    override
    {
        cmp::uout << cmp::to_u8string(ev.get_key()) << " down.\n" << cmp::flush;
    } // function -----

    void
    handle_key_up_event (
        cmp::key_event& ev
    )
    override
    {
        cmp::uout << cmp::to_u8string(ev.get_key()) << " up.\n" << cmp::flush;
    } // function -----
}; // class -----

int main (
    int argc,
    char** argv
) {
    cmp::desktop_gui_application app{argc, argv};
    app_window w;
    w.show();
    return app.run();
} // function -----
```

This code just gets the string representation of the key that the user presses or releases, and prints it to standard output. Figure 4-1 shows the output you get on Windows when you run the program and press the A key followed by the 5 key on the numpad:




Figure 4-1: The GUI application shown earlier running on Windows 10.

Don't mind the non-default font in the command line window, the default font doesn't support some Unicode characters so I just use a font that does support the Unicode characters I care about.

There are other events you can handle in a `cmp::window` subclass. The complete list is shown in Table 4-1.

Member Function	Description
<code>handle_key_down_event</code>	Called when a key is pressed.
<code>handle_key_up_event</code>	Called when a key is released.
<code>handle_resize_event</code>	Called when the window is resized.
<code>handle_close_event</code>	Called when the window is asked to be closed.

Table 4-1: The event handler functions you can override from `cmp::window`.

The resize event handler is pretty straightforward, it doesn't take any arguments. It just signals that the window was resized. The close event handler, however, is a bit more sophisticated. In the close event handler, you can decide whether the window is actually going to be closed or not. You do this by calling `set_should_close` on the argument that the close event handler receives. That kind of thing could look like this:

```
...

void
handle_close_event (
    cmp::close_event& ev
)
override
{
    // somehow decide whether to close the window
    if (!confirm_exit()) {
        // make it so that the window does not get closed
        ev.set_should_close(false);
    }
} // function -----
...

```

CCL does not yet provide a way to display a message box, so the code shown above hides that with a call to a fictitious function called `confirm_exit`, but the point of this code is to show that you can control whether the window will actually close by calling `set_should_close` on the `cmp::close_event` argument. The default (if you don't explicitly do anything) is for the window to be closed. The window will **not** be closed if you call `ev.set_should_close(false)`.

Wrapping Up

CCL's GUI features are very promising. While you can't do very much just yet, you can already do simple things like create a window and respond to keyboard events. A lot more is planned for future versions.