

BARTŁOMIEJ FILIPEK

C++17 IN DETAIL

LEARN THE EXCITING FEATURES OF
THE NEW C++ STANDARD!

(BF)
C++ STORIES

BFILIPEK.COM

C++17 in Detail

Learn the Exciting Features of The New C++ Standard!

Bartłomiej Filipek

This book is for sale at <http://leanpub.com/cpp17indetail>

This version was published on 2019-03-01



This is a [Leanpub](#) book. Leanpub empowers authors and publishers with the Lean Publishing process. [Lean Publishing](#) is the act of publishing an in-progress ebook using lightweight tools and many iterations to get reader feedback, pivot until you have the right book and build traction once you do.

© 2018 - 2019 Bartłomiej Filipek

for Viola and Mikołaj

Contents

About the Author	i
Technical Reviewer	ii
Additional Reviewers & Supporters	iii
Revision History	v
Preface	vi
About the Book	vii
Who This Book is For	vii
Overall Structure of the Book	viii
Reader Feedback	ix
Example Code	ix
 Part 1 - The Language Features	 1
Quick Start	2
 1. Fixes and Deprecation	 5
Removed Things	6
Fixes	10
Compiler support	13
 2. Language Clarification	 14
Stricter Expression Evaluation Order	15
Guaranteed Copy Elision	19
Dynamic Memory Allocation for Over-Aligned Data	24
Exception Specifications as Part of the Type System	25
Compiler Support	25
 3. General Language Features	 26
Structured Binding Declarations	27
Init Statement for <code>if</code> and <code>switch</code>	34
Inline Variables	36
<code>constexpr</code> Lambda Expressions	38

Nested Namespaces	40
Compiler support	41
4. Templates	42
Template Argument Deduction for Class Templates	43
Fold Expressions	47
if constexpr	50
Declaring Non-Type Template Parameters With auto	58
Other Changes	59
Compiler Support	61
5. Standard Attributes	62
Why Do We Need Attributes?	63
Before C++11	63
Attributes in C++11 and C++14	64
C++17 additions	66
Section Summary	71
Compiler support	72

Part 2 - The Standard Library Changes 73

6. std::optional	74
Introduction	75
std::optional Creation	77
Returning std::optional	81
Accessing The Stored Value	83
std::optional Operations	84
Examples of std::optional	86
Performance & Memory Consideration	88
Migration from boost::optional	90
Special case: optional<bool> and optional<T*>	90
Summary	91
Compiler Support	91
7. std::variant	92
The Basics	93
std::variant Creation	96
Changing the Values	99
Accessing the Stored Value	101
Visitors for std::variant	102
Other std::variant Operations	107
Exception Safety Guarantees	107
Performance & Memory Considerations	108

CONTENTS

Migration From <code>boost::variant</code>	109
Examples of <code>std::variant</code>	110
Wrap Up	118
Compiler Support	118
8. <code>std::any</code>	119
The Basics	120
<code>std::any</code> Creation	122
Changing the Value	124
Accessing The Stored Value	125
Performance & Memory Considerations	126
Migration from <code>boost::any</code>	127
Examples of <code>std::any</code>	127
Wrap Up	130
Compiler Support	130
9. <code>std::string_view</code>	131
The Basics	132
The <code>std::basic_string_view</code> Type	133
<code>std::string_view</code> Creation	134
Other Operations	135
Risks Using <code>string_view</code>	137
Initializing <code>string</code> Members from <code>string_view</code>	141
Handling Non-Null Terminated Strings	145
Performance & Memory Considerations	147
Migration from <code>boost::string_ref</code> and <code>boost::string_view</code>	148
Examples	149
Wrap Up	152
10. String Conversions	153
Elementary String Conversions	154
Converting From Characters to Numbers: <code>from_chars</code>	155
Converting Numbers into Characters: <code>to_chars</code>	158
The Benchmark	161
Summary	165
Compiler support	166
11. Searchers & String Matching	167
Overview of String Matching Algorithms	168
New Algorithms Available in C++17	169
Examples	170
Summary	176
Compiler support	176

CONTENTS

12. Filesystem	177
Filesystem Overview	178
Demo	178
The Path Object	181
The Directory Entry & Directory Iteration	189
Supporting Functions	190
Error Handling & File Races	196
Examples	197
Chapter Summary	202
Compiler Support	204
13. Parallel STL Algorithms	205
Introduction	206
Overview	207
Execution Policies	208
Algorithm Update	213
New Algorithms	214
Performance of Parallel Algorithms	218
Examples	219
Chapter Summary	230
Compiler Support	232
14. Other Changes In The Library	233
std::byte	234
Improvements for Maps and Sets	235
Return Type of Emplace Methods	241
Sampling Algorithms	242
New Mathematical Functions	243
Shared Pointers and Arrays	245
Non-member size(), data() and empty()	246
constexpr Additions to the Standard Library	247
std::scoped_lock	249
std::iterator Is Deprecated	250
Polymorphic Allocator, pmr	252
Compiler support	254

Part 3 – More Examples and Use Cases255

15. Refactoring with std::optional and std::variant	256
The Use Case	257
The Tuple Version	258
A Separate Structure	259
With std::optional	260

CONTENTS

With <code>std::variant</code>	261
Wrap up	263
16. Enforcing Code Contracts With <code>[[nodiscard]]</code>	264
Introduction	265
Where Can It Be Used?	265
How to Ignore <code>[[nodiscard]]</code>	268
Before C++17	269
Summary	269
17. Replacing <code>enable_if</code> with <code>if constexpr</code> - Factory with Variable Arguments	270
The Problem	271
Before C++17	273
With <code>if constexpr</code>	274
Summary	275
18. How to Parallelise CSV Reader	276
Introduction and Requirements	277
The Serial Version	278
Using Parallel Algorithms	284
Wrap up & Discussion	290
Appendix A - Compiler Support	293
GCC	293
Clang	293
VisualStudio - MSVC	293
Compiler Support of C++17 Features	294
Appendix B - Resources and References	297

About the Author

Bartłomiej Filipek is a C++ software developer with more than 11 years of professional experience. In 2010 he graduated from Jagiellonian University in Cracow with a Masters Degree in Computer Science.

Bartek currently works at [Xara](#), where he develops features for advanced document editors. He also has experience with desktop graphics applications, game development, large-scale systems for aviation, writing graphics drivers and even biofeedback. In the past, Bartek has also taught programming (mostly game and graphics programming courses) at local universities in Cracow.

Since 2011 Bartek has been regularly blogging at his website: bfilipek.com. In the early days the topics revolved around graphics programming, and now the blog focuses on Core C++. He also helps as co-organizer at [C++ User Group in Krakow](#). You can hear Bartek in one [@CppCast episode](#) where he talks about C++17, blogging and text processing.

Since October 2018, Bartek has been a C++ Expert for Polish National Body that works directly with ISO/IEC JTC 1/SC 22 (C++ Standard Committee). In the same month, Bartek was also awarded by Microsoft and got his first MVP title for years 2019/2020.

In his spare time, he loves assembling trains and Lego with his little son. And he's a collector of large Lego models.

Technical Reviewer

Jacek Galowicz is a Software Engineer with roughly a decade of professional experience in C++. He got his master of science degree in electrical engineering at RWTH Aachen University in Germany.

Jacek co-founded the Cyberus Technology GmbH in early 2017 and works on products around low-level cybersecurity, virtualization, microkernels, and advanced testing infrastructure. At former jobs, he implemented performance- and security-sensitive microkernel operating systems for Intel x86 virtualization at Intel and FireEye in Germany. In general, he gained experience with kernel driver development, 3D graphics programming, databases, network communication, physics simulation, mostly in C or C++.

In his free time, Jacek maintains a little C++ blog, which has seen some lack of love while he wrote the C++17 STL Cookbook. He is a regular visitor of the C++ Usergroups in Hannover and Braunschweig. In order to do meta programming and generic programming better, he also learned and uses Haskell, which in turn sparked his interest to generally bring the advantages of purely functional programming to C++.

Additional Reviewers & Supporters

Without the support of many good people, this book would have been far less than it is. It is a great pleasure to thank them. A lot of people read drafts, found errors, pointed out confusing explanations, suggested rewordings, tested programs, and offered support and encouragement. Many reviewers generously supplied insights and comments that I was able to incorporate into the book. Any mistakes that remain are, of course, my own.

Thanks especially to the following reviewers, who either commented on large sections of the book, smaller parts or gave me a general direction for the whole project.

Patrice Roy - Patrice has been playing with C++, either professionally, for pleasure or (most of the time) both for over 20 years. After a few years doing R&D and working on military flight simulators, he moved on to academics and has been teaching computer science since 1998. Since 2005, he's been involved more specifically in helping graduate students and professionals from the fields of real-time systems and game programming develop the skills they need to face today's challenges. The rapid evolution of C++ in recent years has made his job even more enjoyable.

Jonathan Boccara - Jonathan is a passionate C++ developer working on large codebase of financial software. His interests revolve around making code expressive. He created and regularly blogs on [Fluent C++](#), where he explores how to use the C++ language to write expressive code, make existing code clearer, and also about how to keep your spirits up when facing unclear code.

Lukasz Rachwalski - Software engineer - founder C++ User Group Krakow.

Michał Czaja - C++ software developer and network engineer. Works in telecommunication industry since 2009.

Arne Mertz - Software Engineer from Hamburg, Germany. He is a C++ and clean code enthusiast. He's the author of the [Simplify C++](#) blog.

JFT - Has been involved with computer programming and "computer techy stuff" for over 45 years - including OS development and teaching c++ in the mid 1990's.

Victor Ciura - Senior Software Engineer at CAPHYON and Technical Lead on the [Advanced Installer team](#). For over a decade, he designed and implemented several core components and libraries of Advanced Installer. He's a regular guest at Computer Science Department of his Alma Mater, University of Craiova, where he gives student lectures & workshops on "Using C++STL for Competitive Programming and Software Development". Currently, he spends most of his time working with his team on improving and extending the repackaging and virtualization technologies in Advanced Installer IDE, helping clients migrate their Win32 desktop apps to the Windows Store (MSIX).

Karol Gasiński - Tech Lead on macOS VR in Apple's GPU SW Architecture Team. Previously Senior Graphics Software Engineer at Intel. As a member of KHRONOS group, he contributed to OpenGL 4.4 and OpenGL ES 3.1 Specifications. Except for his work, Karol's biggest passion is game development and low-level programming. He conducts research in the field of VR, new interfaces and game engines. In game industry is known from founding WGK conference. In the past, he was working on mobile versions of such games as Medieval Total War, Pro Evolution Soccer and Silent Hill.

Marco Arena - Software Engineer and C++ Specialist building mission critical and high performance software products in the Formula 1 Industry. Marco is very active in the C++ ecosystem as a blogger, speaker and organizer: he founded the Italian C++ Community in 2013, he joined isocpp.org editors in 2014 and he has been involved in community activities for many years. Marco has held the Microsoft MVP Award since 2016. Discover more at marcoarena.com.

Konrad Jaśkowiec - C++ expert with almost 8 years of professional experience at the time with prime focus on system design and optimization. You can find his profile at [Linkedin](#).

Revision History

- 10th August 2018 - the first release! The book is 90% ready!
- 31st August - minor release
 - Added section about [nested namespaces](#) in the General Language Features chapter
 - Added section about [using statement in folding expressions](#) in the Template chapter
 - Added more information about the overload pattern
 - A useful example of [std::visit with multiple variants](#), in the Variant chapter
 - improved “[Enforcing Code Contracts With \[\[nodiscard\]\]](#)” chapter
 - improved “[Refactoring with optional](#)” chapter - added info about `std::variant`
 - Grammar, typos, formatting issues, rewording
- 28th September 2018 - major release
 - New chapter - [String Conversions](#)
 - New chapter - [Searchers & String Matching](#)
 - Updated chapter about Parallel Algorithms Chapter, perf results, better explanations
 - Added notes about `gcd`, `lcm`, `clamp` in the Other STL Changes Chapter
 - Better explanations in many chapters like Variant, `string_view`, General Language
 - Typos, Grammar, formatting issues
- 3rd October 2018 - hot fixes
 - fixed line numbering in the demo section, grammar, typos, punctuation, clarification of the benchmark in [String Conversions](#)
- 4th November 2018 - major release
 - [Parallel Algorithms](#) was rewritten and is 3X larger, new examples and descriptions
- 21st December 2018 - major release
 - New chapter - [How to Parallelise CSV Reader](#)
- 18th January 2019 - major release - the book is 99% ready!
 - [Filesystem](#) chapter was rewritten and is 5X larger, new examples and descriptions
 - Grammar and style fixes for the [How to Parallelise CSV Reader](#) chapter
- 1st February 2019
 - Small additions to the [Filesystem](#) chapter: [file permissions](#) and [the regex example](#) updated
- 15th February 2019 - smaller updated
 - Better explanation in “[Structured Binding](#)” and “[if constexpr](#)”, style and grammar
- 1st March 2019 - the book is 100% ready!
 - Added [scoped_lock](#), [std::iterator deprecation](#) and [polymorphic memory allocator](#) sections in [Chapter 14](#), style and grammar

Preface

After the long awaited C++11, the C++ Committee has made changes to the standardisation process and we can now expect a new language standard every three years. In 2014 the ISO Committee delivered C++14. Now it's time for C++17, which was published at the end of 2017. As I am writing these words, in the middle of 2018, we're in the process of preparing C++20.

As you can see, the language and the Standard Library evolves quite fast! Since 2011 you've got a set of new library modules and language features every three years. Thus staying up to date with the whole state of the language has become quite a challenging task, and this is why this book will help you.

The chapters of this book describe all the significant changes in C++17 and will give you the essential knowledge to stay at the edge of the latest features. What's more, each section contains lots of practical examples and also compiler-specific notes to give you a more comfortable start.

It's a pleasure for me to write about new and exciting things in the language and I hope you'll have fun discovering C++17 as well!

Best regards,

Bartek

About the Book

C++11 was a major update for the language. With all the modern features like lambdas, constexpr, variadic templates, threading, range based for loops, smart pointers and many more powerful elements, it was enormous progress for the language. Even now, in 2018, lots of teams struggle to modernise their projects to leverage all the modern features. Later there was a minor update - C++14, which improved some things from the previous standard and added a few smaller elements. With C++17 we got a lot of mixed emotions.

Although C++17 is not as big as C++11, it's larger than C++14. Everyone expected modules, co-routines, concepts and other powerful features, but it wasn't possible to prepare everything on time.

Is C++17 weak?

Far from it! And this book will show you why!

The book brings you exclusive content about C++17 and draws from the experience of many articles that have appeared on bfilipek.com. The chapters were rewritten from the ground-up and updated with the latest information. All of that equipped with lots of new examples and practical tips. Additionally, the book provides insight into the current implementation status, compiler support, performance issues and other relevant knowledge to boost your current projects.

Who This Book is For

This book is intended for all C++ developers who have at least basic experience with C++11/14.

The principal aim of the book is to make you equipped with practical knowledge about C++17. After reading the book, you'll be able to move past C++11 and leverage the latest C++ techniques in your day to day tasks.

Please don't worry if you're not an expert in C++11/14. The book will give you necessary background, so you'll get the information in a proper context.

Overall Structure of the Book

C++17 brings a lot of changes to the language and to the Standard Library. In this book, all the new features were categorised into a few segments, so that they are easier to comprehend.

As a result, the book has the following sections:

- Part One - C++17 Language features:
 - Fixes and Deprecation
 - Language Clarification
 - General Language Features
 - Templates
 - Attributes
- Part Two - C++17 The Standard Library:
 - `std::optional`
 - `std::variant`
 - `std::any`
 - `std::string_view`
 - String Operations
 - Filesystem
 - Parallel STL
 - Other Changes
- Part Three - More Examples and Use Cases
- Appendix A - Compiler Support
- Appendix B - Resources and Links

The **first part** - about the language features - is shorter and will give you a quick run over the most significant changes. You can read it in any order you like.

The **second part** - describes a set of new library types that were added to the Standard. The helper types create a potential new vocabulary for C++ code: like when you use `optional`, `any`, `variant` or `string_view`. And what's more, you have new powerful capabilities especially in the form of parallel algorithms and the standard filesystem. A lot of examples in this part will use many other features from C++17.

The **third part** - brings together all of the changes in the language and shows examples where a lot of new features are used alongside. You'll see discussions about refactoring, simplifying code with new template techniques or working with parallel STL and the filesystem. While the first and the second part can also be used as a reference, the third part shows more of larger C++17 patterns that join many features.

A considerable advantage of the book is the fact that with each new feature you'll get information about the compiler support and the current implementation status. That way you'll be able to check if a particular version of the most popular compilers (MSVC, GCC or Clang) implements it or not. The book also gives practical hints on how to apply new techniques in your current codebase.

Reader Feedback

If you spot an error, typo, a grammar mistake... or anything else (especially some logical issues!) that should be corrected, then please let us know!

Write your feedback to bartlomiej.filipek AT bfilipek.com.

You can also use those two places:

- [Leanpub Book's Feedback Page](#)¹
- [GoodReads Book's Page](#)²

Example Code

You can find the ZIP package with all the example on the book's website: cppindetail.com/data/cpp17indetail.zip³. The same ZIP package should be also attached with the ebook.

A lot of the examples in the book are relatively short. You can copy and paste the lines into your favourite compiler/IDE and then run the code snippet.

Code License

The code for the book is available under the Creative Commons License.

Compiling

To use C++17 make sure you provide a proper flag for your compiler:

- in GCC (at least 7.1 or 8.0 or newer): use `-std=c++17` or `-std=c++2a`
- in Clang (at least 4.0 or newer): use `-std=c++17` or `-std=c++2a`
- in MSVC (Visual Studio 2017 or newer): use `/std:c++17` or `/std:c++latest` in project options -> C/C++ -> Language -> C++ Language Standard.

Formatting

The code is presented in a monospace font, similarly to the following example:

¹<https://leanpub.com/cpp17indetail/feedback>

²<https://www.goodreads.com/book/show/41447221-c-17-in-detail>

³<https://www.cppindetail.com/data/cpp17indetail.zip>

```
// Chapter Example/example_one.cpp
#include <iostream>

int main()
{
    std::cout << "Hello World\n";
}
```

Snippets of longer programs were usually shortened to present only the core mechanics. In that case, you'll find their full version in the separate ZIP package that comes with the book.

The corresponding file for the code snippet is mentioned in the first line - for example, “// Chapter Example/example_one.cpp”.

Usually, source code uses full type names with namespaces, like `std::string`, `std::filesystem::*`. However, to make code compact and present it nicely on a book page the namespaces sometimes might be removed, so they don't use space. Also, to avoid line wrapping longer lines might be manually split into two. In some case the code in the book might skip `include` statements.

Syntax Highlighting Limitations

The current version of the book might show some limitations regarding syntax highlighting.

For example:

- `if constexpr` - Link to Pygments issue: [#1432 - C++ if constexpr not recognized \(C++17\)](#)⁴
- The first method of a class is not highlighted - [#1084 - First method of class not highlighted in C++](#)⁵
- Template method is not highlighted [#1434 - C++ lexer doesn't recognize function if return type is templated](#)⁶
- Modern C++ attributes are sometimes not recognised properly

Other issues for C++ and Pygments: [issues C++](#)⁷.

Online Compilers

Instead of creating local projects you can also leverage some online compilers. They offer a basic text editor and usually allow you to compile only one source file (the code that you edit). They are very handy if you want to play with a simple code example.

For example, many of the the code samples for this book were created in Coliru Online compiler and then adapted adequately for the book content.

Here's a list of some of the useful services:

⁴<https://bitbucket.org/birkenfeld/pygments-main/issues/1432/c-if-constexpr-not-recognized-c-17>

⁵<https://bitbucket.org/birkenfeld/pygments-main/issues/1084/first-method-of-class-not-highlighted-in-c>

⁶<https://bitbucket.org/birkenfeld/pygments-main/issues/1434/c-lexer-doesnt-recognize-function-if>

⁷<https://bitbucket.org/birkenfeld/pygments-main/issues?q=c%2B%2B>

- [Coliru](#)⁸ - uses GCC 8.1.0 (as of July 2018), offers link sharing and a basic text editor, it's simple but very effective.
- [Wandbox](#)⁹ - it offers a lot of compilers - for example, most of Clang and GCC versions - and also you can use boost libraries. Also offers link sharing.
- [Compiler Explorer](#)¹⁰ - shows the compiler output from your code! Has many compilers to pick from.
- [CppBench](#)¹¹ - run a simple C++ performance tests (using google benchmark library).
- [C++ Insights](#)¹² - it's a Clang-based tool which does a source to source transformation. It shows how the compiler sees the code, for example by expanding lambdas, auto, structured bindings or range-based for loops.

There's also a nice list of online compilers gathered on this website: [List of Online C++ Compilers](#)¹³.

⁸<http://coliru.stacked-crooked.com/>

⁹<https://wandbox.org/>

¹⁰<https://gcc.godbolt.org/>

¹¹<http://quick-bench.com/>

¹²<https://cppinsights.io/>

¹³<https://arnemertz.github.io/online-compilers/>

Part 1 - The Language Features

C++17 is a major update of C++, and it brings a lot of language features. Most of the new additions make the language cleaner and more straightforward.

In this part you'll learn:

- What was removed from the language and what is deprecated.
- How the language is more precise: for example, thanks to expression evaluation order guarantees.
- What are new features of templates: like `if constexpr` or fold expressions
- What are the new standard attributes
- How you can write cleaner and more expressive code thanks to structured binding, inline variables, compile-time if or template argument deduction for classes

Quick Start

To make you more curious about the new Standard, below there are a few code samples that present combined language features.

Don't worry if you find those examples too complicated because they mix too many new things at once. All the new features are individually explained in depth in the next chapters.

Working With Maps

```
1 // Example: Part I/demo_map.cpp
2 #include <iostream>
3 #include <map>
4
5 int main() {
6     std::map<std::string, int> mapUsersAge { { "Alex", 45 }, { "John", 25 } };
7
8     std::map mapCopy{mapUsersAge};
9
10    if (auto [iter, wasAdded] = mapCopy.insert_or_assign("John", 26); !wasAdded)
11        std::cout << iter->first << " reassigned...\n";
12
13    for (const auto& [key, value] : mapCopy)
14        std::cout << key << ", " << value << '\n';
15 }
```

The code will output:

```
John reassigned...
Alex, 45
John, 26
```

The above example uses the following features:

- Line 8: Template Argument Deduction for Class Templates - `mapCopy` type is deduced from the type of `mapUsersAge`. No need to declare `std::map<std::string, int> mapCopy{...}`.
- Line 10: New inserting method for maps - `insert_or_assign`.
- Line 10: Structured Bindings - captures a returned pair from `insert_or_assign` into separate names.
- Line 10: init if statement - `iter` and `wasAdded` are visible only in the scope of the surrounding `if` statement.
- Line 13: Structured Bindings inside a range-based for loop - we can iterate using `key` and `value` rather than `pair.first` and `pair.second`.

Debug Printing

```

1  // Example: Part I/demo_print.cpp
2  #include <iostream>
3
4  template<typename T> void linePrinter(const T& x) {
5      if constexpr (std::is_integral_v<T>)
6          std::cout << "num: " << x << '\n';
7      else if constexpr (std::is_floating_point_v<T>) {
8          const auto frac = x - static_cast<long>(x);
9          std::cout << "flt: " << x << ", frac " << frac << '\n';
10     }
11     else if constexpr (std::is_pointer_v<T>) {
12         std::cout << "ptr, ";
13         linePrinter(*x);
14     }
15     else
16         std::cout << x << '\n';
17 }
18
19 template<typename ... Args> void printWithInfo(Args ... args) {
20     (linePrinter(args), ...); // fold expression over the comma operator
21 }
22
23 int main () {
24     int i = 10;
25     float f = 2.56f;
26     printWithInfo(&i, &f, 30);
27 }

```

The code will output:

```

ptr, num: 10
ptr, flt: 2.56, frac 0.56
num: 30

```

Here you can see the following features:

- Line 5, 7, 11: `if constexpr` - to discard code at compile time, used to match the template parameter.
- Line 5, 7, 11: `_v` variable templates for type traits - no need to write `std::trait_name<T>::value`.
- Line 20: Fold Expressions inside `printWithInfo`. This feature simplifies variadic templates. In the example we invoke `linePrinter()` over all input arguments.

Let's Start!

The code you've seen so far is just the tip of the iceberg! Read the next chapters to see much more: fixes in the language, clarifications, removed things (like `auto_ptr`), and of course the new stuff: `constexpr` `lambda`, `if constexpr`, fold expressions, structured bindings, `template<auto>`, inline variables, template argument deduction for class templates and much more!

1. Fixes and Deprecation

The Standard for C++17 now contains over 1600 pages and has grown over 200 pages compared to C++14¹! Fortunately, the language specification was cleaned up in a few places, and some old or potentially bad features could be cleared out.

In this chapter you'll learn:

- What was removed from the language like: `register` keyword, `auto_ptr` and `std::random_shuffle`.
- What's now deprecated and will be removed in future versions of the language.
- What was fixed, notably the auto type deduction with brace initialisation.

¹For example the draft from July 2017 [N4687](#) compared to C++14 draft [N4140](#)

Removed Things

One of the core concepts behind each iteration of C++ is its compatibility with the previous versions. We'd like to have new things in the language, but at the same time, our old projects should still compile. However, sometimes, there's a chance to remove parts that are rarely used or are wrong.

This section briefly explains what was removed from the Standard.

Removing `auto_ptr`

Probably the best news of all!

C++98 added `auto_ptr` as a way to support basic RAII features for raw pointers. However, due to the lack of move semantics in the language, this smart pointer could be easily misused and cause runtime errors.

Here's an example where `auto_ptr` might cause a crash:

```
// Chapter Fixes And Deprecation/auto_ptrCrash.cpp
```

```
void doSomething(std::auto_ptr<int> myPtr)
{
    *myPtr = 11;
}

void AutoPtrTest() {
    std::auto_ptr<int> myTest(new int(10));
    doSomething(myTest);
    *myTest = 12;
}
```

`doSomething()` takes the pointer by value, but since it's not a shared pointer, it takes the unique ownership of the managed object. Later when the function is done, the copy of the pointer goes out of scope, and the object is deleted.

In `AutoPtrTest()` when `doSomething` is finished the pointer is already cleaned up, and you'll get undefined behavior, maybe even a crash!

In C++11 we got smart pointers: `unique_ptr`, `shared_ptr` and `weak_ptr`. With the move semantics the language could finally support proper unique resource transfers. Also, new smart pointers can be stored in standard containers, which was not possible with `auto_ptr`. You should replace `auto_ptr` with `unique_ptr` as it's the direct and the best equivalent for `auto_ptr`.

New smart pointers are much more powerful and safer than `auto_ptr` so it has been deprecated since C++11. Compilers should report a warning:

```
warning: 'template<class> class std::auto_ptr' is deprecated
```

Now, when you compile with a conformant C++17 compiler, you'll get an error.

Here's the error from MSVC 2017 when using `/std:c++latest`:

```
error C2039: 'auto_ptr': is not a member of 'std'
```

If you need help with the conversion from `auto_ptr` to `unique_ptr` you can check Clang Tidy, as it provides auto conversion: [Clang Tidy: modernize-replace-auto-ptr²](#).



Extra Info

The change was proposed in: [N4190³](#).

Removing the `register` Keyword

The `register` keyword was deprecated in 2011 (C++11), and it has no meaning since that time. Now it's being removed. This keyword is reserved and might be repurposed in the future revisions of the Standard (for example `auto` keyword was reused and now is an entirely new and powerful feature).

If you use `register` to declare a variable:

```
register int a;
```

You might get the following warning (GCC8.1 below)

```
warning: ISO C++17 does not allow 'register' storage class specifier
```

or error in Clang (Clang 7.0)

```
error: ISO C++17 does not allow 'register' storage class specifier
```



Extra Info

The change was proposed in: [P0001R1⁴](#).

²<https://clang.llvm.org/extra/clang-tidy/checks/modernize-replace-auto-ptr.html>

³<https://wg21.link/n4190>

⁴<https://wg21.link/p0001r1>

Removing Deprecated `operator++(bool)`

This operator has been deprecated for a very long time! The committee recommended against its use back in 1998 (C++98), but they only now finally agreed to remove it from the language.

If you try to write the following code:

```
bool b;  
b++;
```

You should get a similar error like this from GCC (GCC 8.1):

```
error: use of an operand of type 'bool' in 'operator++' is forbidden in C++17
```



Extra Info

The change was proposed in: [P0002R1⁵](https://ericniebler.com/2017/05/01/cplusplus-bool-operators/).

Removing Deprecated Exception Specifications

In C++17, exception specification will be part of the type system (as discussed [in the next chapter about Language Clarification](#)). However, the standard contains old and deprecated exception specification that appeared to be impractical and unused.

For example:

```
void fooThrowsInt(int a) throw(int) {  
    printf_s("can throw ints\n");  
    if (a == 0)  
        throw 1;  
}
```

Pay special attention to that `throw(int)` part.

The above code has been deprecated since C++11. The only practical exception declaration is `throw()` which means - this code won't throw anything. Since C++11 it's been advised to use `noexcept`.

For example in clang 4.0 you'll get the following error:

⁵<https://ericniebler.com/2017/05/01/cplusplus-bool-operators/>

error: ISO C++1z does not allow dynamic exception specifications
 [-Wdynamic-exception-spec] note: use 'noexcept(false)' instead



Extra Info

The change was proposed in: [P0003R5](#)⁶.

Other Removed Features

Here's a list of other smaller things removed from the language:

`std::random_shuffle`

The algorithm was marked already as deprecated in C++14. The reason was that in most cases it used the `rand()` function which is considered inefficient and even error-prone (as it uses global state). If you need the same functionality use:

```
template< class RandomIt, class URBG >
void shuffle( RandomIt first, RandomIt last, URBG&& g );
```

`std::shuffle` takes a random number generator as the third argument. More in [N4190](#)⁷.

“Removing Old functional Stuff”

Functions like `bind1st()`/`bind2nd()`/`mem_fun()`, ... were introduced in the C++98-era and are not needed now as you can apply a lambda. What's more, the other functions were not updated to handle perfect forwarding, `decltype` and other modern techniques from C++11. Thus it's best not to use them in modern code. More in [N4190](#)⁸.

Removing Trigraphs

Trigraphs are special character sequences that could be used when a system doesn't support 7-bit ASCII (like ISO 646). For example `??=` generated `#`, `??~` produced `~`. (All of C++'s basic source character set fits in 7-bit ASCII). Today, trigraphs are rarely used, and by removing them from the translation phase, the compilation process can be more straightforward. More in [N4086](#)⁹.

⁶<http://wg21.link/p0003r5>

⁷<https://wg21.link/n4190>

⁸<https://wg21.link/n4190>

⁹<https://wg21.link/n4086>

Fixes

We can argue what is a fix in a language standard and what is not. Below there are three things that might look like a fix for something that was missing or not working in the previous rules.

New auto rules for direct-list-initialisation

Since C++11 there's been a strange problem where:

```
auto x { 1 };
```

Is deduced as `std::initializer_list<int>`. Such behaviour is not intuitive as in most cases you should expect it to work like `int x { 1 };`.

Brace initialisation is the preferred pattern in modern C++, but such exceptions make the feature weaker.

With the new standard, we can fix this so that it will deduce `int`.

To make this happen, we need to understand two ways of initialisation - copy and direct:

```
auto x = foo();    // copy-initialisation
auto x{foo()};     // direct-initialisation, initializes an
                  // initializer_list (until C++17)

int x = foo();     // copy-initialisation
int x{foo()};      // direct-initialisation
```

For the direct initialisation, C++17 introduces new rules:

- For a braced-init-list with only a single element, auto deduction will deduce from that entry;
- For a braced-init-list with more than one element, auto deduction will be ill-formed.

For example:

```
auto x1 = { 1, 2 }; // decltype(x1) is std::initializer_list<int>
auto x2 = { 1, 2.0 }; // error: cannot deduce element type
auto x3{ 1, 2 };      // error: not a single element
auto x4 = { 3 };      // decltype(x4) is std::initializer_list<int>
auto x5{ 3 };         // decltype(x5) is int
```



Extra Info

The change was proposed in: [N3922¹⁰](http://wg21.link/n3922) and [N3681¹¹](http://wg21.link/n3681). The compilers fixed this issue quite early, as the improvement is available in GCC 5.0 (Mid 2015), Clang 3.8 (Early 2016) and MSVC 2015 (Mid 2015). Much earlier than C++17 was approved.

static_assert With no Message

This feature adds a new overload for `static_assert`. It enables you to have the condition inside `static_assert` without passing the message.

It will be compatible with other asserts like `BOOST_STATIC_ASSERT`. Programmers with boost experience will now have no trouble switching to C++17 `static_assert`.

```
static_assert(std::is_arithmetic_v<T>, "T must be arithmetic");
static_assert(std::is_arithmetic_v<T>); // no message needed since C++17
```

In many cases, the condition you check is expressive enough and doesn't need to be mentioned in the message string.



Extra Info

The change was proposed in: [N3928¹²](http://wg21.link/n3928).

Different begin and end Types in Range-Based For Loop

Since C++11 range-based for loop:

```
for (for-range-declaration : for-range-initializer)
{
    statement;
}
```

According to the standard such loop expression was equivalent to the following code:

¹⁰<http://wg21.link/n3922>

¹¹<http://wg21.link/n3681>

¹²<https://wg21.link/n3928>

```

{
    auto && __range = for-range-initializer;
    for ( auto __begin = begin-expr, __end = end-expr;
        __begin != __end;
        ++__begin ) {
        for-range-declaration = *__begin;
        statement
    }
}

```

As you can see, `__begin` and `__end` have the same type. This works nicely but is not scalable enough. For example, you might want to iterate until some sentinel that is a different type than the start of the range.

In C++17 it's changed into:

```

{
    auto && __range = for-range-initializer;
    auto __begin = begin-expr;
    auto __end = end-expr;
    for ( ; __begin != __end; ++__begin ) {
        for-range-declaration = *__begin;
        statement
    }
}

```

Types of `__begin` and `__end` might be different; only the comparison operator is required. Such change has no visible outcome for the users of for loops but enables more options for libraries. For example, this little change allows Range TS (and Ranges in C++20) to work with the range for loop.



Extra Info

The change was proposed in: [P0184R0](https://wg21.link/p0184r0)¹³.

¹³<https://wg21.link/p0184r0>

Compiler support

Feature	GCC	Clang	MSVC
Removing <code>register</code> keyword	7.0	3.8	VS 2017 15.3
Remove Deprecated <code>operator++(bool)</code>	7.0	3.8	VS 2017 15.3
Removing Deprecated Exception Specifications	7.0	4.0	VS 2017 15.5
Removing <code>auto_ptr</code> , <code>random_shuffle</code> , old <code><functional></code>	No ¹⁴	not yet	VS 2015
Removing trigraphs	5.1	3.5	VS 2010
New auto rules for direct-list-initialisation	5.0	3.8	VS 2015
<code>static_assert</code> with no message	6.0	2.5	VS 2017
Different begin and end types in range-based for	6.0	3.6	VS 2017

¹⁴Kept for compatibility.

2. Language Clarification

C++ is a challenging language to learn and fully understand, and some parts might be confusing for programmers. One of the reasons for the lack of clarity might be the free choice for the implementation/compiler. For example, to allow for more aggressive optimisations or for the requirement to be backward (or C) compatible. C++17 reviews some of most popular 'holes' and addresses them.

In this chapter you'll learn:

- What Evaluation Order is and why it might generate unexpected results.
- Copy elision (an optional optimisation that seems to be implemented across all of the popular compilers).
- Exceptions as part of the function declaration.
- Memory allocations for (over)aligned data.

Stricter Expression Evaluation Order

Until C++17 the language hasn't specified any evaluation order for function parameters. **Period.**

For example, that's why in C++14 `make_unique` is not just a syntactic sugar, but it guarantees memory safety:

Let's have a look at the following example:

```
foo(make_unique<T>(), otherFunction());
```

And with explicit new.

```
foo(unique_ptr<T>(new T), otherFunction());
```

In C++14, in the above code, we know that `new T` is guaranteed to happen before `unique_ptr` construction, but that's all. For example, `new T` might happen first, then `otherFunction()`, and then `unique_ptr` constructor.

When `otherFunction()` throws, then `new T` generates a leak (as the unique pointer is not yet created). When you use `make_unique`, then it's not possible to leak, even when the order of execution is unknown.

C++17 addresses this issue, and now the evaluation order of attributes is "practical" and predictable.

Examples

In an expression:

```
f(a, b, c);
```

The order of evaluation of `a`, `b`, `c` is still unspecified, but any parameter is fully evaluated before the next one is started. It's especially crucial for complex expressions like:

```
f(a(x), b, c(y));
```

When the compiler chooses to evaluate the first argument - `a(x)` - then it also needs to evaluate `x` before processing `b` or `c(y)`.

This fixes a problem with `make_unique` vs `unique_ptr<T>(new T())` - as the function argument must be fully evaluated before other arguments are started.

Consider the following case:

```
#include <iostream>

class Query
{
public:
    Query& addInt(int i)
    {
        std::cout << "addInt: " << i << '\n';
        return *this;
    }

    Query& addFloat(float f)
    {
        std::cout << "addFloat: " << f << '\n';
        return *this;
    }
};

float computeFloat()
{
    std::cout << "computing float... \n";
    return 10.1f;
}

float computeInt()
{
    std::cout << "computing int... \n";
    return 8;
}

int main()
{
    Query q;
    q.addFloat(computeFloat()).addInt(computeInt());
}
```

You probably expect that using C++14 `computeInt()` happens after `addFloat`. Unfortunately, that might not be the case. For instance here's an output from GCC 4.7.3:

```

computing int...
computing float...
addFloat: 10.1
addInt: 8

```

The chaining of functions is already specified to work from left to right, but the order of evaluation of inner expressions might be different. To be precise:

The expressions are indeterminately sequenced with respect to each other.

Now, with C++17, function chaining will work as expected when they contain inner expressions, i.e., they are evaluated from left to right:

In an expression:

```
`a(expA).b(expB).c(expC)`
```

expA is evaluated before calling b.

The above example, compiled using a conformant C++17 compiler yields the following result:

```

computing float...
addFloat: 10.1
computing int...
addInt: 8

```

Another result of this change is that when using operator overloading, the order of evaluation is determined by the order associated with the corresponding built-in operator.

That's why `std::cout << a() << b() << c()` is evaluated as a, b, c. Before C++17, it could have been in any order.

Here are more rules described in the standard:

the following expressions are evaluated in the order a, then b:

1. `a.b`
2. `a->b`
3. `a->*b`
4. `a(b1, b2, b3)` // b1, b2, b3 - in any order
5. `b @= a` // '@' means any operator
6. `a[b]`
7. `a << b`
8. `a >> b`

If you're not sure how your code might be evaluated, then it's better to make it simple and split it into several clear statements. You can find some guides in the Core C++ Guidelines for example [ES.44](#)¹ and [ES.44](#)².

**Extra Info**

The change was proposed in: [P0145R3](#)³.

¹<http://isocpp.github.io/CppCoreGuidelines/CppCoreGuidelines#Res-order>

²<http://isocpp.github.io/CppCoreGuidelines/CppCoreGuidelines#es44-dont-depend-on-order-of-evaluation-of-function-arguments>

³<https://wg21.link/p0145r3>

Guaranteed Copy Elision

Copy Elision is a popular optimisation that avoids creating unnecessary temporary objects.

For example:

```
// Chapter Clarification/copy_elision.cpp
#include <iostream>

struct Test
{
    Test() { std::cout << "Test::Test\n"; }
    Test(const Test&) { std::cout << "Test(const Test&)\n"; }
    Test(Test&&) { std::cout << "Test(Test&&)\n"; }
    ~Test() { std::cout << "~Test\n"; }
};

Test Create()
{
    return Test();
}

int main()
{
    auto n = Create();
}
```

In the above call, you might assume a temporary copy is used - to store the return value of `Create`. In C++14, most compilers can see that the temporary object can be easily optimised and they can create `n` “directly” from the call of `Create()`. So you’ll probably see the following output:

```
Test::Test // create n
~Test // destroy n when main finishes
```

In its basic form, such optimisation - copy elision - is called Return Value Optimisation - **RVO**.

As an experiment, in GCC you can add a compiler flag `-fno-elide-constructors` and use `-std=c++14` (or some earlier language standard). In that case you’ll see a different output:

```
// compiled as "g++ CopyElision.cpp -std=c++14 -fno-elide-constructors"
Test::Test
Test(Test&&)
~Test
Test(Test&&)
~Test
~Test
```

In this case, we have two extra copies that the compiler uses to pass the return value into `n`;

Compilers are even smarter, and they can elide in cases when you return a named object - it's called **Named Return Value Optimisation - NRVO**:

```
Test Create()
{
    Test t;
    // several instruction to initialize 't'...
    return t;
}

auto n = Create(); // temporary will be usually elided
```

Currently, the standard allows eliding in cases like:

- when a temporary object is used to initialise another object (including the object returned by a function, or the exception object created by a throw-expression)
- when a variable that is about to go out of scope is returned or thrown
- when an exception is caught by value

However, it's up to the compiler/implementation to elide or not. In practice, all the constructors' definitions are required. Sometimes elision might happen only in release builds (optimised), while Debug builds (without any optimisation) won't elide anything.

With C++17 we get clear rules on when elision has to happen, and thus constructors might be entirely omitted.

Why might this be useful?

- to allow returning objects that are not movable/copyable - because we could now skip copy/move constructors
- to improve code portability - as every conformant compiler supports the same rule
- to support the 'return by value' pattern rather than using output arguments
- to improve performance

Below you can see another example with non-movable/non-copyable type:

```
// Chapter Clarification/copy_elision_non_moveable.cpp
#include <array>

// based on P0135R0
struct NonMoveable
{
    NonMoveable(int x) : v(x) { }
    NonMoveable(const NonMoveable&) = delete;
    NonMoveable(NonMoveable&&) = delete;

    std::array<int, 1024> arr;
    int v;
};

NonMoveable make(int val)
{
    if (val > 0)
        return NonMoveable(val);

    return NonMoveable(-val);
}

int main()
{
    auto largeNonMoveableObj = make(90); // construct the object
    return largeNonMoveableObj.v;
}
```

The above code wouldn't compile under C++14 as it lacks copy and move constructors. But with C++17 the constructors are not required - because the object `largeNonMovableObj` will be constructed in place.

Please notice that you can also use many return statements in one function and copy elision will still work.

Moreover, it's important to remember, that in C++17 copy elision works only for temporary objects, not for Named RVO.

But how is mandatory copy elision defined in the standard? The functionality is based on value categories, so read on to the next section to understand how it works.

Updated Value Categories

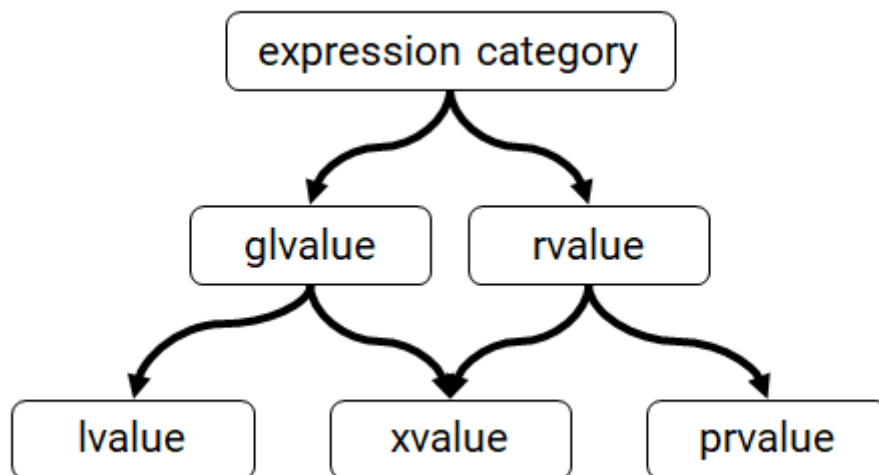
In C++98/03 we had two basic categories of expressions:

- `lvalue`
- `rvalue`

Since C++11 this taxonomy has been extended (because of the move semantics), and now we have five classes:

- `lvalue`
- `glvalue`
- `xvalue`
- `rvalue`
- `prvalue`

Here's a diagram that gives a better overview of the categories:



Value Categories

Remember that we have three core categories (below with colloquial “definitions”):

- `lvalue` - an expression that has an identity, and which we can take the address of
- `xvalue` - “eXpiring lvalue” - an object that we can move from, which we can reuse. Usually, its lifetime ends soon
- `prvalue` - pure rvalue - something without a name, which we cannot take the address of, we can move from such expression

To support Copy Elision in the Standard, the authors of the proposal suggested simplification of `glvalue` and `prvalue`:

- `glvalue` - “generalised” `lvalue` - A `glvalue` is an expression whose evaluation computes the location of an object, bit-field, or function
- `prvalue` - “pure” `rvalue` - A `prvalue` is an expression whose evaluation initialises an object, bit-field, or operand of an operator, as specified by the context in which it appears

For example:

```
class X { int a; };
X{10}    // this expression is prvalue
X x;     // x is lvalue
x.a      // it's lvalue (location)
```

In short: `prvalues` perform initialisation, `glvalues` describe locations.

The C++17 Standard specifies that when there's a `prvalue` of some class or array there's no need to create any temporary copies, the initialisation from that `prvalue` can happen directly. There's no move or copy involved (so there's no need to have a copy and move constructors); the compiler can safely do the elision.

It can happen:

- in initialisation of an object from a `prvalue`: `Type t = T()`
- in a function call where the function returns a `prvalue` - like in our examples.

There are several exceptions where the temporary is still needed:

- when a `prvalue` is bound to a reference
- when member access is performed on a class `prvalue`
- when array subscripting is performed on an array `prvalue`
- when an array `prvalue` is decayed to a pointer
- when a derived-to-base conversions performed on a class `prvalue`
- when a `prvalue` is used as a discarded value expression



Extra Info

The change was proposed in: [P0135R0](https://wg21.link/p0135r0)⁴(reasoning) - and [P0135R1](https://wg21.link/p0135r1)⁵(wording).

⁴<https://wg21.link/p0135r0>

⁵<https://wg21.link/p0135r1>

Dynamic Memory Allocation for Over-Aligned Data

When you work with SIMD (Single instruction, multiple data) or when you have some other memory layout requirements, you might need to align objects specifically. For example, in SSE (Streaming SIMD Extensions) you need the 16-byte alignment (for AVX 256 there's 32-byte alignment requirement). You could define a `vec4` like:

```
class alignas(16) vec4
{
    float x, y, z, w;
};
auto pVectors = new vec4[1000];
```

Note: the `alignas`⁶ specifier is available since C++11.

However, in C++11/14 you have no guarantee how the memory will be aligned. Often you have to use some special routines like `_aligned_malloc/_aligned_free` to be sure the alignment is preserved. That's not ideal as it's not working with C++ smart pointers and also makes memory management visible in the code⁷.

C++17 fixes that hole by introducing additional memory allocation functions that use align parameters:

```
void* operator new(size_t, align_val_t);
void* operator new[](size_t, align_val_t);
void operator delete(void*, align_val_t);
void operator delete[](void*, align_val_t);
void operator delete(void*, size_t, align_val_t);
void operator delete[](void*, size_t, align_val_t);
```

now, you can allocate that `vec4` array as:

```
auto pVectors = new vec4[1000];
```

No code changes, but now the alignment of `vec4` is properly handled:

```
operator new[](sizeof(vec4), align_val_t(alignof(vec4)))
```

In other words, `new` is now aware of the alignment of the object.



Extra Info

The change was proposed in: [P0035R4](#)⁸.

⁶<http://en.cppreference.com/w/cpp/language/alignas>

⁷According to Core Guidelines we should stop using raw `new` and `delete` explicitly in our code.

⁸<http://wg21.link/p0035r4>

Exception Specifications as Part of the Type System

Exception Specification for a function didn't use to belong to the type of the function, but now it will be part of it.

You can now have two function overloads: one with `noexcept` and one without it.

For example, you'll get an error in the case:

```
void (*p)();
void (**pp)() noexcept = &p; // error: cannot convert to
                             // pointer to noexcept function

struct S { typedef void (*p)(); operator p(); };
void (*q)() noexcept = S(); // error: cannot convert to
                             // pointer to noexcept
```

One of the reasons for adding the feature is the possibility of allowing for better optimisation. That can happen when you have a guarantee that a function is `noexcept` and that it doesn't throw any exceptions.

Also, as described in the previous chapter [about Language Fixes](#), in C++17 Exception Specification is cleaned up by removing deprecated exception specifications. Effectively, you can only use `noexcept specifier`⁹ for declaring that a function might throw something or not.



Extra Info

The change was proposed in: [P0012R1](#)¹⁰.

Compiler Support

Feature	GCC	Clang	MSVC
Stricter expression evaluation order	7.0	4.0	VS 2017
Guaranteed copy elision	7.0	4.0	VS 2017 15.6
Dynamic memory allocation for over-aligned data	7.0	4.0	VS 2017 15.5
Exception specifications part of the type system	7.0	4.0	VS 2017 15.5

⁹http://en.cppreference.com/w/cpp/language/noexcept_spec

¹⁰<http://wg21.link/p0012r1>

3. General Language Features

Having finished the chapters on language fixes and clarifications, we're now ready to look at wide-spread features. Improvements described in this section also have the potential to make your code more compact and expressive.

For example, with Structured bindings, you can leverage much easier syntax tuples (and tuple-like expressions). Something that was easy in other languages like Python is also possible with good-old C++!

In this chapter you'll learn:

- Structured bindings/Decomposition declarations
- How to provide Structured Binding interface for your custom classes
- Init-statement for if/switch
- Inline variables and their impact on header-only libraries
- Lambda expressions that might be used in a `constexpr` context
- Simplified use of nested namespaces

Structured Binding Declarations

Do you often work with tuples or pairs?

If not, then you should probably start looking into those handy types. Tuples enable you to bundle data ad-hoc with nice library support instead of creating your own types for everything or using output parameters. And language support like structured binding makes them nice to handle.

Consider a function that returns two results in a pair:

```
std::pair<int, bool> InsertElement(int el) { ... }
```

You can use `auto ret = InsertElement(...)` and then refer to `ret.first` or `ret.second`. Alternatively you can leverage `std::tie` which will unpack the tuple/pair into custom variables:

```
int index { 0 };  
bool flag { false };  
std::tie(index, flag) = InsertElement(10);
```

Such code might be useful when you work with `std::set::insert` which returns `std::pair`:

```
std::set<int> mySet;  
std::set<int>::iterator iter;  
bool inserted;  
  
std::tie(iter, inserted) = mySet.insert(10);  
  
if (inserted)  
    std::cout << "Value was inserted\n";
```

With C++17 the code can be more compact:

```
std::set<int> mySet;  
  
auto [iter, inserted] = mySet.insert(10);
```

Now, instead of `pair.first` and `pair.second`, you can use variables with concrete names. In addition you have one line instead of three and the code is easier to read. The code is also safer as `iter` and `inserted` are initialised in the expression.

Such syntax is called structured binding expression.

The Syntax

The basic syntax is as follows:

```
auto [a, b, c, ...] = expression;  
auto [a, b, c, ...] { expression };  
auto [a, b, c, ...] ( expression );
```

The compiler introduces all identifiers from the `a, b, c, ...` list as names in the surrounding scope and binds them to sub-objects or elements of the object denoted by `expression`.

Behind the scenes, the compiler might generate the following **pseudo code**:

```
auto tempTuple = expression;  
using a = tempTuple.first;  
using b = tempTuple.second;  
using c = tempTuple.third;
```

Conceptually, the expression is copied into a tuple-like object (`tempTuple`) with member variables that are exposed through `a`, `b` and `c`. However, the variables `a`, `b` and `c` are not references, they are aliases (or bindings) to the hidden object member variables. The temporary object has a unique name assigned by the compiler.

For example:

```
std::pair a(0, 1.0f);  
auto [x, y] = a;
```

`x` binds to `int` stored in the hidden object that is a copy of `a`. And similarly, `y` binds to `float`.

Modifiers

Several modifiers can be used with structured bindings:

`const` modifiers:

```
const auto [a, b, c, ...] = expression;
```

References:

```
auto& [a, b, c, ...] = expression;  
auto&& [a, b, c, ...] = expression;
```

For example:

```
std::pair a(0, 1.0f);
auto& [x, y] = a;
x = 10; // write access
// a.first is now 10
```

In the above example `x` binds to the element in the hidden object that is a reference to `a`.

Now it's also quite easy to get a reference to a tuple member:

```
auto& [ refA, refB, refC, refD ] = myTuple;
```

You can also add `[[attribute]]`:

```
[[maybe_unused]] auto& [a, b, c, ...] = expression;
```



Structured Bindings or Decomposition Declaration?

For this feature, you might have seen another name “decomposition declaration” in use. During the standardisation process those two names were considered, but now the final version sticks with “Structured Bindings.”

Structured Binding Declaration cannot be declared `constexpr`

There's one limitation that is worth remembering.

```
constexpr auto [x, y] = std::pair(0, 0);
```

This generates error:

```
error: structured binding declaration cannot be 'constexpr'
```

This limitation might be solved in C++20 by the proposal [P1481](https://wg21.link/P1481)¹

¹<https://wg21.link/P1481>

Binding

Structured Binding is not only limited to tuples, we have three cases from which we can bind from:

1. If the initializer is an array:

```
// works with arrays:
double myArray[3] = { 1.0, 2.0, 3.0 };
auto [a, b, c] = myArray;
```

In this case an array is copied into a temporary object and a, b and c refers to copied elements from the array.

The number of identifier must match the number of elements in the array.

2. If the initializer supports `std::tuple_size<>` and provides `get<N>()` and `std::tuple_element` functions:

```
std::pair myPair(0, 1.0f);
auto [a, b] = myPair; // binds myPair.first/second
```

In the above snippet we bind to `myPair`. But this also means that you can provide support for your classes, assuming you add `get<N>` interface implementation. See an example in the later section.

3. If the initializer's type contains only non static, public members:

```
struct Point {
    double x;
    double y;
};

Point GetStartPoint() {
    return { 0.0, 0.0 };
}

const auto [x, y] = GetStartPoint();
```

x and y refer to `Point::a` and `Point::b` from the `Point` structure.

The class doesn't have to be POD, but the number of identifiers must equal the number of non-static data members.



Note: In C++17, you could use structured bindings to bind to class members as long as they were public. This might be a problem when you want to access private members in the implementation of the class. With C++20 it should be fixed. See [P0969R0](https://wg21.link/P0969R0)².

²<https://wg21.link/P0969R0>

Examples

One of the **coolest use cases** - binding inside a range based for loop:

```
std::map<KeyType, ValueType> myMap;
for (const auto & [key, val] : myMap)
{
    // use key/value rather than iter.first/iter.second
}
```

In the above example, we bind to a pair of [key, val] so we can use those names in the loop. Before C++17 you had to operate on an iterator from the map - which is a pair <first, second>. Using the real names key/value is more expressive.

The above technique can be used in:

```
#include <map>
#include <iostream>
#include <string>

int main()
{
    const std::map<std::string, int> mapCityPopulation {
        { "Beijing", 21'707'000 },
        { "London", 8'787'892 },
        { "New York", 8'622'698 }
    };

    for (auto&[city, population] : mapCityPopulation)
        std::cout << city << ": " << population << '\n';
}
```

In the loop body, you can safely use city and population variables.

Providing Structured Binding Interface for Custom Class

As mentioned earlier you can provide Structured Binding support for a custom class.

To do that you have to define `get<N>`, `std::tuple_size` and `std::tuple_element` specialisations for your type.

For example, if you have a class with three members, but you'd like to expose only its public interface:

```

class UserEntry {
public:
    void Load() { }

    std::string GetName() const { return name; }
    unsigned GetAge() const { return age; }
private:
    std::string name;
    unsigned age { 0 };
    size_t cacheEntry { 0 }; // not exposed
};

```

The interface for Structured Bindings:

```

// with if constexpr:
template <size_t I> auto get(const UserEntry& u) {
    if constexpr (I == 0) return u.GetName();
    else if constexpr (I == 1) return u.GetAge();
}

namespace std {
    template <> struct tuple_size<UserEntry> : std::integral_constant<size_t, 2> { };

    template <> struct tuple_element<0,UserEntry> { using type = std::string; };
    template <> struct tuple_element<1,UserEntry> { using type = unsigned; };
}

```

`tuple_size` specifies how many fields are available, `tuple_element` defines the type for a specific element and `get<N>` returns the values.

Alternatively, you can also use explicit `get<>` specialisations rather than `if constexpr`:

```

template<> string get<0>(const UserEntry &u) { return u.GetName(); }
template<> unsigned get<1>(const UserEntry &u) { return u.GetAge(); }

```

For a lot of types, writing two (or several) functions might be simpler than using `if constexpr`.

Now you can use `UserEntry` in a structured binding, for example:

```
UserEntry u;  
u.Load();  
auto [name, age] = u; // read access  
std::cout << name << ", " << age << '\n';
```

This example only allows read access of the class. If you want write access, then the class should also provide accessors that return references to members. Later you have to implement `get` with references support.

As you’ve seen `if constexpr` was used to implement `get<N>` functions, read more in the [if constexpr](#) chapter.



Extra Info

The change was proposed in: [P0217](#)³(wording), [P0144](#)⁴(reasoning and examples), [P0615](#)⁵(renaming “decomposition declaration” with “structured binding declaration”).

³<https://wg21.link/p0217>

⁴<https://wg21.link/p0144>

⁵<https://wg21.link/p0615>

Init Statement for `if` and `switch`

C++17 provides new versions of the `if` and `switch` statements:

`if (init; condition)` and `switch (init; condition)`.

In the `init` section you can specify a new variable and then check it in the `condition` section. The variable is visible only in `if/else` scope.

To achieve a similar result, before C++17 you had to write:

```
{
    auto val = GetValue();
    if (condition(val))
        // on success
    else
        // on false...
}
```

Please notice that `val` has a separate scope, without that it ‘leaks’ to enclosing scope.

Now, in C++17, you can write:

```
if (auto val = GetValue(); condition(val))
    // on success
else
    // on false...
```

Now, `val` is visible only inside the `if` and `else` statements, so it doesn’t ‘leak.’ `condition` might be any boolean condition.

Why is this useful?

Let’s say you want to search for a few things in a string:

```
const std::string myString = "My Hello World Wow";

const auto pos = myString.find("Hello");
if (pos != std::string::npos)
    std::cout << pos << " Hello\n"

const auto pos2 = myString.find("World");
if (pos2 != std::string::npos)
    std::cout << pos2 << " World\n"
```

You have to use different names for `pos` or enclose it with a separate scope:

```
{
    const auto pos = myString.find("Hello");
    if (pos != std::string::npos)
        std::cout << pos << " Hello\n"
}

{
    const auto pos = myString.find("World");
    if (pos != std::string::npos)
        std::cout << pos << " World\n"
}
```

The new `if` statement will make that additional scope in one line:

```
if (const auto pos = myString.find("Hello"); pos != std::string::npos)
    std::cout << pos << " Hello\n";

if (const auto pos = myString.find("World"); pos != std::string::npos)
    std::cout << pos << " World\n";
```

As mentioned before, the variable defined in the `if` statement is also visible in the `else` block. So you can write:

```
if (const auto pos = myString.find("World"); pos != std::string::npos)
    std::cout << pos << " World\n";
else
    std::cout << pos << " not found!!\n";
```

Plus, you can use it with structured bindings (following Herb Sutter code⁶):

```
// better together: structured bindings + if initializer
if (auto [iter, succeeded] = mymap.insert(value); succeeded) {
    use(iter); // ok
    // ...
} // iter and succeeded are destroyed here
```



Extra Info

The change was proposed in: [P0305R1](http://wg21.link/p0305r1)⁷.

⁶<https://herbsutter.com/2016/06/30/trip-report-summer-iso-c-standards-meeting-oulu/>

⁷<http://wg21.link/p0305r1>

Inline Variables

With Non-Static Data Member Initialisation introduced in C++11, you can now declare and initialise member variables in one place:

```
class User
{
    int _age {0};
    std::string _name {"unknown"};
};
```

Still, with static variables (or `const static`) you usually need to define it in some `cpp` file.

C++11 and `constexpr` keyword allow you to declare and define static variables in one place, but it's limited to constant expressions only.

Previously, only methods/functions could be specified as `inline`, but now you can do the same with variables, inside a header file.

From the proposal P0386R2⁸: A variable declared inline has the same semantics as a function declared inline: it can be defined, identically, in multiple translation units, must be defined in every translation unit in which it is used, and the behaviour of the program is as if there was exactly one variable.

For example:

```
// inside a header file:
struct MyClass
{
    static const int sValue;
};

// later in the same header file:
inline int const MyClass::sValue = 777;
```

Or even declaration and definition in one place:

⁸<http://wg21.link/p0386r2>

```
struct MyClass
{
    inline static const int sValue = 777;
};
```

Also, note that constexpr variables are inline implicitly, so there's no need to use constexpr inline myVar = 10;.

An inline variable is also more flexible than a constexpr variable as it doesn't have to be initialised with a constant expression. For example, you can initialise an inline variable with rand(), but it's not possible to do the same with constexpr variable.

How Can it Simplify the Code?

A lot of header-only libraries can limit the number of hacks (like using inline functions or templates) and switch to using inline variables.

For example:

```
class MyClass
{
    static inline int Seed(); // static method
};

inline int MyClass::Seed() {
    static const int seed = rand();
    return seed;
}
```

Can be changed into:

```
class MyClass
{
    static inline int seed = rand();
};
```

C++17 guarantees that MyClass::seed will have the same value (generated at runtime) across all the compilation units!



Extra Info

The change was proposed in: [P0386R2](http://wg21.link/p0386r2)⁹.

⁹<http://wg21.link/p0386r2>

constexpr Lambda Expressions

Lambda expressions were introduced in C++11, and since that moment they've become an essential part of modern C++. Another significant feature of C++11 is "constant expressions" - declared mainly with `constexpr`. In C++17 the two elements are allowed to exist together - so your lambda can be invoked in a constant expression context.

In C++11/14 the following code wouldn't compile:

```
auto SimpleLambda = [] (int n) { return n; };  
static_assert(SimpleLambda(3) == 3, "");
```

GCC compiled with the `-std=c++11` flag reports the following error:

```
error: call to non-constexpr function 'main()::<lambda(int)>'  
    static_assert(SimpleLambda(3) == 3, "");
```

However, with the `-std=c++17` the code compiles! This is because since C++17 lambda expressions that follow the rules of standard `constexpr` functions are implicitly declared as `constexpr`.

What does that mean? What are the limitations?

Quoting the Standard - 10.1.5 *The constexpr specifier* [dcl.constexpr]

The definition of a `constexpr` function shall satisfy the following requirements:

- it shall not be virtual (13.3);
- its return type shall be a literal type; A >- each of its parameter types shall be a literal type; A >- its function-body shall be = delete, = default, or a compound-statement that does not contain:
 - an asm-definition,
 - a goto statement,
 - an identifier label (9.1),
 - a try-block, or
 - a definition of a variable of non-literal type or of static or thread storage duration or for which no initialisation is performed.

Practically, if you want your function or lambda to be executed at compile-time then the body of this function shouldn't invoke any code that is not `constexpr`. For example, you cannot allocate memory dynamically.

Lambda can be also explicitly declared `constexpr`:

```
auto SimpleLambda = [] (int n) constexpr { return n; };
```

And if you violate the rules of constexpr functions, you'll get a compile-time error.

```
auto FaultyLeakyLambda = [] (int n) constexpr {  
    int *p = new int(10);  
  
    return n + (*p);  
};
```

operator new is not constexpr so that won't compile.

Having constexpr lambdas will be a great feature combined with the constexpr standard algorithms that are coming in C++20.



Extra Info

The change was proposed in: [P0170¹⁰](http://wg21.link/p0170).

¹⁰<http://wg21.link/p0170>

Nested Namespaces

Namespaces allow grouping types and functions into separate logical units.

For example, it's common to see that each type or function from a library XY will be stored in a namespace xy. Like in the below case, where there's `SuperCompressionLib` and it exposes functions called `Compress()` and `Decompress()`:

```
namespace SuperCompressionLib {  
    bool Compress();  
    bool Decompress();  
}
```

Things get interesting if you have two or more nested namespaces.

```
namespace MySuperCompany {  
    namespace SecretProject {  
        namespace SafetySystem {  
            class SuperArmor {  
                // ...  
            };  
            class SuperShield {  
                // ...  
            };  
        } // SafetySystem  
    } // SecretProject  
} // MySuperCompany
```

With C++17 nested namespaces can be written in a more compact way:

```
namespace MySuperCompany::SecretProject::SafetySystem {  
    class SuperArmor {  
        // ...  
    };  
    class SuperShield {  
        // ...  
    };  
}
```

Such syntax is comfortable, and it will be easier to use for developers that have experience in languages like C# or Java.

In C++17 also the Standard Library was “compacted” in several places by using the new nested namespace feature:

For example for `regex`.

In C++17 it’s defined as:

```
namespace std::regex_constants {
    typedef T1 syntax_option_type;
    // ...
}
```

Before C++17 the same was declared as:

```
namespace std {
    namespace regex_constants {
        typedef T1 syntax_option_type;
        // ...
    }
}
```

The above nested declarations appear in the C++ Specification, but it might look different in an STL implementation.



Extra Info

The change was proposed in: [N4230¹¹](http://ericniebler.com/2016/05/24/n4230/).

Compiler support

Feature	GCC	Clang	MSVC
Structured Binding Declarations	7.0	4.0	VS 2017 15.3
Init-statement for if/switch	7.0	3.9	VS 2017 15.3
Inline variables	7.0	3.9	VS 2017 15.5
constexpr Lambda Expressions	7.0	5.0	VS 2017 15.3
Nested namespaces	6.0	3.6	VS 2015

¹¹<http://wg21.link/N4230>

4. Templates

Do you work with templates and/or meta-programming?

If your answer is “YES,” then you might be very happy with the updates from C++17.

The new standard introduces many enhancements that make template programming much easier and more expressive.

In this chapter you’ll learn:

- Template argument deduction for class templates
- `template<auto>`
- Fold expressions
- `if constexpr` - the compile-time if for C++!
- Plus some smaller, detailed improvements and fixes

Template Argument Deduction for Class Templates

Do you often use `make_Type` functions to construct a templated object (like `std::make_pair`)?

With C++17 you can forget about (most of) them and just use a regular constructor. C++17 filled a gap in the deduction rules for templates. Now template deduction can occur for standard class templates and not just for functions. That also means that a lot of your code - those `make_Type` functions can now be removed.

For instance, to create a pair it was usually more convenient to write:

```
auto myPair = std::make_pair(42, "hello world");
```

Rather than:

```
std::pair<int, std::string> myPair(42, "hello world");
```

Because `std::make_pair` is a template function, the compiler can perform the deduction of function template arguments and there's no need to write:

```
auto myPair = std::make_pair<int, std::string>(42, "hello world");
```

Now, since C++17, the conformant compiler will nicely deduce the template parameter types for class templates too!

In our example, you can now write:

```
using namespace std::string_literals;  
std::pair myPair(42, "hello world"s); // deduced automatically!
```

This can substantially reduce complex constructions like:

```
// lock guard:  
std::shared_timed_mutex mut;  
std::lock_guard<std::shared_timed_mutex> lck(mut);  
  
// array:  
std::array<int, 3> arr {1, 2, 3};
```

Can now become:

```
std::shared_timed_mutex mut;
std::lock_guard lck(mut);

std::array arr { 1, 2, 3 };
```

Note, that partial deduction cannot happen, you have to specify all the template parameters or none:

```
std::tuple t(1, 2, 3); // OK: deduction
std::tuple<int,int,int> t(1, 2, 3); // OK: all arguments are provided
std::tuple<int> t(1, 2, 3); // Error: partial deduction
```

With this feature, a lot of `make_Type` functions might not be needed - especially those that ‘emulate’ template deduction for classes.

Still, there are factory functions that do additional work. For example `std::make_shared` - it not only creates `shared_ptr`, but also makes sure the control block, and the pointed object are allocated in one memory region:

```
// control block and int might be in different places in memory
std::shared_ptr<int> p(new int{10});

// the control block and int are in the same contiguous memory section
auto p2 = std::make_shared<int>(10);
```

How does template argument deduction for classes work?

Let’s enter the “Deduction Guides” area.

Deduction Guides

The compiler uses special rules called “Deduction Guides” to work out the template class types.

We have two types of rules: compiler generated (implicitly generated) and user-defined.

To understand how the compiler uses the guides, let’s look at an example.

Here’s a custom deduction guide for `std::array`:

```
template <class T, class... U>
array(T, U...) -> array<T, 1 + sizeof...(U)>;
```

The syntax looks like a template function with a trailing return type. The compiler treats such “imaginary” function as a candidate for the parameters. If the pattern matches then the proper type is returned from the deduction.

In our case when you write:

```
std::array arr {1, 2, 3, 4};
```

Then, assuming `T` and `U...` are of the same type, we can build up an array object of `std::array<int, 4>`.

In most cases, you can rely on the compiler to generate automatic deduction guides. They will be created for each constructor (also copy/move) of the primary class template. Please note, that classes that are specialised or partially specialised won't work here.

As mentioned, you might also write your own deduction guides:

A classic example of where you might add your custom deduction guides is a deduction of `std::string` rather than `const char*`:

```
template<typename T>
struct MyType
{
    T str;
};

// custom deduction guide
MyType(const char *) -> MyType<std::string>;

MyType t{"Hello World"};
```

Without the custom deduction `T` would be deduced as `const char*`.

Another example of custom deduction guide is `overload`:

```
template<class... Ts>
struct overload : Ts... { using Ts::operator()...; };

template<class... Ts>
overload(Ts...) -> overload<Ts...>; // deduction guide
```

The `overload` class inherits from other classes `Ts...` and then exposes their `operator()`. The custom deduction guide is used here to “transform” a list of lambdas into the list of classes that we can derive from.

**Extra Info**

The change was proposed in: [P0091R3](http://wg21.link/p0091r3)¹ and [P0433 - Deduction Guides in the Standard Library](https://wg21.link/p0433)².

Please note that while a compiler might declare full support for Template Argument Deduction for Class Templates, its corresponding STL implementation might still lack of custom deduction guides for some STL types. See the Compiler Support section at the end of the chapter.

¹<http://wg21.link/p0091r3>

²<https://wg21.link/p0433>

Fold Expressions

C++11 introduced variadic templates which is a powerful feature, especially if you want to work with a variable number of input template parameters to a function. For example, previously (pre C++11) you had to write several different versions of a template function (one for one parameter, another for two parameters, another for three params...).

Still, variadic templates required some additional code when you wanted to implement ‘recursive’ functions like `sum`, `all`. You had to specify rules for the recursion.

For example:

```
auto SumCpp11(){
    return 0;
}

template<typename T1, typename... T>
auto SumCpp11(T1 s, T... ts){
    return s + SumCpp11(ts...);
}
```

And with C++17 we can write much simpler code:

```
template<typename ...Args> auto sum(Args ...args)
{
    return (args + ... + 0);
}

// or even:
template<typename ...Args> auto sum2(Args ...args)
{
    return (args + ...);
}
```

The following variations of [fold expressions](https://en.cppreference.com/w/cpp/language/fold)³ with binary operators (op) exist:

Expression	Name	Expansion
<code>(... op e)</code>	unary left fold	<code>((e1 op e2) op ...) op eN</code>
<code>(init op ... op e)</code>	binary left fold	<code>((init op e1) op e2) op ... op eN</code>
<code>(e op ...)</code>	unary right fold	<code>e1 op (... op (eN-1 op eN))</code>
<code>(e op ... op init)</code>	binary right fold	<code>e1 op (... op (eN-1 op (eN op init)))</code>

³<https://en.cppreference.com/w/cpp/language/fold>

op is any of the following 32 binary operators: + - * / % ^ & | = < > << >> += -= *= /= %= ^= &= |= <<= >>= == != <= >= && || , .* ->*. In a binary fold, both ops must be the same.

For example, when you write:

```
template<typename ...Args> auto sum2(Args ...args)
{
    return (args + ...); // unary right fold over '+'
}
```

```
auto value = sum2(1, 2, 3, 4);
```

The template function is expanded into:

```
auto value = 1 + (2 + (3 + 4));
```

Also by default we get the following values for empty parameter packs:

Operator	default value
&&	true
	false
,	void()
any other	ill-formed code

That's why you cannot call `sum2()` without any parameters, as the unary fold over operator + doesn't have any default value for the empty parameter list.

More Examples

Here's a quite nice implementation of a `printf` using folds [P0036R0](http://ericniebler.com/2015/05/20/folds/)⁴:

```
template<typename ...Args>
void FoldPrint(Args&&... args)
{
    (std::cout << ... << std::forward<Args>(args)) << '\n';
}
```

```
FoldPrint("hello", 10, 20, 30);
```

However, the above `FoldPrint` will print arguments one by one, without any separator. So for the above call, you'll see "hello102030" on the output.

If you want separators and more formatting options you have to alter the printing technique and use fold over comma:

⁴<http://wg21.link/p0036r0>

```

template<typename ...Args>
void FoldSeparateLine(Args&&... args)
{
    auto separateLine = [](const auto& v) {
        std::cout << v << '\n';
    };
    (... , separateLine (std::forward<Args>(args))); // over comma operator
}

```

The technique with fold over the comma operator is handy. Another example of it might be a special version of `push_back`:

```

template<typename T, typename... Args>
void push_back_vec(std::vector<T>& v, Args&&... args)
{
    (v.push_back(std::forward<Args>(args)), ...);
}

```

```

std::vector<float> vf;
push_back_vec(vf, 10.5f, 0.7f, 1.1f, 0.89f);

```

In general, fold expression allows you to write cleaner, shorter and probably more comfortable to read the code.



Extra Info

The change was proposed in: [N4295](https://wg21.link/n4295)⁵ and [P0036R0](https://wg21.link/p0036r0)⁶.

⁵<https://wg21.link/n4295>

⁶<https://wg21.link/p0036r0>

if constexpr

This is a big one!

The compile-time if for C++!

The feature allows you to discard branches of an if statement at compile-time based on a constant expression condition.

```
if constexpr (cond)
    statement1; // Discarded if cond is false
else
    statement2; // Discarded if cond is true
```

For example:

```
template <typename T>
auto get_value(T t)
{
    if constexpr (std::is_pointer_v<T>)
        return *t;
    else
        return t;
}
```

if constexpr has the potential to simplify a lot of template code - especially when tag dispatching, SFINAE or preprocessor techniques are used.

Why Compile Time If?

At first, you may ask, why do we need if constexpr and those complex templated expressions... wouldn't normal if work?

Here's a code example:

```

template <typename Concrete, typename... Ts>
unique_ptr<Concrete> constructArgs(Ts&&... params)
{
    if (is_constructible_v<Concrete, Ts...>) // normal `if`
        return make_unique<Concrete>(forward<Ts>(params)...);
    else
        return nullptr;
}

```

The above routine is an “updated” version of `make_unique`: it returns `unique_ptr` when the parameters allow it to construct the wrapped objects, or it returns `nullptr`.

Below there’s simple code that tests `constructArgs`:

```

class Test
{
public:
    Test(int, int) { }
};

int main()
{
    auto p = constructArgs<Test>(10, 10, 10); // 3 args!
}

```

The code tries to build `Test` out of three parameters, but please notice that `Test` has only constructor that takes two `int` arguments.

When compiling you might get a similar compiler error:

```

In instantiation of 'typename std::_MakeUniq<Tp>::__single_object std::make_unique(\
_Args&& ...) [with _Tp = Test; _Args = {int, int, int}; typename std::_MakeUniq<Tp>\
::__single_object = std::unique_ptr<Test, std::default_delete<Test> >]':

```

```

main.cpp:8:40:   required from 'std::unique_ptr<Tp> constructArgs(Ts&& ...) [with C\
oncrete = Test; Ts = {int, int, int}]'

```

Let’s try to understand this error message. After the template deduction the compiler compiles the following code:

```

if (std::is_constructible_v<Concrete, 10, 10, 10>)
    return std::make_unique<Concrete>(10, 10, 10);
else
    return nullptr;

```

During the runtime the `if` branch won't be executed - as `is_constructible_v` returns false, yet the code in the branch must compile.

That's why we need `if constexpr`, to "discard" code and compile only the matching statement.

To fix the code you have to add `constexpr`:

```

template <typename Concrete, typename... Ts>
unique_ptr<Concrete> constructArgs(Ts&&... params)
{
    if constexpr (is_constructible_v<Concrete, Ts...>) // fixed!
        return make_unique<Concrete>(forward<Ts>(params)...);
    else
        return nullptr;
}

```

Now, the compiler evaluates the `if constexpr` condition at compile time and for the expression `auto p = constructArgs<Test>(10, 10, 10);` the whole `if` branch will be "removed" from the second step of the compilation process.

To be precise, the code in the discarded branch is not entirely removed from the compilation phase. Only expressions that are dependent on the template parameter used in the condition are not instantiated. The syntax must always be valid.

For example:

```

template <typename T>
void Calculate(T t)
{
    if constexpr (is_integral_v<T>)
    {
        // ...
        static_assert(sizeof(int) == 100);
    }
    else
    {
        execute(t);
        strange syntax
    }
}

```

In the above artificial code, if the type `T` is `int`, then the `else` branch is discarded, which means `execute(t)` won't be instantiated. But the line `strange syntax` will still be compiled (as it's not dependent on `T`) and that's why you'll get a compile error about that.

Furthermore, another compilation error will come from `static_assert`, the expression is also not dependent on `T`, and that's why it will always be evaluated.

Template Code Simplification

Before C++17 if you had several versions of an algorithm - depending on the type requirements - you could use SFINAE or tag dispatching to generate a dedicated overload resolution set.

For example:

```
// Chapter Templates/sfinae_example.cpp
template <typename T>
std::enable_if_t<std::is_integral_v<T>, T> simpleTypeInfo(T t)
{
    std::cout << "foo<integral T> " << t << '\n';
    return t;
}

template <typename T>
std::enable_if_t<!std::is_integral_v<T>, T> simpleTypeInfo(T t)
{
    std::cout << "not integral \n";
    return t;
}
```

In the above example we have two function implementations, but only one of them will end up in the overload resolution set. If `std::is_integral_v` is true for the `T` type then the top function is taken, and the second one rejected due to SFINAE.

The same thing can happen when using tag dispatching:


```
// Chapter Templates/tag_dispatching_example.cpp
template <typename T>
T simpleTypeInfoTagImpl(T t, std::true_type)
{
    std::cout << "foo<integral T> " << t << '\n';
    return t;
}

template <typename T>
T simpleTypeInfoTagImpl(T t, std::false_type)
{
    std::cout << "not integral \n";
    return t;
}

template <typename T>
T simpleTypeInfoTag(T t)
{
    return simpleTypeInfoTagImpl(t, std::is_integral<T>{});
}
```

Now, instead of SFINAE, we generate a unique type tag for the condition: `true_type` or `false_type`. Depending on the result only one implementation is selected.

We can now simplify this pattern with `if constexpr`:

```
template <typename T>
T simpleTypeInfo(T t)
{
    if constexpr (std::is_integral_v<T>)
    {
        std::cout << "foo<integral T> " << t << '\n';
    }
    else
    {
        std::cout << "not integral \n";
    }
    return t;
}
```

Writing template code becomes more “natural” and doesn’t require that many “tricks”.

Examples

Let's see a couple of examples:

Line Printer

You might have already seen the below example in the Jump Start section at the beginning of this Part of the book. Lets dive into the details and see how the code works.

```
template<typename T> void linePrinter(const T& x)
{
    if constexpr (std::is_integral_v<T>)
    {
        std::cout << "num: " << x << '\n';
    }
    else if constexpr (std::is_floating_point_v<T>)
    {
        const auto frac = x - static_cast<long>(x);
        std::cout << "flt: " << x << ", frac " << frac << '\n';
    }
    else if constexpr (std::is_pointer_v<T>)
    {
        std::cout << "ptr, ";
        linePrinter(*x);
    }
    else
    {
        std::cout << x << '\n';
    }
}
```

linePrinter uses `if constexpr` to check the input type. Based on that we can output additional messages. An interesting thing happens with the pointer type - when a pointer is detected the code dereferences it and then calls `linePrinter` recursively.

Declaring Custom `get<N>` Functions

The structured binding expression works for simple structures that have all public members, like

```

struct S
{
    int n;
    std::string s;
    float d;
};

S s;
auto [a, b, c] = s;

```

However, if you have a custom type (with private members), then it's also possible to override `get<N>` functions so that structured binding can work. Here's some code to demonstrate this idea:

```

class MyClass
{
public:
    int GetA() const { return a; }
    float GetB() const { return b; }

private:
    int a;
    float b;
};

template <std::size_t I> auto get(MyClass& c)
{
    if constexpr (I == 0)      return c.GetA();
    else if constexpr (I == 1) return c.GetB();
}

// specialisations to support tuple-like interface
namespace std
{
    template <> struct tuple_size<MyClass> : std::integral_constant<size_t, 2> { };

    template <> struct tuple_element<0,MyClass> { using type = int; };
    template <> struct tuple_element<1,MyClass> { using type = float; };
}

```

In the above code you have the advantage of having everything in one function. It's also possible to do it as template specialisations:

```
template <> auto& get<0>(MyClass &c) { return c.GetA(); }  
template <> auto& get<1>(MyClass &c) { return c.GetB(); }
```

For more examples you can read the chapter about [Replacing `std::enable_if` with `if constexpr`](#) and also the chapter [Structured Bindings](#) - the section about custom `get<N>` specialisations.

You can also see the following blog post at bfilipek.com: [Simplify code with ‘if constexpr’ in C++17](#)⁷



Extra Info

The change was proposed in: [P0292R2](#)⁸.

⁷<https://www.bfilipek.com/2018/03/ifconstexpr.html>

⁸<https://wg21.link/p0292r2>

Declaring Non-Type Template Parameters With `auto`

This is another part of the strategy to use `auto` everywhere. With C++11 and C++14 you can use it to deduce variables or even return types automatically, plus there are also generic lambdas. Now you can also use it for deducing non-type template parameters.

For example:

```
template <auto value> void f() { }  
f<10>();           // deduces int
```

This is useful, as you don't have to have a separate parameter for the type of non-type parameter. Like in C++11/14:

```
template <typename Type, Type value> constexpr Type TConstant = value;  
constexpr auto const MySuperConst = TConstant<int, 100>;
```

With C++17 it's a bit simpler:

```
template <auto value> constexpr auto TConstant = value;  
constexpr auto const MySuperConst = TConstant<100>;
```

No need to write `Type` explicitly.

As one of the advanced uses a lot of papers, and articles point to an example of heterogeneous compile time list:

```
template <auto ... vs> struct HeterogenousValueList {};  
using MyList = HeterogenousValueList<'a', 100, 'b'>;
```

Before C++17 it was not possible to declare such list directly, some wrapper class would have had to be provided first.



Extra Info

The change was proposed in: [P0127R2](https://wg21.link/p0127r2)⁹. In [P0127R1](https://wg21.link/p0127r1)¹⁰ you can find some more examples and reasoning.

⁹<https://wg21.link/p0127r2>

¹⁰<https://wg21.link/p0127r1>

Other Changes

In C++17 there are also other language features related to templates that are worth mentioning:

Allow `typename` in a template template parameters.

Allows you to use `typename` instead of `class` when declaring a template template parameter. Normal type parameters can use them interchangeably, but template template parameters were restricted to `class`.

More information in [N4051](https://ericniebler.com/2016/05/11/n4051/)¹¹.

Allow constant evaluation for all non-type template arguments

Remove syntactic restrictions for pointers, references, and pointers to members that appear as non-type template parameters.

More information in [N4268](https://ericniebler.com/2016/05/11/n4268/)¹².

Variable templates for traits

All the type traits that yields `::value` got accompanying `_v` variable templates. For example:

`std::is_integral<T>::value` can become `std::is_integral_v<T>`

`std::is_class<T>::value` can become `std::is_class_v<T>`

This improvement already follows the `_t` suffix additions in C++14 (template aliases) to type traits that returns `::type`.

More information in [P0006R0](https://ericniebler.com/2016/05/11/p0006r0/)¹³.

Pack expansions in using-declarations

The feature is an enhancement for variadic templates and parameter packs.

The compiler will now support the `using` keyword in pack expansions:

```
template<class... Ts> struct overloaded : Ts... {  
    using Ts::operator()...;  
};
```

¹¹<https://wg21.link/n4051>

¹²<https://wg21.link/n4268>

¹³<https://wg21.link/p0006r0>

The overloaded class exposes all overloads for `operator()` from the base classes. Before C++17 you would have to use recursion for parameter packs to achieve the same result. The overloaded pattern is a very useful enhancement for `std::visit`, read more in the “Overload” section in the Variant chapter.

More information in [P0195](https://wg21.link/P0195)¹⁴.

Logical operation metafunctions

C++17 adds handy template metafunctions:

- `template<class... B> struct conjunction;` - logical AND
- `template<class... B> struct disjunction;` - logical OR
- `template<class B> struct negation;` - logical negation

Here’s an example, based on the code from the proposal:

```
template<typename... Ts>
std::enable_if_t<std::conjunction_v<std::is_same<int, Ts>...> >
PrintIntegers(Ts ... args)
{
    (std::cout << ... << args) << '\n';
}
```

The above function `PrintIntegers` works with variable number of arguments, but they all have to be of type `int`.

The helper metafunctions can increase readability of advanced template code.

More information in [P0013](https://wg21.link/P0013)¹⁵.

¹⁴<https://wg21.link/P0195>

¹⁵<https://wg21.link/P0013>

Compiler Support

Feature	GCC	Clang	MSVC
Template argument deduction for class templates	7.0/8.0 ¹⁶	5.0	VS 2017 15.7
Deduction Guides in the Standard Library	8.0 ¹⁷	7.0/in progress ¹⁸	VS 2017 15.7
Declaring non-type template parameters with auto	7.0	4.0	VS 2017 15.7
Fold expressions	6.0	3.9	VS 2017 15.5
if constexpr	7.0	3.9	VS 2017

¹⁶Additional improvements for Template Argument Deduction for Class Templates happened in GCC 8.0, [P0512R0](#).

¹⁷Deduction Guides are not listed in the [status pages of LibSTDC++](#), so we can assume they were implemented as part of Template argument deduction for class templates.

¹⁸The [status page for LibC++](#) mentions that `<string>`, sequence containers, container adaptors and `<regex>` portions have been implemented so far.

5. Standard Attributes

Code annotations - attributes - are probably not the best known feature of C++. However, they might be handy for expressing additional information for the compiler and also for other programmers. Since C++11 there has been a standard way of specifying attributes. And in C++17 you got even more useful additions related to attributes.

In this chapter you'll learn:

- What are the attributes in C++
- Vendor-specific code annotations vs the Standard form
- In what cases attributes are handy
- C++11 and C++14 attributes
- New additions in C++17

Why Do We Need Attributes?

Have you ever used `__declspec`, `__attribute` or `#pragma` directives in your code?

For example:

```
// set an alignment
struct S { short f[3]; } __attribute__ ((aligned (8)));

// this function won't return
void fatal () __attribute__ ((noreturn));
```

Or for DLL import/export in MSVC:

```
#if COMPILING_DLL
    #define DLLEXPORT __declspec(dllexport)
#else
    #define DLLEXPORT __declspec(dllimport)
#endif
```

Those are existing forms of compiler specific attributes/annotations.

So what is an attribute?

An attribute is additional information that can be used by the compiler to produce code. It might be utilised for optimisation or some specific code generation (like DLL stuff, OpenMP, etc.). In addition, annotations allow you to write more expressive syntax and help other developers to reason about code.

Contrary to other languages such as C#, in C++, the compiler fixes meta information. You cannot add user-defined attributes. In C# you can ‘derive’ from `System.Attribute`.

What’s best about Modern C++ attributes?

With the modern C++, we get more and more standardised attributes that will work with other compilers. So we’re moving from compiler specific annotation to standard forms.

In the next section you’ll see how attributes used to work before C++11.

Before C++11

Each compiler introduced its own set of annotations, usually with a different keyword.

Often, you could see code with `#pragma`, `__declspec`, `__attribute` spread throughout the code.

Here’s the list of the common syntax from GCC/Clang and MSVC:

GCC Specific Attributes

GCC uses annotation in the form of `__attribute__((attr_name))`. For example:

```
int square (int) __attribute__((pure)); // pure function
```

Documentation:

- [Attribute Syntax - Using the GNU Compiler Collection \(GCC\)](#)¹
- [Using the GNU Compiler Collection \(GCC\): Common Function Attributes](#)²

MSVC Specific Attributes

Microsoft mostly used `__declspec` keyword, as their syntax for various compiler extensions. See the documentation here: [__declspec Microsoft Docs](#)³.

```
__declspec(deprecated) void LegacyCode() { }
```

Clang Specific Attributes

Clang, as it's straightforward to customise, can support different types of annotations, so look at the documentation to find more. Most of GCC attributes work with Clang.

See the documentation here: [Attributes in Clang — Clang documentation](#)⁴.

Attributes in C++11 and C++14

C++11 took one big step to minimise the need to use vendor specific syntax. By introducing the standard format, we can move a lot of compiler-specific attributes into the universal set.

C++11 provides a nicer format of specifying annotations over our code.

The basic syntax is just `[[attr]]` or `[[namespace::attr]]`.

You can use `[[attr]]` over almost anything: types, functions, enums, etc., etc.

For example:

¹<https://gcc.gnu.org/onlinedocs/gcc-4.8.1/gcc/Attribute-Syntax.html>

²<https://gcc.gnu.org/onlinedocs/gcc/Common-Function-Attributes.html#Common-Function-Attributes>

³<https://docs.microsoft.com/en-us/cpp/cpp/declspec>

⁴<https://clang.llvm.org/docs/AttributeReference.html>

```
[[abc]] void foo()  
{  
  
}
```

In C++11 we have the following attributes:

[[noreturn]] :

A function does not return. The behaviour is undefined if the function with this attribute returns.

- for example `[[noreturn]] void terminate() noexcept;`
- functions like `std::abort` or `std::exit` are marked with this attribute.

[[carries_dependency]] :

Indicates that the dependency chain in release-consume `std::memory_order` propagates in and out of the function, which allows the compiler to skip unnecessary memory fence instructions. Mostly to help to optimise multi-threaded code and when using different memory models.

C++14 added:

[[deprecated]] and [[deprecated("reason")]] :

Code marked with this attribute will be reported by the compiler. You can set a “reason” why.

Example of `[[deprecated]]`:

```
[[deprecated("use AwesomeFunc instead")]] void GoodFunc() { }  
  
// call somewhere:  
GoodFunc();
```

GCC might report a warning:

```
warning: 'void GoodFunc()' is deprecated: use AwesomeFunc instead  
[-Wdeprecated-declarations]
```

You know a bit about the old approach, new way in C++11/14... so what's the deal with C++17?

C++17 additions

With C++17 we get three more standard attributes:

- `[[fallthrough]]`
- `[[nodiscard]]`
- `[[maybe_unused]]`

Plus three supporting features.



Extra Info

The new attributes were specified in [P0188](https://wg21.link/p0188)⁵ and [P0068](https://wg21.link/p0068)⁶(reasoning).

Let's go through the new attributes first:

`[[fallthrough]]` attribute

Indicates that a fall-through in a switch statement is intentional and a warning should not be issued for it.

```
switch (c) {  
case 'a':  
    f(); // Warning! fallthrough is perhaps a programmer error  
case 'b':  
    g();  
[[fallthrough]]; // Warning suppressed, fallthrough is ok  
case 'c':  
    h();  
}
```

With this attribute, the compiler can understand the intentions of a programmer. It's also much more readable than using a comment.

`[[maybe_unused]]` attribute

Suppresses compiler warnings about unused entities when they are declared with `[[maybe_unused]]`.

⁵<https://wg21.link/p0188>

⁶<https://wg21.link/p0068>

```
static void impl1() { ... } // Compilers may warn about this
[[maybe_unused]] static void impl2() { ... } // Warning suppressed
```

```
void foo() {
    int x = 42; // Compilers may warn about this
    [[maybe_unused]] int y = 42; // Warning suppressed
}
```

[[nodiscard]] attribute

[[nodiscard]] can be applied to a function or a type declaration to mark the importance of the returned value:

```
[[nodiscard]] int Compute();
void Test() {
    Compute(); // Warning! return value of a
               //nodiscard function is discarded
}
```

If you forget to assign the result to a variable, then the compiler should emit a warning.

What it means is that you can force users to handle errors. For example, what happens if you forget about using the return value from `new` or `std::async()`?

Additionally, the attribute can be applied to types. One use case for it might be error codes:

```
enum class [[nodiscard]] ErrorCode {
    OK,
    Fatal,
    System,
    FileIssue
};

ErrorCode OpenFile(std::string_view fileName);
ErrorCode SendEmail(std::string_view sendto,
                   std::string_view text);
ErrorCode SystemCall(std::string_view text);
```

Now every time you'd like to call such functions you're "forced" to check the return value. For important functions checking return codes might be crucial and using [[nodiscard]] might save you from a few bugs.

You might also ask what it means to ‘use’ a return value?

In the Standard, it’s defined as `Discarded-value expressions`⁷ so if your function is called only for side effects (there’s no if statement around or assignment) the compiler is encouraged to report a warning.

However to suppress the warning you can explicitly cast the return value to `void` or use `[[maybe_unused]]`:

```
[[nodiscard]] int Compute();
void Test() {
    static_cast<void>(Compute()); // fine...

    [[maybe_unused]] auto ret = Compute();
}
```

In addition, in C++20 the Standard Library will apply `[[nodiscard]]` in a few places: operator `new`, `std::async()`, `std::allocate()`, `std::launder()`, and `std::empty()`.

This feature was already merged into C++20 with [P0600R1](#)⁸.

Attributes for namespaces and enumerators

The idea for attributes in C++11 was to be able to apply them to all sensible places like: classes, functions, variables, typedefs, templates, enumerations... But there was an issue in the specification that blocked attributes when they were applied on namespaces or enumerators.

This is now fixed in C++17. We can now write:

```
namespace [[deprecated("use BetterUtils")]] GoodUtils {
    void DoStuff() { }
}
```

⁷http://en.cppreference.com/w/cpp/language/expressions#Discarded-value_expressions

⁸<http://www.open-std.org/jtc1/sc22/wg21/docs/papers/2017/p0600r1.pdf>

```
namespace BetterUtils {  
    void DoStuff() { }  
}  
  
int main()  
{  
    GoodUtils::DoStuff();  
}
```

Clang reports:

```
warning: 'GoodUtils' is deprecated: use BetterUtils  
[-Wdeprecated-declarations]
```

Another example is the use of deprecated attribute on enumerators:

```
enum class ColorModes  
{  
    RGB [[deprecated("use RGB8")]],  
    RGBA [[deprecated("use RGBA8")]],  
    RGBA16F,  
    RGB8,  
    RGBA8  
};  
  
// use:  
auto colMode = ColorModes::RGBA;
```

Under GCC we'll get:

```
warning: 'RGBA' is deprecated: use RGBA8  
[-Wdeprecated-declarations]
```



Extra Info

The change was described in [N4266](https://wg21.link/n4266)⁹(wording) and [N4196](https://wg21.link/n4196)¹⁰(reasoning).

⁹<https://wg21.link/n4266>

¹⁰<https://wg21.link/n4196>

Ignore unknown attributes

The feature is mostly for clarification.

Before C++17, if you tried to use some compiler specific attribute, you might even get an error when compiling in another compiler that doesn't support it. Now, the compiler omits the attribute specification and won't report anything (or just a warning). This wasn't mentioned in the standard, and it needed clarification.

```
// compilers which don't
// support MyCompilerSpecificNamespace will ignore this attribute
[[MyCompilerSpecificNamespace::do_special_thing]]
void foo();
```

For example in GCC 7.1 there's a warnings:

```
warning: 'MyCompilerSpecificNamespace::do_special_thing'
scoped attribute directive ignored [-Wattributes]
void foo();
```



Extra Info

The change was described in [P0283R2](https://wg21.link/p0283r2)¹¹(wording) and [P0283R1](https://wg21.link/p0283r1)¹²(reasoning).

Using attribute namespaces without repetition

The feature simplifies the case where you want to use multiple attributes, like:

```
void f() {
    [[rpr::kernel, rpr::target(cpu,gpu)]] // repetition
    doTask();
}
```

Proposed change:

¹¹<https://wg21.link/p0283r2>

¹²<https://wg21.link/p0283r1>

```
void f() {  
    [[using rpr: kernel, target(cpu,gpu)]]  
    doTask();  
}
```

That simplification might help when building tools that automatically translate annotated code of that type into different programming models.



Extra Info

More details in: [P0028R4¹³](#).

Section Summary



All Attributes available in C++17:

- `[[noreturn]]`
- `[[carries_dependency]]`
- `[[deprecated]]`
- `[[deprecated("reason")]]`
- `[[fallthrough]]`
- `[[nodiscard]]`
- `[[maybe_unused]]`

Each compiler vendor can specify their own syntax and list of available attributes and extensions. In Modern C++ the committee tried to extract common parts for attributes. The target is to stick only to the Standard attributes, but that depends on your platform. For example, in the embedded environment there might be a lot of essential and platform-specific annotations that glue code and hardware together.

There's also quite an important [quote from Bjarne Stroustrup's C++11 FAQ/Attributes¹⁴](#):

There is a reasonable fear that attributes will be used to create language dialects. The recommendation is to use attributes to only control things that do not affect the meaning of a program but might help detect errors (e.g. `[[noreturn]]`) or help optimizers (e.g. `[[carries_dependency]]`).

¹³<http://wg21.link/p0028r4>

¹⁴<http://stroustrup.com/C++11FAQ.html#attributes>

Compiler support

Feature	GCC	Clang	MSVC
<code>[[fallthrough]]</code>	7.0	3.9	15.0
<code>[[nodiscard]]</code>	7.0	3.9	15.3
<code>[[maybe_unused]]</code>	7.0	3.9	15.3
Attributes for namespaces and enumerators	4.9/6 ¹⁵	3.4	14.0
Ignore unknown attributes	All versions	3.9	14.0
Using attribute namespaces without repetition	7.0	3.9	15.3

All of the above compilers also support C++11/14 attributes.

¹⁵GCC 4.9 (namespaces) / GCC 6 (enumerations)

Part 2 - The Standard Library Changes

While new language features allow you to write more compact code you also need the tools - in the form of the Standard Library types. The classes and systems that you can find in the Library can significantly enhance your productivity. C++17 offers even more handy instruments: for example the filesystem, new vocabulary types and even parallel algorithms!

In this part you'll learn:

- How to represent nullable types with `std::optional`
- What's a tagged union? And why we need a type-safe union in the form of `std::variant`
- How to represent any type with `std::any`
- How to use `string_view` to gain performance and not break your application
- What are the new string operations available in the Standard Library
- How to work with the filesystem using the Standard Library
- What are the parallel algorithms

6. `std::optional`

C++17 adds a few wrapper types that make it possible to write more expressive code. In this chapter, you'll see `std::optional`, which models a nullable type.

In this chapter you'll learn:

- Why we need nullable types
- How does `std::optional` work and what does it do
- Operations on `std::optional`
- The performance cost of using the type
- Example use cases

Introduction

By adding a boolean flag to other types, you can achieve something called “nullable types.” This kind of wrapper represents an object that might be empty in an expressive way.

While you can achieve “null-ability” by using unique values (`-1`, `infinity`, `nullptr`), it’s not as clear as the separate wrapper type. Alternatively, you could even use `std::unique_ptr<Type>` and treat the empty pointer as not initialised. That works but comes with the cost of allocating memory for the object and is not a recommended technique.

Optional types that come from the functional programming world bring type safety and expressiveness. Most other languages have something similar: for example `std::option` in Rust, `Optional<T>` in Java, `Data.Maybe` in Haskell.

`std::optional` was added in C++17 and brings a lot of experience from `boost::optional` that has been available for many years. With C++17 you can just `#include <optional>` and use the type.



What’s more `std::optional` was available also in Library Fundamentals TS, so there’s a chance that your C++14 compiler could also support it in the `<experimental/optional>` header.

`std::optional` is still a value type (so it can be copied, via deep copy). What’s more, `std::optional` doesn’t need to allocate any memory on the free store.

`std::optional` is a part of C++ **vocabulary types** along with `std::any`, `std::variant` and `std::string_view`.

When to Use

You can usually use an optional wrapper in the following scenarios:

If you want to represent a nullable type.

- Rather than using unique values (like `-1`, `nullptr`, `NO_VALUE` or something)
- For example, a user’s middle name is optional. You could assume that an empty string would work here, but knowing if a user entered something or not might be important. `std::optional<std::string>` gives you more information.

Return a result of some computation (processing) that fails to produce a value and is not an error.

For example, finding an element in a dictionary: if there’s no element under a key it’s not an error, but we need to handle the situation.

To perform lazy-loading of resources.

For example, a resource type has no default constructor, and the construction is substantial. So you can define it as `std::optional<Resource>` (and you can pass it around the system), and then load if needed later.

To pass optional parameters into functions.

There's a description from `boost.optional` which summarises when we should use the type:

From the `boost::optional` documentation: [When to use Optional](https://www.boost.org/doc/libs/1_67_0/libs/optional/doc/html/boost_optional/tutorial/when_to_use_optional.html)¹

It is recommended to use `optional<T>` in situations where there is exactly one, clear (to all parties) reason for having no value of type `T`, and where the lack of value is as natural as having any regular value of `T`.

While sometimes the decision to use `optional` might be blurry, it best suits the cases when the value is empty, and it's a normal state of the program.

Basic Example

Here's a simple example of what you can do with `optional`:

```
// UI class...
std::optional<std::string> UI::FindUserNick()
{
    if (IsNickAvailable())
        return mStrNickName; // return a string

    return std::nullopt; // same as return { };
}

// use:
std::optional<std::string> UserNick = UI->FindUserNick();
if (UserNick)
    Show(*UserNick);
```

In the above code, we define a function that returns an `optional` containing a string. If the user's nickname is available, then it will return a string. If not, then it returns `nullopt`. Later we can assign it to an `optional` and check (it converts to `bool`) if it contains any value or not. `Optional` defines operator`*` so we can easily access the contained value.

In the following sections you'll see how to create `std::optional`, operate on it, pass it around, and even what is the performance cost you might want to consider.

¹https://www.boost.org/doc/libs/1_67_0/libs/optional/doc/html/boost_optional/tutorial/when_to_use_optional.html

std::optional Creation

There are several ways to create std::optional:

- Initialise as empty
- Directly with a value
- With a value using deduction guides
- By using make_optional
- With std::in_place
- From other optional

See code below:

```
// empty:
std::optional<int> oEmpty;
std::optional<float> oFloat = std::nullopt;

// direct:
std::optional<int> oInt(10);
std::optional oIntDeduced(10); // deduction guides

// make_optional
auto oDouble = std::make_optional(3.0);
auto oComplex = std::make_optional<std::complex<double>>(3.0, 4.0);

// in_place
std::optional<std::complex<double>> o7{std::in_place, 3.0, 4.0};

// will call vector with direct init of {1, 2, 3}
std::optional<std::vector<int>> oVec(std::in_place, {1, 2, 3});

// copy from other optional:
auto oIntCopy = oInt;
```

As you can see in the above code sample, you have a lot of flexibility with the creation of optional. It's straightforward for primitive types, and this simplicity is extended even to complex types.

If you want the full control over the creation and efficiency, it's also good to know in_place helper types.

in_place Construction

std::optional is a wrapper type, so you should be able to create optional objects almost in the same way as the wrapped object. And in most cases you can:


```
std::optional<std::string> ostr{"Hello World"};
std::optional<int> oi{10};
```

You can write the above code without stating the constructor such as:

```
std::optional<std::string> ostr{std::string{"Hello World"}};
std::optional<int> oi{int{10}};
```

Because `std::optional` has a constructor that takes U&& (r-value reference to a type that converts to the type stored in the optional). In our case it's recognised as `const char*` and strings can be initialised from it.

So what's the advantage of using `std::in_place_t` in `std::optional`?

There are at least two important reasons:

- Default constructor
- Efficient construction for constructors with many arguments

Default Construction

If you have a class with a default constructor, like:

```
class UserName
{
public:
    UserName() : mName("Default")
    {

    }
    // ...
};
```

How would you create an optional that contains `UserName{}`?

You can write:

```
std::optional<UserName> u0; // empty optional
std::optional<UserName> u1{}; // also empty

// optional with default constructed object:
std::optional<UserName> u2{UserName()};
```

That works but it creates an additional temporary object. If we traced each different constructor and destructor call, we would get the following output:

```

UserName::UserName('Default')
UserName::UserName(move 'Default') // move temp object
UserName::~~UserName('')           // delete the temp object
UserName::~~UserName('Default')

```

The code creates a temporary object and then moves it into the object stored in `optional`.

Here we can use a more efficient constructor - specifically by leveraging `std::in_place_t`:

```
std::optional<UserName> opt{std::in_place};
```

With constructor and destructor traces you'd get the following output:

```

UserName::UserName('Default')
UserName::~~UserName('Default')

```

The object stored in the `optional` is created in place, in the same way as you'd call `UserName{}`. No additional copy or move is needed.

See the example in Chapter `Optional/optional_in_place_default.cpp`. In the file, you'll also see the traces for constructors and destructor.

Non Copyable/Movable Types

As you saw in the example from the previous section, if you use a temporary object to initialise the contained value inside `std::optional` then the compiler will have to use a move or a copy constructor.

But what if your type doesn't allow that? For example `std::mutex` is not movable or copyable.

In that case `std::in_place` is the only way to work with such types.

Constructors With Many Arguments

Another use case is a situation where your type has more arguments in a constructor. By default `optional` can work with a single argument (r-value ref), and efficiently pass it to the wrapped type. But what if you'd like to initialise `Point(x, y)`?

You can always create a temporary copy and then pass it in the construction:

```

struct Point
{
    Point(int a, int b) : x(a), y(b) { }

    int x;
    int y;
};

std::optional<Point> opt{Point{0, 0}};

```

or use `in_place` and the version of the constructor that handles variable argument list:

```

template< class... Args >
constexpr explicit optional( std::in_place_t, Args&&... args );

// or initializer_list:

template< class U, class... Args >
constexpr explicit optional( std::in_place_t,
                             std::initializer_list<U> ilist,
                             Args&&... args );

std::optional<Point> opt{std::in_place_t, 0, 0};

```

The second option is quite verbose and omits to create temporary objects. Temporaries - especially for containers or larger objects, are not as efficient as constructing in place.

Try playing with the example that is located in:

Chapter Optional/optional_point.cpp.

std::make_optional()

If you don't like `std::in_place` then you can look at `make_optional` factory function.

The code:

```

auto opt = std::make_optional<UserName>();

auto opt = std::make_optional<Point>(0, 0);

```

Is as efficient as:

```
std::optional<UserName> opt{std::in_place};
```

```
std::optional<Point> opt{std::in_place_t, 0, 0};
```

make_optional implements in place construction equivalent to:

```
return std::optional<T>(std::in_place, std::forward<Args>(args)...);
```

And also thanks to [mandatory copy elision from C++17](#) there is no temporary object involved.

Returning std::optional

If you return an optional from a function, then it's very convenient to return just std::nullopt or the computed value.

```
std::optional<std::string> TryParse(Input input)
{
    if (input.valid())
        return input.asString();

    return std::nullopt;
}
```

In the above example you can see that the function returns std::string computed from input.asString() and it's wrapped in optional. If the value is unavailable, then you can return std::nullopt.

Of course, you can also declare an empty optional at the beginning of your function and reassign if you have the computed value. So we could rewrite the above example as:

```
std::optional<std::string> TryParse(Input input)
{
    std::optional<std::string> oOut; // empty

    if (input.valid())
        oOut = input.asString();

    return oOut;
}
```

However due to [mandatory copy elision from C++17](#) it's more optimal to use the first version - not named optional - if possible. That way you'll avoid creating temporaries.

Be Careful With Braces when Returning

You might be surprised by the following code²:

```
std::optional<std::string> CreateString()
{
    std::string str {"Hello Super Awesome Long String"};
    return {str}; // this one will cause a copy
//    return str; // this one moves
}
```

According to the Standard if you wrap a return value into braces {} then you prevent move operations from happening. The returned object will be copied only.

This is similar to the case with non-copyable types:

```
std::unique_ptr<int> foo()
{
    std::unique_ptr<int> p;
    return {p}; // uses copy of unique_ptr and so it breaks...
//    return p; // this one moves, so it's fine with unique_ptr
}
```

The Standard says [\[class.copy.elision\]/3](#)³

In the following copy-initialisation contexts, a move operation might be used instead of a copy operation: - If the expression in a return statement ([stmt.return]) is a **(possibly parenthesised) id-expression** that names an object with automatic storage duration declared in the body or parameter-declaration-clause of the innermost enclosing function or lambda-expression, or - if the operand of a throw-expression is the name of a non-volatile automatic object (other than a function or catch-clause parameter) whose scope does not extend beyond the end of the innermost enclosing try-block (if there is one),

Try playing with the example that is located in:

Chapter Optional/optional_return.cpp.

The code shows a few examples with `std::unique_ptr`, `std::vector`, `std::string` and a custom type.

²Thanks to JFT for pointing that problem out.

³<http://www.eel.is/c++draft/class.copy.elision#3>

Accessing The Stored Value

Probably the most important operation for optional (apart from creation) is the way you can fetch the contained value.

There are several options:

- `operator*` and `operator->` - if there's no value the behaviour is **undefined!**
- `value()` - returns the value, or throws `std::bad_optional_access`
- `value_or(defaultVal)` - returns the value if available, or `defaultVal` otherwise

To check if the value is present you can use the `has_value()` method or just check `if (optional)` as optional is contextually convertible to `bool`.

Here's an example:

```
// by operator*
std::optional<int> oint = 10;
std::cout<< "oint " << *oint << '\n';

// by value()
std::optional<std::string> ostr("hello");
try
{
    std::cout << "ostr " << ostr.value() << '\n';
}
catch (const std::bad_optional_access& e)
{
    std::cout << e.what() << '\n';
}

// by value_or()
std::optional<double> odouble; // empty
std::cout<< "odouble " << odouble.value_or(10.0) << '\n';
```

So the most useful way is probably just to check if the value is there and then access it:

```
// compute string function:
std::optional<std::string> maybe_create_hello();
// ...

if (auto ostr = maybe_create_hello(); ostr)
    std::cout << "ostr " << *ostr << '\n';
else
    std::cout << "ostr is null\n";
```

std::optional Operations

Let's see what other operations are available for the type.

Changing the Value

If you have an existing optional object, then you can quickly change the contained value by using several operations like `emplace`, `reset`, `swap`, `assign`. If you assign (or reset) with a `nullopt` then if the optional contains a value its destructor will be called.

Here's a quick summary:

```
// chapter Optional/optional_reset.cpp
#include <optional>
#include <iostream>
#include <string>

class UserName
{
public:
    explicit UserName(std::string str) : mName(std::move(str))
    {
        std::cout << "UserName::UserName(' " << mName << "')\n";
    }
    ~UserName()
    {
        std::cout << "UserName::~~UserName(' " << mName << "')\n";
    }

private:
    std::string mName;
};
```

```

int main()
{
    std::optional<UserName> oEmpty;

    // emplace:
    oEmpty.emplace("Steve");

    // calls ~Steve and creates new Mark:
    oEmpty.emplace("Mark");

    // reset so it's empty again
    oEmpty.reset(); // calls ~Mark
    // same as:
    //oEmpty = std::nullopt;

    // assign a new value:
    oEmpty.emplace("Fred");
    oEmpty = UserName("Joe");
}

```

Comparisons

std::optional allows you to compare contained objects almost “normally”, but with a few exceptions when the operands are nullopt. See below:

```

// chapter Optional/optional_comparision.cpp
#include <optional>
#include <iostream>

int main()
{
    std::optional<int> oEmpty;
    std::optional<int> oTwo(2);
    std::optional<int> oTen(10);

    std::cout << std::boolalpha;
    std::cout << (oTen > oTwo) << '\n';
    std::cout << (oTen < oTwo) << '\n';
    std::cout << (oEmpty < oTwo) << '\n';
    std::cout << (oEmpty == std::nullopt) << '\n';
}

```



```
std::cout << (oTen == 10) << '\n';
}
```

The above code generates:

```
true  // (oTen > oTwo)
false // (oTen < oTwo)
true  // (oEmpty < oTwo)
true  // (oEmpty == std::nullopt)
true  // (oTen == 10)
```

When operands contain values (of the same type), then you'll see the expected results. But when one operand is `nullopt` then it's always "less" than any optional with some value.

Examples of std::optional

Here are a few more extended examples where `std::optional` fits nicely.

User Name with an Optional Nickname and Age

```
// Chapter Optional/optional_user_name.cpp
#include <optional>
#include <iostream>

using namespace std;

class UserRecord
{
public:
    UserRecord (string name, optional<string> nick, optional<int> age)
        : mName{move(name)}, mNick{move(nick)}, mAge{age}
    {
    }

    friend ostream& operator << (ostream& stream, const UserRecord& user);

private:
    string mName;
    optional<string> mNick;
    optional<int> mAge;
};
```

```
ostream& operator << (ostream& os, const UserRecord& user)
{
    os << user.mName;

    if (user.mNick)
        os << ' ' << *user.mNick;

    if (user.mAge)
        os << ' ' << "age of " << *user.mAge;

    return os;
}

int main()
{
    UserRecord tim { "Tim", "SuperTim", 16 };
    UserRecord nano { "Nathan", nullopt, nullopt };

    cout << tim << '\n';
    cout << nano << '\n';
}
```

The above example shows a simple class with optional fields. While the name is obligatory, the other attributes: “nickname” and “age” are optional.

Parsing ints From the Command Line

```
// Chapter Optional/optional_parsing.cpp
#include <optional>
#include <iostream>
#include <string>

std::optional<int> ParseInt(const char* arg)
{
    try
    {
        return { std::stoi(std::string(arg)) };
    }
    catch (...)
    {
        std::cerr << "cannot convert '" << arg << "' to int!\n";
    }
}
```

```

    }

    return { };
}

int main(int argc, const char* argv[])
{
    if (argc >= 3)
    {
        auto oFirst = ParseInt(argv[1]);
        auto oSecond = ParseInt(argv[2]);

        if (oFirst && oSecond)
        {
            std::cout << "sum of " << *oFirst << " and " << *oSecond;
            std::cout << " is " << *oFirst + *oSecond << '\n';
        }
    }
}

```

The above code uses optional to indicate whether we performed the conversion or not. Note that we actually converted exception handling into optional, so we skip the errors that might appear. Such a technique might look “controversial”.

The code also uses `stoi` which might be replaced with new methods `from_chars`.

Other Examples

Here are a few more ideas where you might use `std::optional`:

- Representing optional configuration values
- Geometry & Math: finding if there’s an intersection between objects
- Return values for `Find*()` functions (assuming you don’t care about errors, like connection drops, database errors or something)

You may find other interesting uses in the post: [A Wall of Your std::optional Examples](https://www.bfilipek.com/2018/06/optional-examples-wall.html)⁴. The blog post contains a lot of examples submitted by blog readers.

Performance & Memory Consideration

When you use `std::optional` you’ll pay with an increased memory footprint.

Conceptually your version of the standard library might implement optional as:

⁴<https://www.bfilipek.com/2018/06/optional-examples-wall.html>

```

template <typename T>
class optional
{
    bool _initialized;
    std::aligned_storage_t<sizeof(t), alignof(T)> _storage;

public:
    // operations
};

```

In short `optional` wraps your type, prepares space for it and then adds one boolean parameter. This means it will extend the size of your Type according to the alignment rules.

Alignment rules are important as The standard defines:

Class template `optional` [optional.optional]: The contained value shall be allocated in a region of the optional storage suitably aligned for the type T.

For example:

```

// sizeof(double) = 8
// sizeof(int) = 4
std::optional<double> od; // sizeof = 16 bytes
std::optional<int> oi; // sizeof = 8 bytes

```

While `bool` type usually takes only one byte, the `optional` type needs to obey the alignment rules, and thus the whole wrapper is larger than just `sizeof(YourType) + 1` byte.

For example, if you have a type like:

```

struct Range
{
    std::optional<double> mMin;
    std::optional<double> mMax;
};

```

it will take up more space than when you use your custom type:

```

struct Range
{
    bool mMinAvailable;
    bool mMaxAvailable;
    double mMin;
    double mMax;
};

```

In the first case, we're using 32 bytes! The second version is 24 bytes.

You can see the full code in Chapter `Optional/optional_sizeof.cpp`.

Migration from `boost::optional`

`std::optional` was adapted directly from `boost::optional`, so you should see the same experience in both versions. Moving from one to the other should be easy, but of course, there are little differences.

aspect	std::optional	boost::optional (as of 1.67.0 ⁵)
Move semantics	yes	yes
noexcept	yes	yes
hash support	yes	no
a throwing value accessor	yes	yes
literal type (can be used in constexpr expressions)	yes	no
in place construction	<code>`emplace`, tag `in_place`</code>	<code>emplace()</code> , tags <code>in_place_init_if_t</code> , <code>in_place_init_t</code> , <code>utility</code> <code>in_place_factory</code>
disengaged state tag	<code>nullopt</code>	<code>none</code>
optional references	no	yes
conversion from <code>optional<U></code> to <code>optional<T></code>	yes	yes
explicit convert to ptr (<code>get_ptr</code>)	no	yes
deduction guides	yes	no

Special case: `optional<bool>` and `optional<T*>`

While you can use `optional` on any type, you need to pay attention when trying to wrap boolean or pointers.

`std::optional<bool>` ob - what does it model? With such a construction, you have a tri-state bool. So if you need it, then maybe it's better to look for a real tri-state bool like `boost::tribool`⁶.

⁶https://www.boost.org/doc/libs/1_67_0/doc/html/tribool.html

What's more, it might be confusing to use such type because `ob` converts to `bool` if there's a value inside and `*ob` returns that stored value (if available).

Likewise you have a similar ambiguity with pointers:

```
// Don't try doing it this way, it's just an example!
std::optional<int*> opi { new int(10) };
if (opi && *opi)
{
    std::cout << **opi << std::endl;
    delete *opi;
}
if (opi)
    std::cout << "opi is still not empty!";
```

The pointer to `int` is naturally “nullable”, so wrapping it into `optional` makes it very hard to use.

Summary

A few core elements to know about `std::optional`:

- `std::optional` is a wrapper type to express “null-able” types
- `std::optional` won't use any dynamic allocation
- `std::optional` contains a value or it's empty
 - use operator `*`, operator `->`, `value()` or `value_or()` to access the underlying value.
- `std::optional` is implicitly converted to `bool` so that you can easily check if it contains a value or not

Compiler Support

Feature	GCC	Clang	MSVC
<code>std::optional</code>	7.1	4.0	VS 2017 15.0

7. `std::variant`

Another handy wrapper type that we get in C++17 is `std::variant`. This is a type-safe union - you can store different type variants with the proper object lifetime guarantee. The new type offers a huge advantage over the C-style union. You can store all of the types inside - no matter if it's something simple like `int`, or `float`, but also complex entities like `std::vector<std::string>`. In all of the cases, objects will be correctly initialised and cleaned up.

What's crucial is the fact that the new type enhances implementation of design patterns. For example, you can now use a visitor, pattern matching and runtime polymorphism for unrelated type hierarchies in a much easier way.

In this chapter you'll learn:

- What problems can occur with unions
- What discriminated unions are, and why we need type-safety with unions
- How `std::variant` works and what it does
- Operations on `std::variant`
- Performance cost and memory requirements
- Example use cases

The Basics

Unions are rarely used in the client code, and most of the time they should be avoided.

For example, there's a “common” trick with floating point operations:

```
union SuperFloat
{
    float f;
    int i;
}

int RawMantissa(SuperFloat f)
{
    return f.i & ((1 << 23) - 1);
}

int RawExponent(SuperFloat f)
{
    return (f.i >> 23) & 0xFF;
}
```

However, while the above code might work in C99, due to stricter aliasing rules it's undefined behaviour in C++!

There's an existing Core Guideline Rule on that:

C.183: Don't use a union for type punning

Reason It is undefined behaviour to read a union member with a different type from the one with which it was written. Such punning is invisible, or at least harder to spot than using a named cast. Type punning using a union is a source of errors.

Before C++17, if you wanted a type-safe union, you could use boost variant or another third-party library. But now you have `std::variant`.

Here's a basic demo of what you can do with this new type:


```
1 // Chapter Variant/variant_demo.cpp
2 #include <string>
3 #include <iostream>
4 #include <variant>
5
6 using namespace std;
7
8 // used to print the currently active type
9 struct PrintVisitor
10 {
11     void operator()(int i) { cout << "int: " << i << '\n'; }
12     void operator()(float f) { cout << "float: " << f << '\n'; }
13     void operator()(const string& s) { cout << "str: " << s << '\n'; }
14 };
15
16 int main()
17 {
18     variant<int, float, string> intFloatString;
19     static_assert(variant_size_v<decltype(intFloatString)> == 3);
20
21     // default initialised to the first alternative, should be 0
22     visit(PrintVisitor{}, intFloatString);
23
24     // index will show the currently used 'type'
25     cout << "index = " << intFloatString.index() << endl;
26     intFloatString = 100.0f;
27     cout << "index = " << intFloatString.index() << endl;
28     intFloatString = "hello super world";
29     cout << "index = " << intFloatString.index() << endl;
30
31     // try with get_if:
32     if (const auto intPtr = get_if<int>(&intFloatString))
33         cout << "int: " << *intPtr << '\n';
34     else if (const auto floatPtr = get_if<float>(&intFloatString))
35         cout << "float: " << *floatPtr << '\n';
36
37     if (holds_alternative<int>(intFloatString))
38         cout << "the variant holds an int!\n";
39     else if (holds_alternative<float>(intFloatString))
40         cout << "the variant holds a float\n";
41     else if (holds_alternative<string>(intFloatString))
42         cout << "the variant holds a string\n";
43 }
```

```

44     // try/catch and bad_variant_access
45     try
46     {
47         auto f = get<float>(intFloatString);
48         cout << "float! " << f << '\n';
49     }
50     catch (bad_variant_access&)
51     {
52         cout << "our variant doesn't hold float at this moment...\n";
53     }
54 }

```

The output:

```

int: 0
index = 0
index = 1
index = 2
the variant holds a string
our variant doesn't hold float at this moment...

```

Several points are worth examining in the example above:

- Line 18, 22: If you don't initialise a variant with a value, then the variant is initialised with the first type. In that case, the first alternative type must have a default constructor. Line 22 will print the value 0.
- Line: 25, 27, 29, 31, 33, 35: You can check what the currently used type is via `index()` or `holds_alternative`.
- Line 39, 41, 47: You can access the value by using `get_if` or `get` (might throw `bad_variant_access` exception).
- Type Safety - the variant doesn't allow you to get a value of the type that's not active.
- No extra heap allocation occurs.
- Line 9, 22: You can use a visitor to invoke an action on a currently active type. The example uses `PrintVisitor` to print the currently active value. It's a simple structure with overloads for `operator()`. The visitor is then passed to `std::visit` which performs the visitation.
- The variant class calls destructors and constructors of non-trivial types, so in the example, the string object is cleaned up before we switch to new variants.

When to Use

Unless you're doing some low-level stuff, possibly only with simple types, then unions might be a valid option¹. But for all other use cases, where you need alternative types, `std::variant` is the way to go.

¹See [C++ Core Guidelines - Unions](#) for examples of a valid use cases for unions.

Some possible uses:

- All the places where you might get a few types for a single field: so things like parsing command lines, ini files, language parsers, etc.
- Expressing several possible outcomes of a computation efficiently: like finding roots of equations.
- Error handling - for example you can return `variant<Object, ErrorCode>`. If the value is available, then you return `Object` otherwise you assign some error code.
- Finite State Machines.
- Polymorphism without `vtables` and inheritance (thanks to the visitor pattern).

A Functional Background

It's also worth mentioning that variant types (also called a tagged union, a discriminated union, or a sum type) come from the functional language world and [Type Theory](#)².

std::variant Creation

There are several ways you can create and initialize `std::variant`:

```
// Chapter Variant/variant_creation.cpp

// default initialisation: (the first type has to have a default ctor)
std::variant<int, float> intFloat;
std::cout << intFloat.index() << ", val: " << std::get<int>(intFloat) << '\n';

// monostate for default initialisation:
class NotSimple
{
public:
    NotSimple(int, float) { }
};

// std::variant<NotSimple, int> cannotInit; // error
std::variant<std::monostate, NotSimple, int> okInit;
std::cout << okInit.index() << '\n';

// pass a value:
std::variant<int, float, std::string> intFloatString { 10.5f };
std::cout << intFloatString.index()
```

²https://en.wikipedia.org/wiki/Algebraic_data_type

```

    << ", value " << std::get<float>(intFloatString) << '\n';

// ambiguity
// double might convert to float or int, so the compiler cannot decide

//std::variant<int, float, std::string> intFloatString { 10.5 };

// ambiguity resolved by in_place
variant<long, float, std::string> longFloatString {
    std::in_place_index<1>, 7.6 // double!
};
std::cout << longFloatString.index() << ", value "
    << std::get<float>(longFloatString) << '\n';

// in_place for complex types
std::variant<std::vector<int>, std::string> vecStr {
    std::in_place_index<0>, { 0, 1, 2, 3 }
};
std::cout << vecStr.index() << ", vector size "
    << std::get<std::vector<int>>(vecStr).size() << '\n';

// copy-initialize from other variant:
std::variant<int, float> intFloatSecond { intFloat };
std::cout << intFloatSecond.index() << ", value "
    << std::get<int>(intFloatSecond) << '\n';

```

Play with the code here [@Coliru³](https://coliru.stacked-crooked.com/a/1b930be2da3bd969).

- By default, a variant object is initialised with the first type,
 - if that's not possible when the type doesn't have a default constructor, then you'll get a compiler error
 - you can use `std::monostate` to pass it as the first type in that case
 - `monostate` allows you to build a variant with “no-value” so it can behave similarly to `optional`
- You can initialise it with a value, and then the best matching type is used
 - if there's an ambiguity, then you can use a version `std::in_place_index` to explicitly mention what type should be used
- `std::in_place` also allows you to create more complex types and pass more parameters to the constructor

³[http://coliru.stacked-crooked.com/a/1b930be2da3bd969](https://coliru.stacked-crooked.com/a/1b930be2da3bd969)

About std::monostate

In the example, you might notice a special type called `std::monostate`. This is just an empty type that can be used with variants to represent an empty state. The type might be handy when the first alternative doesn't have a default constructor. In that situation, you can place `std::monostate` as the first alternative (or you can also shuffle the types, and find the type with a default constructor).

In Place Construction

`std::variant` has two `in_place` helpers that you can use:

- `std::in_place_type` - used to specify which type you want to change/set in the variant
- `std::in_place_index` - used to specify which index you want to change/set. Types are enumerated from 0.
 - In a variant `std::variant<int, float, std::string>` - `int` has the index 0, `float` has index 1 and the string has index of 2. The index is the same value as returned from `variant::index` method.

Fortunately, you don't always have to use the helpers to create a variant. It's smart enough to recognise if it can be constructed from the passed single parameter:

```
// this constructs the second/float:  
std::variant<int, float, std::string> intFloatString { 10.5f };
```

For variant we need the helpers for at least two cases:

- ambiguity - to distinguish which type should be created where several could match
- efficient complex type creation (similar to optional)

Ambiguity

What if you have an initialisation like:

```
std::variant<int, float> intFloat { 10.5 }; // conversion from double?
```

The value `10.5` could be converted to `int` or `float`, and the compiler doesn't know which conversion should be applied. It might report a few pages of compiler errors, however.

But you can easily handle such errors by specifying which type you'd like to create:

```
std::variant<int, float> intFloat { std::in_place_index<0>, 10.5f };
```

```
// or
```

```
std::variant<int, float> intFloat { std::in_place_type<int>, 10.5f };
```

Complex Types

Similarly to `std::optional`, if you want to efficiently create objects that get several constructor arguments - then just use `std::in_place*`:

For example:

```
std::variant<std::vector<int>, std::string> vecStr {
    std::in_place_index<0>, { 0, 1, 2, 3 } // initializer list passed into vector
};
```

Changing the Values

There are four ways to change the current value of the variant:

- the assignment operator
- `emplace`
- `get` and then assign a new value for the currently active type
- a visitor (you cannot change the type, but you can change the value of the current alternative)

The important part is to know that everything is type safe and also that the object lifetime is honoured.

```
// Chapter Variant/variant_changing_values.cpp
```

```
std::variant<int, float, std::string> intFloatString { "Hello" };
```

```
intFloatString = 10; // we're now an int
```

```
intFloatString.emplace<2>(std::string("Hello")); // we're now string again
```

```
// std::get returns a reference, so you can change the value:
```

```
std::get<std::string>(intFloatString) += std::string(" World");
```

```
intFloatString = 10.1f;
```

```
if (auto pFloat = std::get_if<float>(&intFloatString); pFloat)
    *pFloat *= 2.0f;
```

Object Lifetime

When you use union, you need to manage the internal state: call constructors or destructors. This is error-prone and it's easy to shoot yourself in the foot. But `std::variant` handles object lifetime as you expect. That means that if it's about to change the currently stored type, then a destructor of the underlying type is called.

```
std::variant<std::string, int> v { "Hello A Quite Long String" };  
// v allocates some memory for the string  
v = 10; // we call destructor for the string!  
// no memory leak
```

Or see this example with a custom type:

```
// Chapter Variant/variant_lifetime.cpp  
class MyType  
{  
    public:  
        MyType() { std::cout << "MyType::MyType\n"; }  
        ~MyType() { std::cout << "MyType::~~MyType\n"; }  
};  
  
class OtherType  
{  
    public:  
        OtherType() { std::cout << "OtherType::OtherType\n"; }  
        ~OtherType() { std::cout << "OtherType::~~OtherType\n"; }  
};  
  
int main()  
{  
    std::variant<MyType, OtherType> v;  
    v = OtherType();  
  
    return 0;  
}
```

This will produce the output:

```

MyType::MyType
OtherType::OtherType
MyType::~~MyType
OtherType::~~OtherType
OtherType::~~OtherType

```

At the start, we initialise with a default value of type `MyType`; then we change the value with an instance of `OtherType`, and before the assignment, the destructor of `MyType` is called. Later we destroy the temporary object and the object stored in the variant.

Accessing the Stored Value

From all of the examples you've seen so far, you might get an idea of how to access the value. But let's make a summary of this vital operation.

First of all, even if you know what the currently active type is you cannot do:

```

std::variant<int, float, std::string> intFloatString { "Hello" };
std::string s = intFloatString;

```

```

// error: conversion from
// 'std::variant<int, float, std::string>'
// to non-scalar type 'std::string' requested
// std::string s = intFloatString;

```

So you have to use helper functions to access the value.

You have `std::get<Type|Index>(variant)` which is a non member function. It returns a reference to the desired type if it's active (you can pass a `Type` or `Index`). If not then you'll get `std::bad_variant_access` exception.

```

std::variant<int, float, std::string> intFloatString;
try
{
    auto f = std::get<float>(intFloatString);
    std::cout << "float! " << f << '\n';
}
catch (std::bad_variant_access&)
{
    std::cout << "our variant doesn't hold float at this moment...\n";
}

```

The next option is `std::get_if`. This function is also a non-member and won't throw. It returns a pointer to the active type or `nullptr`. While `std::get` needs a reference to the variant, `std::get_if` takes a pointer.


```
if (const auto intPtr = std::get_if<0>(&intFloatString))
    std::cout << "int!" << *intPtr << '\n';
```

However, probably the most important way to access a value inside a variant is by using visitors.

Visitors for std::variant

With the introduction of std::variant we also got a handy STL function called std::visit.

It can call a given “visitor” on all passed variants.

Here’s the declaration:

```
template <class Visitor, class... Variants>
constexpr visit(Visitor&& vis, Variants&&... vars);
```

And it will call vis on the currently active type of variants.

If you pass only one variant, then you have to have overloads for the types from that variant. If you give two variants, then you have to have overloads for all possible *pairs* of the types from the variants.

A visitor is “a Callable that accepts every possible alternative from every variant”.

Let’s see some examples:

```
// a generic lambda:
auto PrintVisitor = [](const auto& t) { std::cout << t << '\n'; };

std::variant<int, float, std::string> intFloatString { "Hello" };
std::visit(PrintVisitor, intFloatString);
```

In the above example, a generic lambda is used to generate all possible overloads. Since all of the types in the variant support << (stream output operator) then we can print them.

In another example we can use a visitor to change the value:

```

auto PrintVisitor = [](const auto& t) { std::cout << t << '\n'; };
auto TwiceMoreVisitor = [](auto& t) { t*= 2; };

```

```

std::variant<int, float> intFloat { 20.4f };
std::visit(PrintVisitor, intFloat);
std::visit(TwiceMoreVisitor, intFloat);
std::visit(PrintVisitor, intFloat);

```

Generic lambdas can work if our types share the same “interface”, but in most cases, we’d like to perform different actions based on an active type.

That’s why we can define a structure with several overloads for the operator ():

```

struct MultiplyVisitor
{
    float mFactor;

    MultiplyVisitor(float factor) : mFactor(factor) { }

    void operator()(int& i) const {
        i *= static_cast<int>(mFactor);
    }

    void operator()(float& f) const {
        f *= mFactor;
    }

    void operator()(std::string& ) const {
        // nothing to do here...
    }
};

std::visit(MultiplyVisitor(0.5f), intFloat);
std::visit(PrintVisitor, intFloat);

```

In the example, you might notice that `MultiplyVisitor` uses a state to hold the desired scaling factor value. That gives a lot of options for visitation.

With lambdas, we got used to declaring things just next to its usage. And when you need to write a separate structure, you need to go out of that local scope. That’s why it might be handy to use overload construction.

Overload

With this utility you can write all lambdas for all matching types in one place:

```
std::variant<int, float, std::string> myVariant;
std::visit(
    overload {
        [](const int& i) { std::cout << "int: " << i; },
        [](const std::string& s) { std::cout << "string: " << s; },
        [](const float& f) { std::cout << "float: " << f; }
    },
    myVariant
);
```

Currently this helper is not a part of the Standard Library (it might be added into with C++20). You can implement it with the following code:

```
template<class... Ts> struct overload : Ts... { using Ts::operator()...; };
template<class... Ts> overload(Ts...) -> overload<Ts...>;
```

The code creates a struct that inherits from lambdas and uses their `Ts::operator()`. The whole structure can now be passed to `std::visit` - it will then select the proper overload.

`overload` uses three C++17 features:

- Pack expansions in using declarations - short and compact syntax with variadic templates.
- Custom template argument deduction rules - this allows the compiler to deduce types of lambdas that are the base classes for the pattern. Without it we'd have to define a "make" function.
- Extension to aggregate Initialisation - the `overload` pattern uses aggregate initialisation to init base classes. Before C++17 it was not possible.

Here's another example of how to use it:

```
// Chapter Variant/variant_overload.cpp
std::variant<int, float, std::string> intFloatString { "Hello" };
std::visit(overload{
    [](int& i) { i*= 2; },
    [](float& f) { f*= 2.0f; },
    [](std::string& s) { s = s + s; }
}, intFloatString);
std::visit(PrintVisitor, intFloatString);
// prints: "HelloHello"
```

Here's the paper for the proposal of `std::overload`: [P0051 - C++ generic overload function](https://wg21.link/p0051r3)⁴.

And you can read more about the mechanics of the `overload` pattern in this blog post at [bfilipek.com](https://bfilipek.com/2019/02/2lines3featuresoverload.html): [2 Lines Of Code and 3 C++17 Features - The overload Pattern](https://bfilipek.com/2019/02/2lines3featuresoverload.html)⁵

⁴<https://wg21.link/p0051r3>

⁵<https://www.bfilipek.com/2019/02/2lines3featuresoverload.html>

Visiting Multiple Variants

`std::visit` allows you not only to visit one variant but many in the same call. However, it's important to know that with multiple variants you have to implement function overloads taking as many arguments as the number of input variants. And you have to provide all the possible combination of types.

For example for:

```
std::variant<int, float, char> v1 { 's' };
std::variant<int, float, char> v2 { 10 };
```

You have to provide 9 function overloads if you call `std::visit` on the two variants:

```
std::visit(overload{
    [](int a, int b) { },
    [](int a, float b) { },
    [](int a, char b) { },
    [](float a, int b) { },
    [](float a, float b) { },
    [](float a, char b) { },
    [](char a, int b) { },
    [](char a, float b) { },
    [](char a, char b) { }
}, v1, v2);
```

If you skip one overload, then the compiler will report an error.

Have a look at the following example, where each variant represents an ingredient and we want to compose two of them together:

```
// Chapter Variant/visit_multiple.cpp
#include <iostream>
#include <variant>

template<class... Ts> struct overload : Ts... { using Ts::operator()...; };
template<class... Ts> overload(Ts...) -> overload<Ts...>;

struct Pizza { };
struct Chocolate { };
struct Salami { };
struct IceCream { };
```

```

int main() {
    std::variant<Pizza, Chocolate, Salami, IceCream> firstIngredient { IceCream() };
    std::variant<Pizza, Chocolate, Salami, IceCream> secondIngredient { Chocolate()};

    std::visit(overload{
        [](const Pizza& p, const Salami& s) {
            std::cout << "here you have, Pizza with Salami!\n";
        },
        [](const Salami& s, const Pizza& p) {
            std::cout << "here you have, Pizza with Salami!\n";
        },
        [](const Chocolate& c, const IceCream& i) {
            std::cout << "Chocolate with IceCream!\n";
        },
        [](const IceCream& i, const Chocolate& c) {
            std::cout << "IceCream with a bit of Chocolate!\n";
        },
        [](const auto& a, const auto& b) {
            std::cout << "invalid composition...\n";
        },
    }, firstIngredient, secondIngredient);

    return 0;
}

```

The code will output: IceCream with a bit of Chocolate!

The above code uses overload and uses multiple lambdas rather than a separate struct with overloads for operator().

What’s interesting is that the example provides implementation only for “valid” ingredient compositions, while the “rest” is handled by generic lambdas (from C++14).

A generic lambda [](const auto& a, const auto& b) { } is equivalent to the following callable type:

```

class UnnamedUniqueClass // << compiler specific name...
{
public:
    template<typename T, typename U>
    auto operator () (const T& a, const T& b) const { }
};

```

The generic lambda used in the example will provide all the remaining function overloads for the

ingredient types. Since it's a template, it will always fall behind the concrete overload (lambda with concrete types) when the best viable function is determined.

Other std::variant Operations

Just for the sake of completeness:

- You can **compare** two variants of the same type:
 - if they contain the same active alternative, then the corresponding comparison operator is called.
 - If one variant has an “earlier” alternative then it's “less than” the variant with the next active alternative.
- Variant is a value type, so you can **move** it.
- std::hash on a variant is also possible.

Exception Safety Guarantees

So far everything looks nice and smooth... but what happens when there's an exception during the creation of the alternative in a variant?

For example:

```
// Chapter Variant/variant_valueless.cpp
class ThrowingClass
{
public:
    explicit ThrowingClass(int i) { if (i == 0) throw int (10); }
    operator int () { throw int(10); }
};

int main(int argc, char** argv)
{
    std::variant<int, ThrowingClass> v;

    // change the value:
    try
    {
        v = ThrowingClass(0);
    }
    catch (...)
    {

```

```

        std::cout << "catch(...)\n";
        // we keep the old state!
        std::cout << v.valueless_by_exception() << '\n';
        std::cout << std::get<int>(v) << '\n';
    }

    // inside emplace
    try
    {
        v.emplace<0>(ThrowingClass(10)); // calls the operator int
    }
    catch (...)
    {
        std::cout << "catch(...)\n";
        // the old state was destroyed, so we're not in invalid state!
        std::cout << v.valueless_by_exception() << '\n';
    }

    return 0;
}

```

In the first case - with the assignment operator - the exception is thrown in the constructor of the type. This happens before the old value is replaced in the variant, so the variant state is unchanged. As you can see we can still access `int` and print it.

However, in the second case - `emplace` - the exception is thrown after the old state of the variant is destroyed. `emplace` calls `operator int` to replace the value, but that throws. After that, the variant is in the wrong state, and we cannot recover the previous state.

Also note that a variant that is “valueless by exception” is in an invalid state. Accessing a value from such variant is not possible. That’s why `variant::index` returns `variant_npos`, and `std::get` and `std::visit` will throw `bad_variant_access`.

Performance & Memory Considerations

`std::variant` uses the memory in a similar way to union: so it will take the max size of the underlying types. But since we need something that will know what the currently active alternative is, then we need to use some more space. Plus everything needs to honour the alignment rules.

Here are some basic sizes:

```
// Chapter Variant/variant_sizeof.cpp
std::cout << "sizeof string: "
           << sizeof(std::string) << '\n';

std::cout << "sizeof variant<int, string>: "
           << sizeof(std::variant<int, std::string>) << '\n';

std::cout << "sizeof variant<int, float>: "
           << sizeof(std::variant<int, float>) << '\n';

std::cout << "sizeof variant<int, double>: "
           << sizeof(std::variant<int, double>) << '\n';
```

On GCC 8.1, 32 bit:

```
sizeof string: 32
sizeof variant<int, string>: 40
sizeof variant<int, float>: 8
sizeof variant<int, double>: 16
```

What's more interesting is that `std::variant` won't allocate **any extra space**! No dynamic allocation happens to hold variants or the discriminator.

To have a safe sum type you pay with an increased memory footprint. The additional bits might influence CPU caches. That's why, you might want to do some benchmarking for the hot spots in your application that uses variants.

Migration From `boost::variant`

Boost Variant was introduced around the year 2004, so it was 13 years of experience before `std::variant` was added into the Standard. The STL type draws from the experience of the boost version and improves it.

Here are the main changes:

Feature	Boost.Variant (1.67.0) ⁶	std::variant
Extra memory allocation	Possible on assignment, see Design Overview - Never Empty ⁷	No
visiting	<code>apply_visitor</code>	<code>std::visit</code>
get by index	no	yes
recursive variant	yes, see make_recursive_variant ⁸	no
duplicated entries	no	yes
empty alternative	<code>boost::blank</code>	<code>std::monostate</code>

⁷https://www.boost.org/doc/libs/1_67_0/doc/html/variant/design.html

⁸https://www.boost.org/doc/libs/1_67_0/doc/html/boost/make_recursive_variant.html

Examples of std::variant

Having learned most of the std::variant details, we can now explore a few examples.

Error Handling

The basic idea is to wrap the possible return type with some ErrorCode, and that way allow functions to output more information about the errors. Without using exceptions or output parameters. This is similar to what std::expected - a new type planned for the future C++ Standard.

```
// Chapter Variant/variant_error_handling.cpp
enum class ErrorCode
{
    Ok,
    SystemError,
    IoError,
    NetworkError
};

std::variant<std::string, ErrorCode> FetchNameFromNetwork(int i)
{
    if (i == 0)
        return ErrorCode::SystemError;

    if (i == 1)
        return ErrorCode::NetworkError;

    return std::string("Hello World!");
}

int main()
{
    auto response = FetchNameFromNetwork(0);
    if (std::holds_alternative<std::string>(response))
        std::cout << std::get<std::string>(response) << "\n";
    else
        std::cout << "Error!\n";

    response = FetchNameFromNetwork(10);
    if (std::holds_alternative<std::string>(response))
```

```
std::cout << std::get<std::string>(response) << "n";
else
    std::cout << "Error!\n";

return 0;
}
```

In the example, `ErrorCode` or a regular type is returned.

Parsing a Command Line

Command line might contain text arguments that could be interpreted in a few ways:

- as an integer
- as a floating point
- as a boolean flag
- as a string (not parsed)
- ...

We can build a variant that will hold all the possible options.

Here's a simple version with `int`, `float` and `string`:

```
// Chapter Variant/variant_parsing_int_float.cpp
class CmdLine
{
public:
    using Arg = std::variant<int, float, std::string>;

private:
    std::map<std::string, Arg> mParsedArgs;

public:
    explicit CmdLine(int argc, const char** argv) { ParseArgs(argc, argv); }

    std::optional<Arg> Find(const std::string& name) const;

    // ...
};
```

And the parsing code:

```

// Chapter Variant/variant_parsing_int_float.cpp
CmdLine::Arg TryParseString(std::string_view sv)
{
    // try with float first
    float fResult = 0.0f;
    const auto last = sv.data() + sv.size();
    const auto res = std::from_chars(sv.data(), last, fResult);
    if (res.ec != std::errc{} || res.ptr != last)
    {
        // if not possible, then just assume it's a string
        return std::string{sv};
    }

    // no fraction part? then just cast to integer
    if (static_cast<int>(fResult) == fResult)
        return static_cast<int>(fResult);

    return fResult;
}

void CmdLine::ParseArgs(int argc, const char** argv)
{
    // the form: -argName value -argName value
    for (int i = 1; i < argc; i+=2)
    {
        if (argv[i][0] != '-') // super advanced pattern matching! :)
            throw std::runtime_error("wrong command name");

        mParsedArgs[argv[i]+1] = TryParseString(argv[i+1]);
    }
}

```

At the moment of writing, `std::from_chars` in GCC/Clang only supports integers. MSVC 2017 15.8 has full support also for floating point numbers. You can read more about `from_chars` in the separate [String Conversions Chapter](#).

The idea of `TryParseString` is to try parsing the input string into the best matching type. So if it looks like an integer, then we try to fetch an integer. Otherwise, we'll return an unparsed string. Of course, we can extend this approach.

After the parsing is complete the client can use `Find()` method to test for existence of a parameter:

```
std::optional<CmdLine::Arg> CmdLine::Find(const std::string& name) const
{
    if (const auto it = mParsedArgs.find(name); it != mParsedArgs.end())
        return it->second;

    return { };
}
```

Find() uses std::optional to return the value. If the argument cannot be found in the map, then the client will get empty optional.

Example of how we can use it:

```
// Chapter Variant/variant_parsing_int_float.cpp
try
{
    CmdLine cmdLine(argc, argv);

    if (auto arg = cmdLine.Find("paramInt"); arg)
        std::cout << "paramInt is " << std::get<int>(*arg) << '\n';

    if (auto arg = cmdLine.Find("paramFloat"); arg)
    {
        if (const auto intPtr = std::get_if<int>(&*arg); intPtr)
            std::cout << "paramFloat is " << *intPtr << " (integer)\n";
        else
            std::cout << "paramFloat is " << std::get<float>(*arg) << '\n';
    }

    if (auto arg = cmdLine.Find("paramText"); arg)
        std::cout << "paramText is " << std::get<std::string>(*arg) << '\n';
}
catch (std::bad_variant_access& err)
{
    std::cerr << " ...err: accessing a wrong variant type, "
              << err.what() << '\n';
}
catch (std::runtime_error &err)
{
    std::cerr << " ...err: " << err.what() << '\n';
}
```

The above example uses cmdLine.Find() to check if there's a given parameter. It returns

`std::optional` so we have to check if it's not empty. When we're sure the parameter is available, we can check its type.

`CmdLine` tries to find the best matching type, so for example, with floats, we have ambiguity - as 90 is also `float` but the code will store it as `int` (as it doesn't have the fraction part).

To solve such ambiguities we could pass some additional information about the desired type, or provide some helper methods.


Parsing a Config File

The idea comes from the previous example of a command line. In the case of a configuration file, we usually work with pairs of `<Name, Value>`. Where `Value` might be a different type: `string`, `int`, `array`, `bool`, `float`, etc.

For such a use case even `void*` could be used to hold such an unknown type. However, this pattern is extremely error-prone. We could improve the design by using `std::variant` if we know all the possible types, or leverage `std::any`.

State Machines

How about modelling a state machine? For example door's state:



Door State Machine

We can use different types of states and the use visitors as events:

```
// Chapter Variant/variant_fsm.cpp
struct DoorState
{
    struct DoorOpened {};
    struct DoorClosed {};
    struct DoorLocked {};

    using State = std::variant<DoorOpened, DoorClosed, DoorLocked>;

    void open()
    {
        m_state = std::visit(OpenEvent{}, m_state);
    }

    void close()
    {
        m_state = std::visit(CloseEvent{}, m_state);
    }

    void lock()
    {
        m_state = std::visit(LockEvent{}, m_state);
    }

    void unlock()
    {
        m_state = std::visit(UnlockEvent{}, m_state);
    }

    State m_state;
};
```

And here are the events:

```
// Chapter Variant/variant_fsm.cpp
struct OpenEvent
{
    State operator()(const DoorOpened&){ return DoorOpened(); }
    State operator()(const DoorClosed&){ return DoorOpened(); }
    // cannot open locked doors
    State operator()(const DoorLocked&){ return DoorLocked(); }
};

struct CloseEvent
{
    State operator()(const DoorOpened&){ return DoorClosed(); }
    State operator()(const DoorClosed&){ return DoorClosed(); }
    State operator()(const DoorLocked&){ return DoorLocked(); }
};

struct LockEvent
{
    // cannot lock opened doors
    State operator()(const DoorOpened&){ return DoorOpened(); }
    State operator()(const DoorClosed&){ return DoorLocked(); }
    State operator()(const DoorLocked&){ return DoorLocked(); }
};

struct UnlockEvent
{
    // cannot unlock opened doors
    State operator()(const DoorOpened&){ return DoorOpened(); }
    State operator()(const DoorClosed&){ return DoorClosed(); }
    // unlock
    State operator()(const DoorLocked&){ return DoorClosed(); }
};
```

Polymorphism

Most of the time in C++ we can safely use runtime polymorphism based on a vtable approach. You have a collection of related types that share the same interface, and you have a well defined virtual method that can be invoked.

But what if you have “unrelated” types that don’t share the same base class? What if you’d like to quickly add new functionality without changing the code of the supported types?

With `std::variant` and `std::visit` we can build the following example:

```
// Chapter Variant/variant_polymorphism.cpp
class Triangle
{
public:
    void Render() { std::cout << "Drawing a triangle!\n"; }
};

class Polygon
{
public:
    void Render() { std::cout << "Drawing a polygon!\n"; }
};

class Sphere
{
public:
    void Render() { std::cout << "Drawing a sphere!\n"; }
};

int main()
{
    std::vector<std::variant<Triangle, Polygon, Sphere>> objects {
        Polygon(),
        Triangle(),
        Sphere(),
        Triangle()
    };

    auto CallRender = [](auto& obj) { obj.Render(); };

    for (auto& obj : objects)
        std::visit(CallRender, obj);
}
```

The above example shows only the first case of invoking a method from unrelated types. It wraps all the possible shape types into a single variant and then uses a visitor to dispatch the call to the proper type.

If you'd like, for example, to sort objects, then you can write another visitor, one that holds some state. And that way you'll get more functionality without changing the types.

Wrap Up

Here are the things to remember about `std::variant`:

- It holds one of several alternatives in a type-safe way
- No extra memory allocation is needed. The variant needs the size of the max of the sizes of the alternatives, plus some little extra space for knowing the currently active value
- By default, it initialises with the default value of the first alternative
- You can access the value by using `std::get`, `std::get_if` or by using a form of a visitor
- To check the currently active type you can use `std::holds_alternative` or `std::variant::index`
- `std::visit` is a way to invoke an operation on the currently active type in the variant. It's a callable object with overloads for all the possible types in the variant(s)
- Rarely `std::variant` might get into invalid state, you can check it via `valueless_by_exception`

Compiler Support

Feature	GCC	Clang	MSVC
<code>std::variant</code>	7.1	4.0	VS 2017 15.0

8. `std::any`

With `std::optional` you can represent some Type or nothing. With `std::variant` you can wrap several variants into one entity. And C++17 gives us one more wrapper type: `std::any` which can hold anything in a type-safe way.

In this chapter you'll learn:

- Why `void*` is a very unsafe pattern
- `std::any` and its basic usage
- `std::any` use cases with examples
- `any_cast` and how to use all its “modes”

The Basics

In C++14 there weren't many options for holding variable types in a variable. You could use `void*`, of course, but this wasn't safe. `'void*` is just a raw pointer, and you have to manage the whole object lifetime and protect it from casting to a different type.

Potentially, `void*` could be wrapped in a class with some type discriminator.

```
class MyAny
{
    void* _value;
    TypeInfo _typeInfo;
};
```

As you see, we have some basic form of the type, but there's a bit of coding required to make sure `MyAny` is type-safe. That's why it's best to use the Standard Library rather than rolling a custom implementation.

And this is what `std::any` from C++17 is in its basic form. It lets you store anything in an object, and it reports errors (or throw exceptions) when you'd like to access a type that is not active.

A little demo:

```
// Chapter Any/any_demo.cpp
std::any a(12);

// set any value:
a = std::string("Hello!");
a = 16;
// reading a value:

// we can read it as int
std::cout << std::any_cast<int>(a) << '\n';

// but not as string:
try
{
    std::cout << std::any_cast<std::string>(a) << '\n';
}
catch(const std::bad_any_cast& e)
{
    std::cout << e.what() << '\n';
}
```

```

// reset and check if it contains any value:
a.reset();
if (!a.has_value())
{
    std::cout << "a is empty!" << '\n';
}

// you can use it in a container:
std::map<std::string, std::any> m;
m["integer"] = 10;
m["string"] = std::string("Hello World");
m["float"] = 1.0f;

for (auto &[key, val] : m)
{
    if (val.type() == typeid(int))
        std::cout << "int: " << std::any_cast<int>(val) << '\n';
    else if (val.type() == typeid(std::string))
        std::cout << "string: " << std::any_cast<std::string>(val) << '\n';
    else if (val.type() == typeid(float))
        std::cout << "float: " << std::any_cast<float>(val) << '\n';
}

```

The code will output:

```

16
bad any_cast
a is empty!
float: 1
int: 10
string: Hello World

```

The example above shows us several things:

- `std::any` is not a template class like `std::optional` or `std::variant`.
- by default it contains no value, and you can check it via `.has_value()`.
- you can reset an any object via `.reset()`.
- it works on “decayed” types - so before assignment, initialisation, or emplacement the type is transformed by `std::decay`¹.
- when a different type is assigned, then the active type is destroyed.

¹<http://en.cppreference.com/w/cpp/types/decay>

- you can access the value by using `std::any_cast<T>`. It will throw `bad_any_cast` if the active type is not `T`.
- you can discover the active type by using `.type()` that returns [std::type_info](http://en.cppreference.com/w/cpp/types/type_info)² of the type.

When to Use

While `void*` might be an extremely unsafe pattern with some limited use cases, `std::any` adds type-safety, and that's why it has more applications.

Some possibilities:

- In Libraries - when a library type has to hold or pass anything without knowing the set of available types
- Parsing files - if you really cannot specify what the supported types are
- Message passing
- Bindings with a scripting language
- Implementing an interpreter for a scripting language
- User Interface - controls might hold anything
- Entities in an editor

In many cases, you can limit the number of supported types, and that's why `std::variant` might be a better choice. Of course, it gets tricky when you implement a library without knowing the final applications - so you don't know the possible types that will be stored in an object.

std::any Creation

There are several ways you can create `std::any` object:

- a default initialisation - then the object is empty
- a direct initialisation with a value/object
- in place `std::in_place_type`
- via `std::make_any`

You can see it in the following example:

²http://en.cppreference.com/w/cpp/types/type_info

```
// Chapter Any/any_creation.cpp

// default initialisation:
std::any a;
assert(!a.has_value());

// initialisation with an object:
std::any a2{10}; // int
std::any a3{MyType{10, 11}};

// in_place:
std::any a4{std::in_place_type<MyType>, 10, 11};
std::any a5{std::in_place_type<std::string>, "Hello World"};

// make_any
std::any a6 = std::make_any<std::string>{"Hello World"};
```

In Place Construction

Following the style of `std::optional` and `std::variant`, `std::any` can use `std::in_place_type` to efficiently create objects in place.

Complex Types

In the below example a temporary object will be needed:

```
std::any a{UserName{"hello"}};
```

but with:

```
std::any a{std::in_place_type<UserName>, "hello"};
```

The object is created in place with the given set of arguments.

std::make_any

For convenience `std::any` has a factory function called `std::make_any` that returns

```
return std::any(std::in_place_type<T>, std::forward<Args>(args)...);
```

So in the previous example we could also write:

```
auto a = std::make_any<UserName>{"hello"};
```

make_any is probably more straightforward to use.

Changing the Value

When you want to change the currently stored value in std::any then you have two options: use `emplace` or the assignment:

```
// Chapter Any/any_changing.cpp
```

```
std::any a;

a = MyType(10, 11);
a = std::string("Hello");

a.emplace<float>(100.5f);
a.emplace<std::vector<int>>>({10, 11, 12, 13});
a.emplace<MyType>(10, 11);
```

Object Lifetime

The crucial part of being safe for std::any is not to leak any resources. To achieve this behaviour std::any will destroy any active object before assigning a new value.

```
// Chapter Any/any_lifetime.cpp
std::any var = std::make_any<MyType>();
var = 100.0f;
std::cout << std::any_cast<float>(var) << '\n';
```

If the constructors and destructors were instrumented with prints, we would get the following output:

```
MyType::MyType
MyType::~MyType
100
```

The any object is initialised with MyType, but before it gets a new value (of 100.0f) it calls the destructor of MyType.

Accessing The Stored Value

To read the currently active value in `std::any` you have mostly one option - `template<class ValueType> ValueType std::any_cast`. This function returns the value of the requested type if it's in the object.

However, this function template is quite powerful, as it has three modes:

- read access - returns a copy of the value, and throws `std::bad_any_cast` when it fails
- read/write access - returns a reference, and throws `std::bad_any_cast` when it fails
- read/write access - returns a pointer to the value (const or not) or `nullptr` on failure

See the example:

// Chapter Any/any_access.cpp

```
struct MyType
{
    int a, b;

    MyType(int x, int y) : a(x), b(y) { }

    void Print() { std::cout << a << ", " << b << '\n'; }
};

int main()
{
    std::any var = std::make_any<MyType>(10, 10);
    try
    {
        std::any_cast<MyType&>(var).Print();
        std::any_cast<MyType&>(var).a = 11; // read/write
        std::any_cast<MyType&>(var).Print();
        std::any_cast<int>(var); // throw!
    }
    catch(const std::bad_any_cast& e)
    {
        std::cout << e.what() << '\n';
    }

    int* p = std::any_cast<int>(&var);
    std::cout << (p ? "contains int... \n" : "doesn't contain an int...\n");
}
```



```
    if (MyType* pt = std::any_cast<MyType>(&var); pt)
    {
        pt->a = 12;
        std::any_cast<MyType&>(var).Print();
    }
}
```

As you see, there are two options regarding error handling: via exceptions (`std::bad_any_cast`) or by returning a pointer (or `nullptr`). The function overloads for `std::any_cast` pointer access are also marked with `noexcept`.

Performance & Memory Considerations

`std::any` looks quite powerful, and you might use it to hold variables of variable types... but you might ask what the price is for such flexibility.

The main issue: extra dynamic memory allocations.

`std::variant` and `std::optional` don't require any extra memory allocations but this is because they know which type (or types) will be stored in the object. `std::any` does not know which types might be stored and that's why it might use some heap memory.

Will it always happen, or sometimes? What are the rules? Will it happen even for a simple type like `int`?

Let's see what the standard says:

From The Standard:

Implementations should avoid the use of dynamically allocated memory for a small contained value. Example: where the object constructed is holding only an `int`. Such small-object optimisation shall only be applied to types `T` for which `is_nothrow_move_constructible_v<T>` is `true`.

To sum up - Implementations are encouraged to use SBO - **Small Buffer Optimisation**. But that also comes at some cost: it will make the type larger - to fit the buffer.

Let's check the size of `std::any`:

Here are the results from the three compilers:

Compiler	sizeof(any)
GCC 8.1 (Coliru)	16
Clang 7.0.0 (Wandbox)	32
MSVC 2017 15.7.0 32-bit	40
MSVC 2017 15.7.0 64-bit	64

In general, as you see, `std::any` is not a “simple” type and it brings a lot of overhead with it. It’s usually not small - due to SBO - it takes 16 or 32 bytes (GCC, Clang, or even 64 bytes in MSVC).

You can see the code in Chapter `Any/any_sizeof.cpp`.

Migration from boost::any

Boost Any was introduced around the year 2001 (Version 1.23.0). Interestingly, the author of the boost library - Kevlin Henney - is also the author of the proposal for `std::any`. So the two types are strongly connected, and the STL version is heavily based on the predecessor.

Here are the main changes:

Feature	Boost.Any (1.67.0) ³	std::any
Extra memory allocation	Yes	Yes
Small buffer optimisation	Yes	Yes
emplace	No	Yes
in_place_type_t in constructor	No	Yes

There are not many differences between the two types. Most of the time you can easily convert from `boost.any` into the STL version.

Examples of std::any

The core of `std::any` is flexibility. In the below examples, you can see some ideas (or concrete implementations) where holding variable type can make an application a bit simpler.

Parsing files

In the examples for `std::variant` you can see how it’s possible to parse configuration files and store the result as an alternative of several types. If you write an entirely generic solution - for example as a part of some library, then you might not know all the possible types.

Storing `std::any` as a value for a property might be good enough from the performance point of view and will give you flexibility.

Message Passing

In Windows API, which is C mostly, there's a message passing system that uses message ids with two optional parameters that store the value of the message. Based on that mechanism you can implement WndProc to handle the messages passed to your window/control:

```
LRESULT CALLBACK WindowProc(
    _In_ HWND    hwnd,
    _In_ UINT    uMsg,
    _In_ WPARAM  wParam,
    _In_ LPARAM  lParam
);
```

The trick here is that the values are stored in wParam or lParam in various forms. Sometimes you have to use only a few bytes of wParam...

What if we changed this system into std::any, so that a message could pass anything to the handling method?

For example:

```
// Chapter Any/any_winapi.cpp
class Message
{
public:
    enum class Type
    {
        Init,
        Closing,
        ShowWindow,
        DrawWindow
    };

public:
    explicit Message(Type type, std::any param) :
        mType(type),
        mParam(param)
    { }
    explicit Message(Type type) :
        mType(type)
    { }

    Type mType;
    std::any mParam;
```

```
};

class Window
{
public:
    virtual void HandleMessage(const Message& msg) = 0;
};
```

For example you can send a message to a window:

```
Message m(Message::Type::ShowWindow, std::make_pair(10, 11));
yourWindow.HandleMessage(m);
```

Then the window can respond to the message with the following message handler:

```
switch (msg.mType) {
// ...
case Message::Type::ShowWindow:
{
    auto pos = std::any_cast<std::pair<int, int>>(msg.mParam);
    std::cout << "ShowWindow: "
                << pos.first << ", "
                << pos.second << '\n';

    break;
}
}
```

Of course, you have to define how the values are specified (what the types of a value of a message are), but now you can use real types rather than doing various tricks with integers.

Properties

The original paper that introduces `any` to C++, [N1939](http://www.open-std.org/jtc1/sc22/wg21/docs/papers/2006/n1939.html)⁴ shows an example of a property class.

⁴<http://www.open-std.org/jtc1/sc22/wg21/docs/papers/2006/n1939.html>

```
struct property
{
    property();
    property(const std::string &, const std::any &);

    std::string name;
    std::any value;
};

typedef std::vector<property> properties;
```

The `properties` object looks quite powerful as it can hold many different types. One of the examples where such structure might be leveraged is a game editor.

Wrap Up

Here are the things to remember about `std::any`:

- `std::any` is not a template class
- `std::any` uses Small Buffer Optimisation, so it will not dynamically allocate memory for simple types like ints, doubles... but for larger types, it will use extra `new`.
- `std::any` might be considered ‘heavy’, but offers a lot of flexibility and type-safety.
- you can access the currently stored value by using `any_cast` that offers a few “modes”: for example it might throw an exception or return `nullptr`.
- use it when you don’t know the possible types - in other cases consider `std::variant`.

Compiler Support

Feature	GCC	Clang	MSVC
<code>std::any</code>	7.1	4.0	VS 2017 15.0

9. `std::string_view`

Since C++11 and move semantics, passing strings has become much faster. Yet you can end up with many temporary string copies. In C++17 you get a new type called `string_view`. It allows you to create a constant, non-owning *view* of a contiguous character sequence. You can manipulate that view and pass it around without the need to copy the referenced data. Nevertheless, the feature comes at some cost: you need to be careful not to end up with “dangling” views, and usually such views might not be null-terminated.

In this chapter you’ll learn:

- What is `string_view`?
- Why might it speed up your code?
- What are the risks involved with using `string_view` objects?
- What is the reference lifetime extension and what does it mean for `string_view`?
- How you can use `string_view` to make your API more generic?

The Basics

Let's try a little experiment:

How many string copies are created in the below example?

```
// string function:
std::string StartFromWordStr(const std::string& strArg, const std::string& word)
{
    return strArg.substr(strArg.find(word)); // substr creates a new string
}

// call:
std::string str {"Hello Amazing Programming Environment" };

auto subStr = StartFromWordStr(str, "Programming Environment");
std::cout << subStr << '\n';
```

Can you count them all?

The answer is 3 or 5 depending on the compiler, but usually, it should be 3.

- The first one is for `str`.
- The second one is for the second argument in `StartFromWordStr` - the argument is `const string&` so since we pass `const char*` it will create a new string.
- The third one comes from `substr` which returns a new string.
- Then we might also have another copy or two - as the object is returned from the function. But usually, the compiler can optimise and elide the copies (especially since C++17 when Copy Elision became mandatory in that case).
- If the string is short, then there might be no heap allocation as Small String Optimisation¹.

The above example is simplistic. However, you might imagine a production code where string manipulations happen very often. In that scenario, it's even hard to count all the temporaries that the compiler creates.

A much better pattern to solve the problem with temporary copies is to use `std::string_view`. As the name suggests, instead of using the original string, you'll only get a non-owning view of it. Most of the time it will be a pointer to the contiguous character sequence and the length. You can pass it around and use most of the conventional string operations.

Views work well with string operations like substring - `substr`. In a typical case, each substring operation creates another, smaller copy of the string. With `string_view`, `substr` will only map a different portion of the original buffer, without additional memory usage, or dynamic allocation.

Here's the updated version of our code that accepts `string_view`:

¹Small String Optimisation is not defined in the C++ Standard, but it's a common optimisation across popular compilers. Currently, it's 15 characters in MSVC (VS 2017)/GCC (8.1) or 22 characters in Clang (6.0).

```
// Chapter string_view/avoiding_copies_string_view.cpp

std::string_view StartFromWord(std::string_view str, std::string_view word)
{
    return str.substr(str.find(word)); // substr creates only a new view
}

// call:
std::string str {"Hello Amazing Programming Environment"};

auto subView = StartFromWord(str, "Programming Environment");
std::cout << subView << '\n';
```

In the above case, we have only one allocation - just for the main string - `str`. None of the `string_view` operations invoke copy or extra memory allocation for a new string. Of course, `string_view` is copied - but since it's only a pointer and a length, it's much more efficient than the copy of the whole string.

One warning: while this example shows the optimisation capability of string views, please read on to see the risks and assumptions with that code! Or maybe you can spot a few now?

Ok, so when you should use `string_view`:

When to Use

- Optimisation: you can carefully review your code and replace various string operations with `string_view`. In most cases, you should end up with faster code and fewer memory allocations.
- As a possible replacement for `const std::string&` parameter - especially in functions that don't need the ownership and don't store the string.
- Handling strings coming from other API: `QString`, `CString`, `const char*`... everything that is placed in a contiguous memory chunk and has a basic `char`-type. You can write a function that accepts `string_view` and no conversion from that other implementation will happen.

In any case, it's important to remember that it's only a **non-owning view**, so if the original object is gone, the view becomes rubbish and you might get into trouble.

Moreover, **`string_view` might not contain null terminator** so your code has to support that as well. For example, it's never a good idea to pass `string_view` to a function that accepts null-terminated strings. More on that in a separate section - about "Risks with `string_view`".

The `std::basic_string_view` Type

Although we talk about `string_view` it's important to know that this is only a specialisation of a template class called `basic_string_view`:


```
template<
    class CharT,
    class Traits = std::char_traits<CharT>
> class basic_string_view;
```

Traits class is used to abstract the operations on the character type. For example on how to compare characters, how to find one character in a sequence.

Such a hierarchy creates a similar relation as there is for `std::string` which is a specialisation of `std::basic_string`.

We have the following specialisations:

```
std::string_view      std::basic_string_view<char>
std::wstring_view    std::basic_string_view<wchar_t>
std::u16string_view  std::basic_string_view<char16_t>
std::u32string_view  std::basic_string_view<char32_t>
```

As you can see the specialisations use a different underlying character type.

For the sake of convenience, in the rest of this chapter, only `string_view` will be considered.

std::string_view Creation

You can create a `string_view` in several ways:

- from `const char*` - providing a pointer to a null-terminated string
- from `const char*` with length
- by using a conversion from `std::string`
- by using `"sv` literal

Here's an example of various creation options:

```
// Chapter string_view/string_view_creation.cpp

const char* cstr = "Hello World";

// the whole string:
std::string_view sv1 { cstr };
std::cout << sv1 << ", len: " << sv1.size() << '\n';

// slice
std::string_view sv2 { cstr, 5 }; // not null-terminated!
```

```
std::cout << sv2 << ", len: " << sv2.size() << '\n';

// from string:
std::string str = "Hello String";
std::string_view sv3 = str;
std::cout << sv3 << ", len: " << sv3.size() << '\n';

// "sv literal
using namespace std::literals;
std::string_view sv4 = "Hello\0 Super World"sv;
std::cout << sv4 << ", len: " << sv4.size() << '\n';
std::cout << sv4.data() << " - till zero\n";
```

The code will print:

```
Hello World, len: 11
Hello, len: 5
Hello String, len: 12
Hello Super World, len: 18
Hello - till zero
```

Please notice the last two lines: sv4 contains '\0' in the middle, but std::cout can still print the whole sequence. In the last line, we try to print with .data() and we end up with a string pointer so the printing breaks at the null terminator.

Other Operations

string_view is modelled to be very similar to std::string. The view however, is non-owning, so any operation that modifies the data cannot go into the API. Here's a brief list of methods that you can use with this new type:

Iterators:

- cbegin()/begin() / crbegin()/rbegin() - both are const and constexpr
- cend()/end() / crend()/rend() - both are const and constexpr

Accessing Elements:

- operator[]
- at
- front

- back
- data

Size & Capacity:

- size/length
- max_size
- empty

Modifiers:

- remove_prefix
- remove_suffix
- swap

Other:

- copy (not constexpr)
- substr - complexity $O(1)$ and not $O(n)$ as in `std::string`
- compare
- find
- rfind
- find_first_of
- find_last_of
- find_first_not_of
- find_last_not_of
- operators for lexicographical comparisons: `==`, `!=`, `<=`, `>=`, `<`, `>`
- operator `<<` for ostream output
- specialisation for `std::hash`

Key things about the above operations:

- All of the above methods (except for `copy`, `operator <<` and `std::hash` specialisation) are also `constexpr`! With this capability, you might now work with contiguous character sequences in constant expressions.
- The above list is almost the same as all non-mutable string operations. However there are two new methods: `remove_prefix` and `remove_suffix` - they are not `const` and modify the `string_view` object. Note, that they still cannot modify the referenced data.
- `operator[]`, `at`, `front`, `back`, `data` - are also `const` - so you cannot change the underlying character sequence (it's only "read access"). In `std::string` there are overloads for those methods that return reference - so you get "write access". That's not possible with `string_view`.



More in C++20!

What's more for C++20 we'll get at least two new methods:

- `starts_with`
- `ends_with`

That are implemented both for `std::string_view` and `std::string`. As of now (July 2018) Clang 6.0 supports those functions. You can experiment with them if you like.

Risks Using `string_view`

`std::string_view` was added into the Standard mostly to allow performance optimisations. Nevertheless, it's not a replacement for strings! That's why when you use views you have to remember about a few potentially risky things:

Taking Care of Not Null-Terminated Strings

`string_view` may not contain `\0` at the end of the string. So you have to be prepared for that.

- `string_view` is problematic with *all* functions that accept traditional C-strings because `string_view` breaks with C-string termination assumptions. If a function accepts only a `const char*` parameter, it's probably a bad idea to pass `string_view` into it. On the other hand, it might be safe when such a function accepts `const char*` and `length` parameters.
- Conversion into strings - you need to specify not only the pointer to the contiguous character sequence but also the length.

References and Temporary Objects

`string_view` doesn't own the memory, so you have to be very careful when working with temporary objects.

In general, the lifetime of a `string_view` must never exceed the lifetime of the string-owning object.

That might be important when:

- Returning `string_view` from a function.
- Storing `string_view` in objects or containers.

To explore all the issues, let's start with the initial example from this chapter.

Problems with the Initial Example

The intro section showed you an example of:

```
std::string_view StartFromWord(std::string_view str, std::string_view word)
{
    return str.substr(str.find(word)); // substr creates only a new view
}
```

The code doesn't have any issues with non-null-terminated strings - as all the functions are from the `string_view` API.

However, how about temporary objects?

What will happen if you call:

```
auto str = "My Super"s;
auto sv = StartFromWord(str + " String", "Super");
// use `sv` later in the code...
```

Code like that might blow!

"Super" is a temporary `const char*` literal and it's passed as `string_view` `word` into the function. That's fine, as the temporary is guaranteed to live as long as the whole function invocation.

However, the result of string concatenation `str + " String"` is a temporary and the function returns a `string_view` of this temporary outside the call!

So the general advice in such cases is that while it's possible to return a `string_view` from a function, you have to be careful and be sure about the state of the underlying string.

To understand issues with temporary values, it's good to have a look at the reference lifetime extension.

Reference Lifetime Extension

What happens in the following case:

```
std::vector<int> GenerateVec()
{
    return std::vector<int>(5, 1);
}

const std::vector<int>& refv = GenerateVec();
```

Is the above code safe?

Yes - the C++ rules say that the lifetime of a temporary object bound to a `const` reference is prolonged to the lifetime of the reference itself.

Here's a full example quoted from the standard ([Draft C++17 - N4687](https://wg21.link/n4687)²) 15.2 Temporary objects [class.temporary]:

²<https://wg21.link/n4687>

[Example:

```

struct S {
    S();
    S(int);
    friend S operator+(const S&, const S&);
    ~S();
};
S obj1;
const S& cr = S(16)+S(23);
S obj2;

```

The expression `S(16) + S(23)` creates three temporaries: a first temporary T1 to hold the result of the expression `S(16)`, a second temporary T2 to hold the result of the expression `S(23)`, and a third temporary T3 to hold the result of the addition of these two expressions. The temporary T3 is then bound to the reference `cr`. It is unspecified whether T1 or T2 is created first. On an implementation where T1 is created before T2, T2 shall be destroyed before T1. The temporaries T1 and T2 are bound to the reference parameters of `operator+`; these temporaries are destroyed at the end of the full-expression containing the call to `operator+`. The temporary T3 bound to the reference `cr` is destroyed at the end of `cr`'s lifetime, that is, at the end of the program. In addition, the order in which T3 is destroyed takes into account the destruction order of other objects with static storage duration. That is, because `obj1` is constructed before T3, and T3 is constructed before `obj2`, `obj2` shall be destroyed before T3, and T3 shall be destroyed before `obj1`. -end example]

While it's better not to write such code for all of your variables, it might be a handy feature in cases like:

```

for (auto &elem : GenerateVec())
{
    // ...
}

```

In the above example, `GenerateVec` is bound to a reference (rvalue reference for the start of the vector) inside the range-based for loop. Without the extended lifetime support, the code would break.

How does it relate to `string_view`?

For `string_view` the below code is usually error-prone:

```
std::string func()
{
    std::string s;
    // build s...
    return s;
}

std::string_view sv = func();
// no temp lifetime extension!
```

This might be not obvious - `string_view` is also a constant view, so should behave almost like a `const` reference. But according to existing C++ rules, it's not- the compiler immediately destroys the temporary object after the whole expression is done. The lifetime cannot be extended in this case.

`string_view` is just a proxy object, similar to another code:

```
std::vector<int> CreateVector() { ... }
std::string GetString() { return "Hello"; }

auto &x = CreateVector()[10]; // arbitrary element!
auto pStr = GetString().c_str();
```

In both cases `x` and `pStr` won't extend the lifetime of the temporary object created in `CreateVector()` or `GetString()`.

You might fix it by:

```
std::string func()
{
    std::string s;
    // build s...
    return s;
}

auto temp = func();
std::string_view sv { temp };
// fine lifetime of temporary is extended through `temp`
```

Every time you assign a return value from some function you have to be sure the lifetime of the object is correct.



There's a proposal to fix the issues with `string_views` and other types that should have extended reference lifetime semantics: see [P0936](https://wg21.link/p0936)³.

³<https://wg21.link/p0936>

Initializing string Members from string_view

Since `string_view` is a potential replacement for `const string&` when passing in functions, we might consider a case of string member initialisation. Is `string_view` the best candidate here?

For example:

```
class UserName {
    std::string mName;

public:
    UserName(const std::string& str) : mName(str) { }
};
```

As you can see a constructor is taking `const std::string& str`.

You could potentially replace a constant reference with `string_view`:

```
UserName(std::string_view sv) : mName(sv) { }
```

Let's compare those alternatives implementations in three cases: creating from a string literal, creating from an l-value and creating from an rvalue reference:

```
// creation from a string literal
UserName u1{"John With Very Long Name"};

// creation from l-value:
std::string s2 {"Marc With Very Long Name"};
UserName u2 { s2 };
// use s2 later...

// from r-value reference
std::string s3 {"Marc With Very Long Name"};
UserName u3 { std::move(s3) };

// third case is also similar to taking a return value:
std::string GetString() { return "some string..."; }
UserName u4 { GetString() };
```

Now we can analyse two versions of `UserName` constructor - with a string reference or a `string_view`.

Please note that allocations/creation of `s2` and `s3` are not taken into account, we're only looking at what happens for the constructor call. For `s2` we can also assume it's used later in the code.

For `const std::string&`:

- u1 - two allocations: the first one creates a temp string and binds it to the input parameter, and then there's a copy into mName.
- u2 - one allocation: we have a no-cost binding to the reference, and then there's a copy into the member variable.
- u3 - one allocation: we have a no-cost binding to the reference, and then there's a copy into the member variable.
- You'd have to write a ctor taking r-value reference to skip one allocation for the u1 case, and also that could skip one copy for the u3 case (since we could move from r-value reference).

For std::string_view:

- u1 - one allocation - no copy/allocation for the input parameter, there's only one allocation when mName is created.
- u2 - one allocation - there's a cheap creation of a string_view for the argument, and then there's a copy into the member variable.
- u3 - one allocation - there's a cheap creation of a string_view for the argument, and then there's a copy into the member variable.
- You'd also have to write a constructor taking r-value reference if you want to save one allocation in the u3 case, as you could move from r-value reference.

While the string_view behaves better when you pass a string literal, it's no better when you use it with existing string, or you move from it.

However, since the introduction of move semantics in C++11, it's usually better, and safer to pass string as a value and then move from it.

For example:

```
class UserName {
    std::string mName;

public:
    UserName(std::string str) : mName(std::move(str)) { }
};
```

Now we have the following results:

For std::string:

- u1 - one allocation - for the input argument and then one move into the mName. It's better than with const std::string& where we got two memory allocations in that case. And similar to the string_view approach.
- u2 - one allocation - we have to copy the value into the argument, and then we can move from it.

- u3 - no allocations, only two move operations - that's better than with `string_view` and `const string&`!

When you pass `std::string` by value not only is the code simpler, there's also no need to write separate overloads for r-value references.

See the full code sample in

Chapter `string_view/initializing_from_string_view.cpp`

The approach with passing by value is consistent with item 41 - "Consider pass by value for copyable parameters that are cheap to move and always copied" from *Effective Modern C++* by Scott Meyers.

However, is `std::string` cheap to move?

Although the C++ Standard doesn't specify that, usually, strings are implemented with **Small String Optimisation (SSO)** - the string object contains extra space to fit characters without additional memory allocation⁴. That means that moving a string is the same as copying it. And since the string is short, the copy is also fast.

Let's reconsider our example of passing by value when the string is short:

```
UserName u1{"John"}; // fits in SSO buffer

std::string s2 {"Marc"}; // fits in SSO buffer
UserName u2 { s2 };

std::string s3 {"Marc"}; // fits in SSO buffer
UserName u3 { std::move(s3) };
```

Remember that each move is the same as copy.

For `const std::string&`:

- u1 - two copies: one copy from the input string literal into a temporary string argument, then another copy into the member variable.
- u2 - one copy: the existing string is bound to the reference argument, and then we have one copy into the member variable.
- u3 - one copy: the rvalue reference is bound to the input parameter at no cost, later we have a copy into the member field.

For `std::string_view`:

- u1 - one copy: no copy for the input parameter, there's only one copy when `mName` is initialised.

⁴SSO is not standardised and prone to change. Currently, it's 15 characters in MSVC (VS 2017)/GCC (8.1) or 22 characters in Clang (6.0). For multiplatform code, it's not a good idea to assume optimisations based on SSO.

- u2 - one copy: no copy for the input parameter, as `string_view` creation is fast, and then one copy into the member variable.
- u3 - one copy: `string_view` is cheaply created, there's one copy from the argument into `mName`.
- Extra risk that `string_view` might point to a deleted string.

For `std::string`:

- u1 - two copies: the input argument is created from a string literal, and then there's copy into `mName`.
- u2 - two copies: one copy into the argument and then there's the second copy into the member.
- u3 - two copies: one copy into the argument (move means copy) and then there's the second copy into the member.

As you see for short strings passing by value might be “slower” when you pass some existing string, simply because you have two copies rather than one. On the other hand, the compiler might optimise the code better when it sees a value. What's more, short strings are cheap to copy so the potential “slowdown” might not be even visible.

All in all, passing by value and then moving from a string argument is the preferred solution. You have simple code and better performance for larger strings.

As always, if your code needs maximum performance, then you have to measure all of the possible cases.



Other Types & Automation

The problem discussed in this section can also be extended to other copyable and movable types. If the move operation is cheap then passing by value might be better than by reference. You can also use automation, like Clang-Tidy which can detect potential improvements. Clang Tidy has a separate rule for that use case, see [clang-tidy - modernize-pass-by-value⁵](https://clang.llvm.org/extra/clang-tidy/checks/modernize-pass-by-value.html).

Here's the summary of string passing and initialisation of a string member:

input parameter	<code>const string&</code>	<code>string_view</code>	<code>string</code> and <code>move</code>
<code>const char*</code>	2 allocations	1 allocation	1 allocation + move
<code>const char*</code> SSO	2 copies	1 copy	2 copies
lvalue	1 allocation	1 allocation	1 allocation + 1 move
lvalue SSO	1 copy	1 copy	2 copies
rvalue	1 allocation	1 allocation	2 moves
rvalue SSO	1 copy	1 copy	2 copies

⁵<https://clang.llvm.org/extra/clang-tidy/checks/modernize-pass-by-value.html>

Handling Non-Null Terminated Strings

If you get a `string_view` from a `string` then it will point to a null-terminated chunk of memory:

```
std::string s = "Hello World";
std::cout << s.size() << '\n';
std::string_view sv = s;
std::cout << sv.size() << '\n';
```

The two `cout` statements will both print 11.

But what if you have just a part of the string:

```
std::string s = "Hello World";
std::cout << s.size() << '\n';
std::string_view sv = s;
auto sv2 = sv.substr(0, 5);
std::cout << sv2.data() << '\n'; // oops?
```

`sv2` should contain only "Hello", but when you access the pointer to the underlying memory, you'll receive the pointer to the whole string. The expression: `cout << sv2.data()` will print the whole string, and not just a part of it! `sv2.data()` returns the pointer to the "Hello World" character array inside the `string` `s` object.

Of course when you print `sv2` you'll get the correct result:

```
std::cout << sv2 << '\n';
// prints "Hello"
```

This is because `std::cout` handles `string_view` type.

The example shows a potential problem with all third-party APIs that assume null-terminated strings. To name a few:

Printing with `printf()`

For example:

```
std::string s = "Hello World";
std::string_view sv = s;
std::string_view sv2 = sv.substr(0, 5);
printf("My String %s", sv2.data()); // oops?
```

Instead you should use:

```
printf("%.*s\n", static_cast<int>(sv2.size()), sv2.data());
```

.*** - describes the precision, see in the [printf specification](#)⁶:

The precision is not specified in the format string, but as an additional integer value argument preceding the argument that has to be formatted.

Conversion Functions Like `atoi()`/`atof()`:

```
std::string number = "123.456";
std::string_view svNum { number.data(), 3 };
auto f = atof(svNum.data()); // should be 123, but is 123.456!
std::cout << f << '\n';
```

`atof` takes only the pointer to a null-terminated string, so `string_view` is not compatible.

To fix this, you might have a look at `from_chars` functions (also added in C++17)

```
// use from_chars (C++17)
std::string number = "123.456";
std::string_view svNum { number.data(), 3 };
int res = 0;
std::from_chars(svNum.data(), svNum.data()+svNum.size(), res);
std::cout << res << '\n';
```

General Solution

If your API supports only null-terminated strings and you cannot switch to a function that takes additional count or size parameter, then you need to convert a view into the string.

For example:

⁶<http://www.cplusplus.com/reference/cstdio/printf/>

```

void ConvertAndShow(const char *str)
{
    auto f = atof(str);
    std::cout << f << '\n';
}

std::string number = "123.456";
std::string_view svNum { number.data(), 3 };
// ... some code
std::string tempStr { svNum.data(), svNum.size() };
ConvertAndShow(tempStr.c_str());

```

ConvertAndShow only works with null terminated strings, so the only way we have is to create a temporary string tempStr and then pass it to the function.



If you want to create a string object from string_view then remember to use .data() and .size() so that you refer to the correct slice of the underlying character sequence.

You can find the working example in:

Chapter string_view/string_view_null.cpp

Performance & Memory Considerations

The core idea behind adding string_view into the Standard Library was performance and memory consumption. By leveraging string_view, you can efficiently skip the creation of many temporary strings which might boost performance.

Regarding the memory: string_view is usually implemented as [ptr, len] - one pointer and usually size_t to represent the possible size.

That's why you should see the size of it as 8 bytes or 16 bytes (depending on whether the architecture is x86 or x64).

If we consider the std::string type, due to common Small String Optimisations std::string is usually 24 or 32 bytes, so double the size of string_view. If a string is longer than the SSO buffer then std::string allocates memory on the heap. If SSO is not supported (which is rare), then std::string would consist of a pointer to the allocated memory and the size.

Regarding the performance of string operations.

string_view has only a subset of string operations, those that don't modify the referenced character sequence. Functions like find() should offer the same performance as the string counterparts.

On the other hand, substr is just a copy of two elements in string_view, while string will perform a copy of a memory range. The complexity is $O(1)$ vs $O(n)$. That's why if you need to split

a larger string and work with those splices the `string_view` implementation should offer a better speed.

Strings in Constant Expressions

The interesting property of `string_view` is that all of the methods are marked as `constexpr` (except for `copy`, `operator <<` and `std::hash` functions specialised for string views). With this capability, you can work on strings at compile time.

For example:

```
// Chapter string_view/string_view_constexpr.cpp
#include <string_view>

int main()
{
    using namespace std::literals;

    constexpr auto strv = "Hello Programming World"sv;
    constexpr auto strvCut = strv.substr("Hello "sv.size());

    static_assert(strvCut == "Programming World"sv);
    return strvCut.size();
}
```

If you use a modern compiler, like GCC 8.1 with the following options `-std=c++1z -Wall -pedantic -O2`. Then the compiled assembler should be in the following form:

```
main:
    movl    $17, %eax
    ret
```

A similar version of such code, but with `std::string` would generate much more code. Since the example uses long strings, then Small String Optimisation is not possible, and then the compiler must generate code for `new/delete` to manage memory of the strings.

Migration from `boost::string_ref` and `boost::string_view`

As with most of the new types in C++17 `string_view` is also inspired by boost libraries. Marshall Clow implemented `boost::string_ref` in the version 1.53 (February 2012) and then it evolved into `boost::string_view` (added into the version 1.61 - May 2016).

The main difference between `boost::string_ref` and `boost::string_view` is `constexpr` support.

`boost::string_view` implements the same functionality as `std::string_view` and also adds a few new functions:

- `starts_with`
- `ends_with`

See the full header file in [boost/doc/libs/1_67_0/boost/utility/string_view.hpp](https://www.boost.org/doc/libs/1_67_0/boost/utility/string_view.hpp)⁷ And in [Boost utility library](#)⁸

The link to the discussion about `string_ref` being deprecated: “[string_view versus string_ref](#)”⁹.

Examples

Below you can find two examples of using `string_view`.

Working with Different String APIs

An interesting use case for `string_view` is when you use it in code that works with different string implementations.

For example, you might have `CString` from MFC, `const char*` from C-APIs, `QString` from QT, and of course `std::string`.

Rather than creating overloads for different string types, you might leverage `string_view`!

For example:

```
void Process(std::string_view sv) { }
```

If you want to use `Process` with different string implementations, then all you have to do is to create a string view from your type. Most of the string types should easily allow that.

For example:

⁷https://www.boost.org/doc/libs/1_67_0/boost/utility/string_view.hpp

⁸https://www.boost.org/doc/libs/1_67_0/boost/utility/

⁹<https://lists.boost.org/Archives/boost/2017/07/236903.php>


```
// MFC Strings:
CString cstr;
Process(std::string_view{cstr.GetString(), cstr.GetLength()});

// QT Strings:
QString qstr;
Process(std::string_view{qstr.toLatin1().constData()});

// Your implementation:
MySuperString myStr;
// MySuperString::GetData() - returns char*
// MySuperString::Length() - returns length of a string
Process(std::string_view{myStr.GetData(), myStr.Length()});
```

Hypothetically, `Process()` could be implemented as `Process(const char*, int len)`, but with `string_view` the code is more explicit and simpler. Additionally, you have all the available methods of `string_view` and such code is more convenient than C-style.

String Split

`string_view` might be a potential optimisation for string splitting. If you own a large persistent string, you might want to create a list of `string_view` objects that maps words of that larger string. Please note that the code is inspired by the article [by Marco Arena - string_view odi et amo](https://marcoarena.wordpress.com/2017/01/03/string_view-odi-et-am/)¹⁰.

```
// Chapter string_view/string_view_split.cpp

std::vector<std::string_view>
splitSV(std::string_view strv, std::string_view delims = " ")
{
    std::vector<std::string_view> output;
    auto first = strv.begin();

    while (first != strv.end())
    {
        const auto second = std::find_first_of(
            first, std::cend(strv),
            std::cbegin(delims), std::cend(delims));

        if (first != second)
        {
            output.emplace_back(strv.substr(std::distance(strv.begin(), first),
```

¹⁰https://marcoarena.wordpress.com/2017/01/03/string_view-odi-et-am/

```

        std::distance(first, second)));
    }

    if (second == strv.end())
        break;

    first = std::next(second);
}

return output;
}

```

Example use case:

```

const std::string str {
    "Hello    Extra,,, Super, Amazing World"
};

for (const auto& word : splitSV(str, " ,"))
    std::cout << word << '\n';

```

This will print:

```

Hello
Extra
Super
Amazing
World

```

The algorithm iterates over the input `string_view` and finds breaks - characters that match delimiters. Then the code extracts part of that sequence - between the last and the new break. The sub-view is stored in the output vector.

Some notes for the implementation:

- The `string_view` version of the algorithm assumes the input string is persistent and not a temporary object. Be careful with the returned vector of `string_view` as it also points to the input string.
- The instruction `if (first != second)` - protects from adding empty “words”, in a case where there are multiple delimiters next to each other (like double spaces).
- The algorithm uses `std::find_first_of` but it’s also possible to use `string_view::find_first_of`. The member method doesn’t return an iterator, but the position in the string. The

member method showed to be slower than the `std::` version in some tests when the number of delimiters is small¹¹

Wrap Up

Here are the things to remember about `std::string_view`:

- It's a specialisation of `std::basic_string_view<charType, traits<charType>>` - with `charType` equal to `char`.
- It's a non-owning view of a contiguous sequence of characters.
- It might not include null terminator at the end.
- It can be used to optimise code and limit the need for temporary copies of strings.
- It contains most of `std::string` operations that don't change the underlying characters.
- Its operations are also marked as `constexpr`.

But:

- Make sure the underlying sequence of characters is still present!
- While `std::string_view` looks like a constant reference to the string, the language doesn't extend the lifetime of returned temporary objects that are bound to `std::string_view`.
- Always remember to use `stringView.size()` when you build a string from `string_view`. The `size()` method properly marks the end of `string_view`.
- Be careful when you pass `string_view` into functions that accept null-terminated strings unless you're sure your `string_view` contains a null terminator.

Compiler support:

Feature	GCC	Clang	MSVC
<code>std::string_view</code>	7.1	4.0	VS 2017 15.0

¹¹See experiments in [Performance of std::string_view vs std::string from C++17](#) and [Bartek's coding blog: Performance of std::string_view vs std::string from C++17](#) for a separate discussion and other possible improvements.

10. String Conversions

`string_view` is not the only feature that we get in C++17 that relates to strings. While views can reduce the number of temporary copies, there's also another feature that is very handy: conversion utilities. In the new C++ Standard you have two set of functions `from_chars` and `to_chars` that are low level and promises impressive performance improvements.

In this chapter you'll learn:

- Why do we need low-level string conversion routines?
- Why the current options in the Standard Library might not be enough?
- How to use C++17's conversion routines
- What performance gains you can expect from the new routines

Elementary String Conversions

The growing number of data formats like JSON or XML require efficient string processing and manipulation. The maximum performance is especially crucial when such formats are used to communicate over the network where high throughput is the key.

For example, you get the characters in a network packet, you deserialise it (convert strings into numbers), then process the data, and finally it's serialised back to the same file format (numbers into strings) and sent over the network as a response.

The Standard Library had bad luck in those areas, as it's usually perceived to be too slow for such advanced string processing. Often developers prefer custom solutions or third-party libraries.

The situation might change as with C++17 we get two sets of functions: `from_chars` and `to_chars` that allow for low-level string conversions.

The feature was introduced in [P0067](https://wg21.link/P0067)¹

In the original paper there's a useful table that summarises all the current solutions:

facility	shortcomings
<code>sprintf</code>	format string, locale, buffer overrun
<code>snprintf</code>	format string, locale
<code>sscanf</code>	format string, locale
<code>atol</code>	locale, does not signal errors
<code>strtol</code>	locale, ignores whitespace and 0x prefix
<code>strstream</code>	locale, ignores whitespace
<code>stringstream</code>	locale, ignores whitespace, memory allocation
<code>num_put / num_get</code> facets	locale, virtual function
<code>to_string</code>	locale, memory allocation
<code>stoi</code> etc.	locale, memory allocation, ignores whitespace and 0x prefix, exceptions

As you can see above, sometimes converting functions do too much work, which makes the whole processing slower. Moreover, often there's no need for the extra features.

First of all, all of them use "locale". So even if you work with universal strings, you have to pay a little price for localisation support. For example, if you parse numbers from XML or JSON, there's no need to apply current system language, as those formats are interchangeable.

The next issue is error reporting. Some functions might throw an exception, some of them return just a basic value. Exceptions might not only be costly (as throwing might involve extra memory allocations) but often a parsing error is not an exceptional situation - usually, an application is prepared for such cases. Returning a simple value - for example, 0 for `atoi`, 0.0 for `atof` is also limited, as in that case you don't know if the parsing was successful or not.

Another problem, especially with older C-style API, is that you have to provide some form of the "format string". Parsing such string might involve some additional cost.

¹<https://wg21.link/P0067>

Another thing is “empty space” support. Functions like `strtol` or `stringstream` might skip empty spaces at the beginning of the string. That might be handy, but sometimes you don’t want to pay for that extra feature.

There’s also another critical factor: safety. Simple functions don’t offer any buffer overrun solutions, and also they work only on null-terminated strings. In that case, you cannot use `string_view` to pass the data.

The new C++17 API addresses all of the above issues. Rather than providing many functionalities, they focus on giving the very low-level support. That way you can have the maximum speed and tailor them to your needs.

The new functions are guaranteed to be:

- non-throwing - in case of some error they won’t throw any exception (as `stoi`)
- non-allocating - the whole processing is done in place, without any extra memory allocation
- no locale support - the string is parsed universally as if used with default (“C”) locale
- memory safety - input and output range are specified to allow for buffer overrun checks
- no need to pass string formats of the numbers
- error reporting additional information about the conversion outcome

All in all, with C++17 you have two set of functions:

- `from_chars` - for conversion from strings into numbers, integer and floating points.
- `to_chars` - for converting numbers into string.

Let’s have a look at the functions in a bit more detail.

Converting From Characters to Numbers: `from_chars`

`from_chars` is a set of overloaded functions: for integral types and floating point types.

For integral types we have the following functions:

```
std::from_chars_result from_chars(const char* first,
                                const char* last,
                                TYPE &value,
                                int base = 10);
```

Where `TYPE` expands to all available signed and unsigned integer types and `char`.

`base` can be a number ranging from 2 to 36.

Then there’s the floating point version:

```
std::from_chars_result from_chars(const char* first,
                                const char* last,
                                FLOAT_TYPE& value,
                                std::chars_format fmt = std::chars_format::general);
```

FLOAT_TYPE expands to float, double or long double.

chars_format is an enum with the following values:

```
enum class chars_format {
    scientific = /*unspecified*/,
    fixed = /*unspecified*/,
    hex = /*unspecified*/,
    general = fixed | scientific
};
```

It's a bit-mask type, that's why the values for enums are implementation specific. By default, the format is set to be general so the input string can use "normal" floating point format with scientific form as well.

The return value in all of those functions (for integers and floats) is from_chars_result :

```
struct from_chars_result {
    const char* ptr;
    std::errc ec;
};
```

from_chars_result holds valuable information about the conversion process.

Here's the summary:

- On **Success** from_chars_result::ptr points at the first character not matching the pattern, or has the value equal to last if all characters match and from_chars_result::ec is value-initialized.
- On **Invalid conversion** from_chars_result::ptr equals first and from_chars_result::ec equals std::errc::invalid_argument. value is unmodified.
- On **Out of range** - The number is too large to fit into the value type. from_chars_result::ec equals std::errc::result_out_of_range and from_chars_result::ptr points at the first character not matching the pattern. value is unmodified.

Examples

To sum up, here are two examples of how to convert a string into number using from_chars, to int and float.

Integral types

```
// Chapter String Conversions/from_chars_basic.cpp
#include <charconv> // from_char, to_char
#include <string>

int main()
{
    const std::string str { "12345678901234" };
    int value = 0;
    const auto res = std::from_chars(str.data(),
                                    str.data() + str.size(),
                                    value);

    if (res.ec == std::errc())
    {
        std::cout << "value: " << value
                  << ", distance: " << res.ptr - str.data() << '\n';
    }
    else if (res.ec == std::errc::invalid_argument)
    {
        std::cout << "invalid argument!\n";
    }
    else if (res.ec == std::errc::result_out_of_range)
    {
        std::cout << "out of range! res.ptr distance: "
                  << res.ptr - str.data() << '\n';
    }
}
```

The example is straightforward, it passes a string `str` into `from_chars` and then displays the result with additional information if possible.

Below you can find an output for various `str` value.

str value	output
12345	value: 12345, distance 5
12345678901234	out of range! res.ptr distance: 14
hfhfyt	invalid argument!

In the case of 12345678901234 the conversion routine could parse the number, but it's too large to fit in `int`.

Floating Point

To get the floating point test, we can replace the top lines of the previous example with:


```
// Chapter String Conversions/from_chars_basic_float.cpp
```

```
const std::string str { "16.78" };
double value = 0;
const auto format = std::chars_format::general;
const auto res = std::from_chars(str.data(),
                                str.data() + str.size(),
                                value,
                                format);
```

Here's the example output that we get:

str value	format value	output
1.01	fixed	value: 1.01, distance: 4
-67.90000	fixed	value: -67.9, distance: 9
1e+10	fixed	value: 1, distance: 1 - scientific notation not supported
1e+10	fixed	value: 1, distance: 1 - scientific notation not supported
20.9	scientific	invalid argument!, res.p distance: 0
20.9e+0	scientific	value: 20.9, distance: 7
-20.9e+1	scientific	value: -209, distance: 8
F.F	hex	value: 15.9375, distance: 3
-10.1	hex	value: -16.0625, distance: 5

The `general` format is a combination of `fixed` and `scientific` so it handles regular floating point string with the additional support for `e+num` syntax.

You have a basic understanding about converting from strings to numbers, so let's have a look at how to do it the opposite way.

Parsing a Command Line

In the `std::variant` chapter there's an example with parsing command line parameters. The example uses `from_chars` to match the best type: `int`, `float` or `std::string` and then stores it in a `std::variant`.

You can find the example here: [Parsing a Command Line, the Variant Chapter](#)

Converting Numbers into Characters: `to_chars`

`to_chars` is a set of overloaded functions for integral and floating point types.

For integral types there's one declaration:

```
std::to_chars_result to_chars(char* first, char* last,
                             TYPE value, int base = 10);
```

Where TYPE expands to all available signed and unsigned integer types and char.

Since base might range from 2 to 36, the output digits that are greater than 9 are represented as lowercase letters: a...z.

For floating point numbers there are more options.

Firstly there's a basic function:

```
std::to_chars_result to_chars(char* first, char* last, FLOAT_TYPE value);
```

FLOAT_TYPE expands to float, double or long double.

The conversion works the same as with printf and in default ("C") locale. It uses %f or %e format specifier favouring the representation that is the shortest.

The next function adds std::chars_format fmt that let's you specify the output format:

```
std::to_chars_result to_chars(char* first, char* last,
                             FLOAT_TYPE value,
                             std::chars_format fmt);
```

Then there's the "full" version that allows also to specify precision:

```
std::to_chars_result to_chars(char* first, char* last,
                             FLOAT_TYPE value,
                             std::chars_format fmt,
                             int precision);
```

When the conversion is successful, the range [first, last) is filled with the converted string.

The returned value for all functions (for integer and floating point support) is to_chars_result, it's defined as follows:

```
struct to_chars_result {
    char* ptr;
    std::errc ec;
};
```

The type holds information about the conversion process:

- On **Success** - ec equals value-initialized std::errc and ptr is the one-past-the-end pointer of the characters written. Note that the string is not NULL-terminated.

- On **Error** - ptr equals first and ec equals `std::errc::invalid_argument`. value is unmodified.
- On **Out of range** - ec equals `std::errc::value_too_large` the range [first, last) in unspecified state.

An Example

To sum up, here's a basic demo of `to_chars`.



At the time of writing there was no support for floating point overloads, so the example uses only integers.

```
// Chapter String Conversions/to_chars_basic.cpp

#include <iostream>
#include <charconv> // from_chars, to_chars
#include <string>

int main()
{
    std::string str { "xxxxxxx" };
    const int value = 1986;

    const auto res = std::to_chars(str.data(),
                                   str.data() + str.size(),
                                   value);

    if (res.ec == std::errc())
    {
        std::cout << str << ", filled: "
                  << res.ptr - str.data() << " characters\n";
    }
    else
    {
        std::cout << "value too large!\n";
    }
}
```

Below you can find a sample output for a set of numbers:

value	output
1986	1986xxxx, filled: 4 characters
-1986	-1986xxx, filled: 5 characters
19861986	19861986, filled: 8 characters
-19861986	value too large! (the buffer is only 8 characters)

The Benchmark

So far the chapter has mentioned huge performance potential of the new routines. It would be best to see some real numbers then!

This section introduces a benchmark that measures performance of `from_chars` and `to_chars` against other conversion methods.

How does the benchmark work:

- Generates vector of random integers of the size `VECSIZE`.
- Each pair of conversion methods will transform the input vector of integers into a vector of strings and then back to another vector of integers. This round-trip will be verified so that the output vector is the same as the input vector.
- The conversion is performed `ITER` times.
- Errors from the conversion functions are not checked.
- The code tests:
 - `from_char/to_chars`
 - `to_string/stoi`
 - `sprintf/atoi`
 - `ostringstream/istringstream`

You can find the full benchmark code in:

“Chapter String Conversions/conversion_benchmark.cpp”

Here’s the code for `from_chars/to_chars`:

```

const auto numIntVec = GenRandVecOfNumbers(vecSize);
std::vector<std::string> numStrVec(numIntVec.size());
std::vector<int> numBackIntVec(numIntVec.size());

std::string strTmp(15, ' ');

RunAndMeasure("to_chars", [&]() {
    for (size_t iter = 0; iter < ITERS; ++iter)
    {
        for (size_t i = 0; i < numIntVec.size(); ++i)
        {
            const auto res = std::to_chars(strTmp.data(),
                                         strTmp.data() + strTmp.size(),
                                         numIntVec[i]);
            numStrVec[i] = std::string_view(strTmp.data(),
                                         res.ptr - strTmp.data());
        }
    }
    return numStrVec.size();
});

RunAndMeasure("from_chars", [&]() {
    for (size_t iter = 0; iter < ITERS; ++iter)

```

```

{
    for (size_t i = 0; i < numStrVec.size(); ++i)
    {
        std::from_chars(numStrVec[i].data(),
                        numStrVec[i].data() + numStrVec[i].size(),
                        numBackIntVec[i]);
    }
}
return numBackIntVec.size();
});

```

CheckVectors(numIntVec, numBackIntVec);

CheckVectors - checks if the two input vectors of integers contains the same values, and prints mismatches on error.



The benchmark converts `vector<int>` into `vector<string>` and we measure the whole conversion process which also includes the string object creation.

Here are the results of running 1000 iterations on a vector with 1000 elements:

Method	GCC 8.2	Clang 7.0 Win	VS 2017 15.8 x64
to_chars	21.94	18.15	24.81
from_chars	15.96	12.74	13.43
to_string	61.84	16.62	20.91
stoi	70.81	45.75	42.40
sprintf	56.85	124.72	131.03
atoi	35.90	34.81	32.50
ostringstream	264.29	681.29	575.95
stringstream	306.17	789.04	664.90

Time is in milliseconds.

The machine: Windows 10 x64, i7 8700 3.2 GHz base frequency, 6 cores/12 threads (although the benchmark uses only one thread for processing).

- GCC 8.2 - compiled with `-std=c++17 -O2 -Wall -pedantic`, [MinGW Distro](https://nuwen.net/mingw.html)²
- Clang 7.0 - compiled with `-std=c++17 -O2 -Wall -pedantic`, [Clang For Windows](http://releases.llvm.org/dow4)³
- Visual Studio 2017 15.8 - Release mode, x64

²<https://nuwen.net/mingw.html>

³<http://releases.llvm.org/dow4>

Some notes:


- On GCC `to_chars` is almost 3x faster than `to_string`, 2.6x faster than `sprintf` and 12x faster than `ostringstream`!
- On Clang `to_chars` is actually a bit slower than `to_string`, but $\sim 7x$ faster than `sprintf` and surprisingly almost 40x faster than `ostringstream`!
- MSVC has also slower performance in comparison with `to_string`, but then `to_chars` is $\sim 5x$ faster than `sprintf` and $\sim 23x$ faster than `ostringstream`

Looking now at `from_chars` :


- On GCC it's $\sim 4,5x$ faster than `stoi`, 2,2x faster than `atoi` and almost 20x faster than `istringstream`.
- On Clang it's $\sim 3,5x$ faster than `stoi`, 2,7x faster than `atoi` and 60x faster than `istringstream`!
- MSVC performs $\sim 3x$ faster than `stoi`, $\sim 2,5x$ faster than `atoi` and almost 50x faster than `istringstream`!

As mentioned earlier, the benchmark also includes the cost of string object creation. That's why `to_string` (optimised for strings) might perform a bit better than `to_chars`. If you already have a char buffer, and you don't need to create a string object, then `to_chars` should be faster.

Here are the two charts built from the table above.



Strings into Numbers, time in milliseconds



Numbers into Strings, time in milliseconds



As always it's encouraged to run the benchmarks on your own before you make the final judgment. You might get different results in your environment, where maybe a different compiler or STL library implementation is available.

Summary

This chapter showed how to use two sets of functions `from_chars` - to convert strings into numbers, and `to_chars` that converts numbers into their textual representations.

The functions might look very raw and even C-style. This is a “price” you have to pay for having such low-level support, performance, safety and flexibility. The advantage is that you can provide a simple wrapper that exposes only the needed parts that you want.



Extra Info

The change was proposed in: [P0067](https://wg21.link/P0067)⁴.

⁴<https://wg21.link/P0067>

Compiler support

Feature	GCC	Clang	MSVC
Elementary String Conversions	8.0 ⁵	7.0 ⁶	VS 2017 15.7/15.8 ⁷

⁵In progress, only integral types are supported

⁶In progress, only integral types are supported

⁷Integer support for `from_chars/to_chars` available in 15.7, floating point support for `from_chars` ready in 15.8. Floating point `to_chars` should be ready with 15.9. See [STL Features and Fixes in VS 2017 15.8](#) | [Visual C++ Team Blog](#).

11. Searchers & String Matching

`std::search` in C++14 offers a generic way to search for a pattern in a given range. The algorithm can be used not only for character containers but also for containers with custom types. This technique was, unfortunately, a bit limited as the performance was usually slow - it uses the naive matching algorithm, with the complexity of the size of the pattern times the size of the text. With C++17 we get new `std::search` overloads that expose new and powerful algorithms like Boyer Moore variations that have linear complexity in the average case.

In this chapter you'll learn:

- How we can beat a naive search algorithm with pattern preprocessing.
- How you can use `std::search` to efficiently search for a pattern in a range.
- How to use `std::search` for custom types.

Overview of String Matching Algorithms

String-matching consists of finding one or all of the occurrences of a string (“pattern”) in a text. The strings are built over a finite set of characters, called “alphabet”.

There are lots of algorithms that solve this problem; here’s a short list from [Wikipedia](https://en.wikipedia.org/wiki/String-searching_algorithm#Single-pattern_algorithms)¹:

Algorithm	Preprocessing	Matching	Space
Naive string-search	none	$O(nm)$	none
Rabin-Karp	$O(m)$	average $O(n + m)$, worst $O((n-m)m)$	$O(1)$
Knuth-Morris-Pratt	$O(m)$	$O(n)$	$O(m)$
Boyer-Moore	$O(m + k)$	best $O(n/m)$, worst $O(mn)$	$O(k)$
Boyer-Moore-Horspool	$O(m + k)$	best $O(n/m)$, worst $O(mn)$	$O(k)$

m - the length of the pattern n - the length of the text k - the size of the alphabet

The naive algorithm tries to match the pattern at each position of the text:

Pattern = Hello

Text = SuperHelloWorld

```

    SuperHelloWorld
1. Hello <- XX
2.  Hello <- XX
3.   Hello <- XX
4.    Hello <- XX
5.     Hello <- XX
6.      Hello <- OK!
```

In the example above we’re looking for “Hello” in “SuperHelloWorld”. As you can see the naive version tries each position until it finds the “Hello” at the 6th iteration.

The main difference between the naive way and the other algorithms is that the faster algorithms use additional knowledge about the input pattern. That way they can skip a lot of fruitless comparisons.

To gain that knowledge, they usually build some lookup tables for the pattern in the preprocessing phase. The size of lookup tables is often tied to the size of the pattern and the alphabet.

In the above case we can skip most of the iterations, as we can observe that when we try to match Hello in the first position, there’s a difference at the last letter o vs r. In fact, since r doesn’t occur in our pattern at all, we can actually move 5 steps further.

¹https://en.wikipedia.org/wiki/String-searching_algorithm#Single-pattern_algorithms

```
Pattern = Hello
Text = SuperAwesomeHelloWorld
```

```
    SuperAwesomeHelloWorld
1. Hello <- XX,
2.     Hello <- OK
```

We have a match with just 2 iterations! The rule that was used in that example comes from the Boyer Moore algorithm - it's called The Bad Character Rule.

In C++ string matching is implemented through `std::search`, which finds a range (the pattern) inside another range:

```
template< class ForwardIt1, class ForwardIt2 >
ForwardIt1 search( ForwardIt1 first, ForwardIt1 last,
                  ForwardIt2 s_first, ForwardIt2 s_last );
```

Before C++17 you had no control over the algorithm selection inside `std::search`. The complexity of `std::search` was specified as $O(mn)$ - so it was usually the naive approach. Now, in C++17, you have a few more options.

To be precise, there's also `string::find` that works exclusively with character sequences. There might be different implementations of this method, and you don't have control over the algorithm that is used inside. In the examples section, you'll see some performance experiments that also compare the new algorithms for `std::search` with `string::find`.

New Algorithms Available in C++17

C++17 updated `std::search` algorithm in two ways:

- you can now use execution policy to run the default version of the algorithm in a parallel way.
- you can provide a Searcher object that handles the search.

In C++17 we have three searchers:

- `default_searcher` - same as the version before C++17, usually meaning the naive approach. Operates on Forward Iterators.
- `boyer_moore_searcher` - uses Boyer Moore Algorithm - the full version, with two rules: bad character rule and good suffix rule. Operates on Random Access Iterators.
- `boyer_moore_horspool_searcher` - Simplified version of Boyer-Moore that uses only Bad Character Rule, but still has good average complexity. Operates on Random Access Iterators.

`std::search` with a searcher cannot be used along with execution policy.

Using Searchers

The `std::search` function uses the following overload for searchers:

```
template<class ForwardIterator, class Searcher>
ForwardIterator search(ForwardIterator first, ForwardIterator last,
                      const Searcher& searcher );
```

For example:

```
string testString = "Hello Super World";
string needle = "Super";
const auto it = search(begin(testString), end(testString),
                      boyer_moore_searcher(begin(needle), end(needle)));

if (it == cend(testString))
    cout << "The string " << needle << " not found\n";
```

Each searcher initialises its state through the constructor. The constructors need to store the pattern range and also perform the preprocessing phase. Then `std::search` calls their `operator()(iter TextFirst, iter TextLast)` method to perform the search in the text range.

Since a searcher is an object, you can pass it around in the application. That might be useful if you'd like to search for the same pattern inside various text objects. In that case, the preprocessing phase will be done only once.

Examples

Performance Experiments

This example builds a performance test to exercise several ways of finding a pattern in a larger text.

Here's how the test works:

- the application loads a text file (configurable via a command line argument), for example, a book sample (like a 500KB text file)
- the whole file content is stored in one string - that will be "text" where we'll be doing the lookups.
- a pattern is selected - N letters from the input text. That way we can be sure the pattern can be found. The position of the string can be located at the front, centre or at the end.
- the benchmark uses several algorithms and runs each search ITER times.

You can find the example in:

Chapter Searchers/searchers_benchmark.cpp

Example benchmarks:

the `std::string::find` version:

```
RunAndMeasure("string::find", [&]() {
    for (size_t i = 0; i < ITERS; ++i)
    {
        std::size_t found = testString.find(needle);
        if (found == std::string::npos)
            std::cout << "The string " << needle << " not found\n";
    }
});
```

The `boyer_moore_horspool` version:

```
RunAndMeasure("boyer_moore_horspool_searcher", [&]() {
    for (size_t i = 0; i < ITERS; ++i)
    {
        auto it = std::search(testString.begin(), testString.end(),
            std::boyer_moore_horspool_searcher(
                needle.begin(), needle.end()));
        if (it == testString.end())
            std::cout << "The string " << needle << " not found\n";
    }
});
```

`RunAndMeasure` is a function that takes a callable object to execute (for example a lambda). It measures the time of that execution and prints the results.

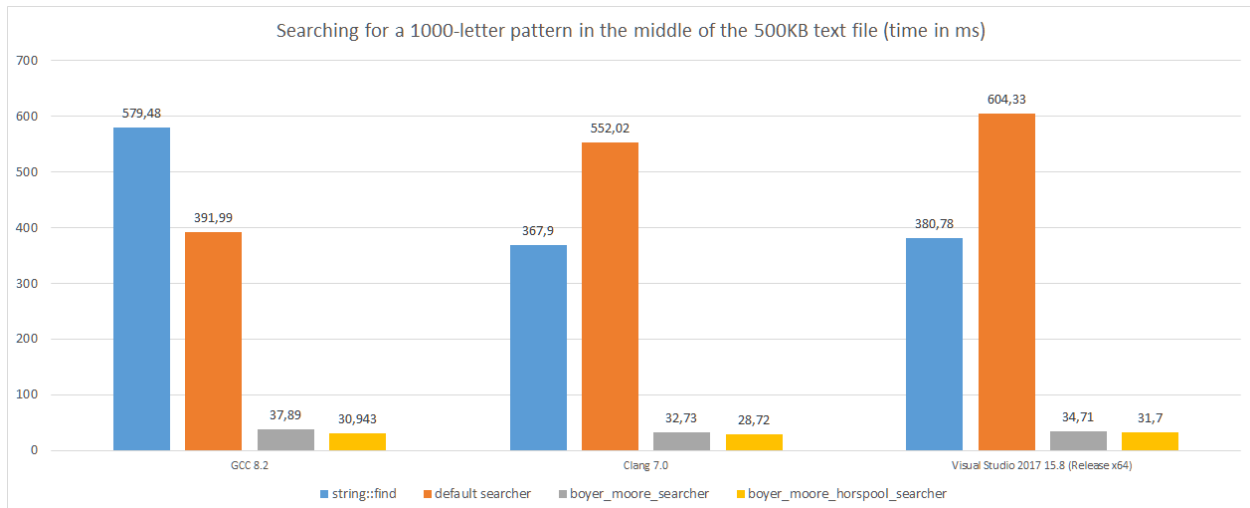
Since the input string is loaded from a file, the compiler cannot trick us and won't optimise code away.

Here are some of the results running the application on Win 10 64bit, i7 8700, 3.20 GHz base frequency, 6 cores/ 12 threads (the application runs on a single thread however).

The string size is 547412 bytes (comes from a 500KB text file), and we run the benchmark 1000 times.

Algorithm	GCC 8.2	Clang 7.0	Visual Studio (Release x64)
<code>string::find</code>	579.48	367.90	380.78
default searcher	391.99	552.02	604.33
<code>boyer_moore_searcher</code>	37.89 (init 3.98)	32.73 (init 3.02)	34.71 (init 3.52)
<code>boyer_moore_horspool_searcher</code>	30.943 (init 0)	28.72 (init 0.5)	31.70 (init 0.69)

When searching for 1000 letters from the centre of the input string, both of the new algorithms were faster than the default searcher and `string::find`. `boyer_moore` uses more time to perform the initialisation than `boyer_moore_horspool` (it creates two lookup tables, rather than one, so it will use more space and preprocessing). The results also show that `boyer_moore` usually takes longer time to preprocess the input pattern than `boyer_moore_horspool`. And also, the second algorithm is faster in our case. But all in all, the new algorithms perform even 10...15x faster than the default versions.




Searching for a 1000-letter pattern in the middle of the 500KB text file

Here are the results from another run, this time we use the same input string (from a 500KB text file), we perform 1000 iterations, but the pattern is only 48 letters. It's a sentence that's located at the end of the file (a single occurrence).

Algorithm	GCC 8.2	Clang 7.0	Visual Studio (Release x64)
<code>string::find</code>	164.58	39.63	40.28
default searcher	102.75	332.98	396.11
<code>boyer_moore_searcher</code>	115.69 (init 0.96)	95.56 (init 0.45)	101.73 (init 0.49)
<code>boyer_moore_horspool_searcher</code>	100.74 (init 0)	97.48 (init 0.21)	105.44 (init 0.23)

In this test Boyer-Moore algorithms in Visual Studio and Clang are 2.5x slower than `string::find`. However, on GCC `string::find` performed worse and `boyer_moore_horspool` is the fastest.



Searching for a 48-letter text at the end of the 500KB text file

You can run the experiments and see how your STL implementation performs. There are many ways to configure the benchmark so you can test various positions (beginning, centre, end) of the text, or check for some string pattern.

DNA Matching

To demonstrate the range of uses for `std::search`, let's have a look at a simple DNA matching demo. The example will match custom types rather than regular characters.

For instance, we'd like to search a DNA sequence to see whether GCTGC occurs in the sequence CTGATGTTAAGTCAACGCTGC.

The application uses a simple data structure for Nucleotides:

```
struct Nucleotide
{
    enum class Type : uint8_t {
        A = 0,
        C = 1,
        G = 3,
        T = 2
    };

    Type mType;

    friend bool operator==(Nucleotide a, Nucleotide b) noexcept
    {
        return a.mType == b.mType;
    }
}
```



```

    static char ToChar(Nucleotide t);
    static Nucleotide FromChar(char ch);
};

```

With the two converting static methods:

```

char Nucleotide::ToChar(Nucleotide t)
{
    switch (t.mType)
    {
        case Nucleotide::Type::A: return 'A';
        case Nucleotide::Type::C: return 'C';
        case Nucleotide::Type::G: return 'G';
        case Nucleotide::Type::T: return 'T';
    }
    return 0;
}

```

```

Nucleotide Nucleotide::FromChar(char ch)
{
    return Nucleotide { static_cast<Nucleotide::Type>((ch >> 1) & 0x03) };
}

```

And the two functions that work on a whole string:

```

std::vector<Nucleotide> FromString(const std::string& s)
{
    std::vector<Nucleotide> out;
    out.reserve(s.length());
    std::transform(std::cbegin(s), std::cend(s),
                   std::back_inserter(out), Nucleotide::FromChar);
    return out;
}

std::string ToString(const std::vector<Nucleotide>& vec)
{
    std::stringstream ss;
    std::ostream_iterator<char> out_it(ss);
    std::transform(std::cbegin(vec), std::cend(vec), out_it, Nucleotide::ToChar);
    return ss.str();
}

```

The demo uses `boyer_moore_horspool_searcher` which requires hashing support. So we have to define it as follows:

```
namespace std
{
    template<> struct hash<Nucleotide>
    {
        size_t operator()(Nucleotide n) const noexcept
        {
            return std::hash<Nucleotide::Type>{}(n.mType);
        }
    };
}
```

`std::hash` has support for enums, so we just have to “redirect” it from the whole class.

And then the test code:

```
const std::vector<Nucleotide> dna = FromString("CTGATGTTAAGTCAACGCTGC");
std::cout << ToString(dna) << '\n';
const std::vector<Nucleotide> s = FromString("GCTGC");
std::cout << ToString(s) << '\n';

std::boyer_moore_horspool_searcher searcher(std::cbegin(s), std::cend(s));
const auto it = std::search(std::cbegin(dna), std::cend(dna), searcher);

if (it == std::cend(dna))
    std::cout << "The pattern " << ToString(s) << " not found\n";
else
{
    std::cout << "DNA matched at position: "
                << std::distance(std::cbegin(dna), it) << '\n';
}
```

You can find the full source code in:

Chapter Searchers/dna_demo.cpp

As you can see the example builds a vector of custom types - Nucleotides. To satisfy the searcher a custom type needs to support `std::hash` interface and also define `operator==`.

The Nucleotide type wastes a bit of space - as we use the full byte just to store four options - C T G A. We could use only 2 bits, though the implementation would be more complicated. Another option is to represent triplets of Nucleotides - Codons. Each codon can be expressed in 6 bits, so that way we'd use the full byte more efficiently.

Summary

In this chapter you've learned about the searchers that can be passed into `std::search` algorithm. They allow you to use more advanced algorithms for string matching - Boyer-Moore and Boyer-Moore-Horspool that offers better complexity than a naive approach.

`std::search` with searchers is a general algorithm that works for most of the containers that expose random access iterators. If you work with strings and characters, then you might also compare it against `std::string::find`, which is usually specialized and optimised for character processing (implementation dependent!).



Extra Info

The change was proposed in: [N3905](https://wg21.link/n3905)²

Compiler support

Feature	GCC	Clang	MSVC
Searchers	7.1	3.9	VS 2017 15.3

²<https://wg21.link/n3905>

12. Filesystem

Since early versions, the Standard Library has included an option to work with files. Through streams - like `fstream` - you can open files, read data, write bytes and perform many other operations. However, what was missing was an ability to work with the filesystem as a whole. For example, in C++14 you had to use some third party libraries to iterate over directories, compose paths, delete directories or read file permissions. Now with C++17, we've taken a big step forward in the form of the `std::filesystem` component!

In this chapter you'll learn:

- How `std::filesystem` got into the Standard
- What the basic types and operations are
- How you can work with the paths
- How to handle errors in `std::filesystem`
- How to iterate over a directory
- How to create new directories and files

Filesystem Overview

While the Standard Library lacked some important features, you could always use Boost with its dozens of sub-libraries and do the work. The C++ Committee decided that the Boost libraries are so important that some of the systems were merged into the Standard. For example smart pointers (although improved with the move semantics in C++11), regular expressions, `std::optional`, `std::any` and much more.

A similar story happened with `std::filesystem`.

The filesystem library is modelled directly from Boost filesystem, which has been available since 2003 (with the version 1.30). In C++ implementation the committee also extended the component with non-POSIX systems. The library was first available as TS (Technical Specification) and later, after long time of improvements and feedback, merged into the C++17 Standard.

The library is located in the `<filesystem>` header, and it uses namespace `std::filesystem`.

Core Parts of The Library

The filesystem library is a rather significant part of the Standard Library. It defines many types with dozens of methods, and also gives us many free functions.

Below we can define the core elements of this module:

- The `std::filesystem::path` object allows you to manipulate paths that represent existing or not existing files and directories in the system.
- `std::filesystem::directory_entry` represents an existing path with additional status information like last write time, file size, or other attributes.
- Directory iterators allow you to iterate through a given directory. The library provides a recursive and non-recursive version.
- Many supporting functions like getting information about the path, file manipulation, permissions, creating directories, and many more.

In the next section, you'll see a demo of all the parts that compose `std::filesystem`.

Demo

Instead of exploring the library piece by piece at the start, on the next page you'll see a demo example: displaying basic information about all the files in a given directory (recursively). This should give you a high-level overview of what the library looks like.

```
1 // Chapter Filesystem/filesystem_list_files.cpp
2
3 #include <filesystem>
4 #include <iomanip>
5 #include <iostream>
6
7 namespace fs = std::filesystem;
8
9 void DisplayDirectoryTree(const fs::path& pathToScan, int level = 0) {
10     for (const auto& entry : fs::directory_iterator(pathToScan)) {
11         const auto filenameStr = entry.path().filename().string();
12         if (entry.is_directory()) {
13             std::cout << std::setw(level * 3) << "" << filenameStr << '\n';
14             DisplayDirectoryTree(entry, level + 1);
15         }
16         else if (entry.is_regular_file()) {
17             std::cout << std::setw(level * 3) << "" << filenameStr
18                 << ", size " << fs::file_size(entry) << " bytes\n";
19         }
20         else
21             std::cout << std::setw(level * 3) << "" << " [?]" << filenameStr << '\n';
22     }
23 }
24
25 int main(int argc, char* argv[]) {
26     try {
27         const fs::path pathToShow{ argc >= 2 ? argv[1] : fs::current_path() };
28
29         std::cout << "listing files in the directory: "
30             << fs::absolute(pathToShow).string() << '\n';
31
32         DisplayDirectoryTree(pathToShow);
33     }
34     catch (const fs::filesystem_error& err) {
35         std::cerr << "filesystem error! " << err.what() << '\n';
36     }
37     catch (const std::exception& ex) {
38         std::cerr << "general exception: " << ex.what() << '\n';
39     }
40 }
```

We can run this program on a temp path `D:\testlist\` and see the following output:

Running on Windows:

```
.\ListFiles.exe D:\testlist\  
listing files in the directory: D:\testlist\  
abc.txt, size 357 bytes  
def.txt, size 430 bytes  
ghi.txt, size 190 bytes  
dirTemp  
    jkl.txt, size 162 bytes  
    mno.txt, size 1728 bytes  
tempDir  
    abc.txt, size 174 bytes  
    def.txt, size 163 bytes  
tempInner  
    abc.txt, size 144 bytes  
    mno.txt, size 1728 bytes  
    xyz.txt, size 3168 bytes
```

The application lists files recursively, and with each indentation, you can see that we enter a new directory.

Running on a Linux (Ubuntu 18.04 on WSL):

```
fenbf@FEN-NODE:/mnt/f/wsl$ ./list_files.out testList/  
listing files in the directory: /mnt/f/wsl/testList/  
a.txt, size 965 bytes  
b.txt, size 1667 bytes  
c.txt, size 1394 bytes  
d.txt, size 1408 bytes  
directoryTemp  
    a.txt, size 1165 bytes  
    b.txt, size 1601 bytes  
tempDir  
    a.txt, size 1502 bytes  
    b.txt, size 1549 bytes  
    x.txt, size 1487 bytes
```

Let's now examine the core elements of this demo.

To work with the library, we have to include relevant headers. For the filesystem library it's:

```
#include <filesystem>
```

All the types, functions and names live in the `std::filesystem` namespace. For convenience it's useful to make a namespace alias:

```
namespace fs = std::filesystem;
```

And now we can refer to the names as `fs::path` rather than `std::filesystem::path`.

Let's start with the `main()` function where the logic of the application starts.

The program takes a single optional argument from the command line. If it's empty then we use the current system path:

```
const fs::path pathToShow{ argc >= 2 ? argv[1] : fs::current_path() };
```

`pathToShow` can be created from strings - from `argv[1]` if available, if not then `current_path()` is a helper free function that returns the current system path.

In the next two lines - line 28 and 29 - we display the starting path for the iteration. The `absolute()` function "expands" the input path and it converts it from a relative form into the absolute form if required.

The core part of the application is the `DisplayDirectoryTree()` function.

Inside, we use `directory_iterator` to examine the directory and find either the files or more directories.

```
for (const auto& entry : fs::directory_iterator(pathToShow))
```

Each loop iteration yields another `directory_entry` that we need to check. We can decide if we should call the function recursively (when `entry.is_directory()` is true) or just show some basic information if it's a regular file.

As you see we have access to many methods of path and directory entry. For example, we use `filename` to return only the filename part of the path so we can display "tree" structure. We also invoke `fs::file_size` to query the size of the file.

After a little demo let's now have a closer look at `filesystem::path`, `filesystem::directory_entry`, supporting, non-member functions and error handling.

The Path Object

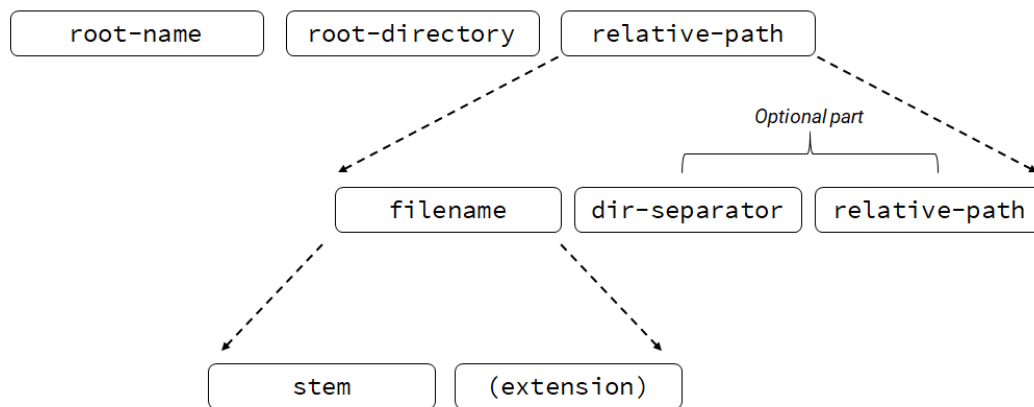
The core part of the library is the path object. It contains a pathname - a string that forms the name of the path. The object doesn't have to point to an existing file in the filesystem. The path might even be in an invalid form.

The path is composed of the following elements:

root-name root-directory relative-path:

- (optional) root-name: POSIX systems don't have a root name. On Windows, it's usually the name of a drive, like "C:"
- (optional) root-directory: distinguishes relative path from the absolute path
- relative-path:
 - filename
 - directory separator
 - relative-path

We can illustrate it with the following diagram:



The Path Structure

The class implements a lot of methods that extracts the parts of the path:

Method	Description
<code>path::root_name()</code>	returns the root-name of the path
<code>path::root_directory()</code>	returns the root directory of the path
<code>path::root_path()</code>	returns the root path of the path
<code>path::relative_path()</code>	returns path relative to the root path
<code>path::parent_path()</code>	returns the path of the parent path
<code>path::filename()</code>	returns the filename path component
<code>path::stem()</code>	returns the stem path component
<code>path::extension()</code>	returns the file extension path component

If a given element is not present, then the above functions return an empty path.

There are also methods that query elements of the path:

Query name	Description
<code>path::has_root_path()</code>	queries if a path has a root
<code>path::has_root_name()</code>	queries if a path has a root name
<code>path::has_root_directory()</code>	checks if a path has a root directory
<code>path::has_relative_path()</code>	checks if a path has a relative path component
<code>path::has_parent_path()</code>	checks if a path has a parent path
<code>path::has_filename()</code>	checks if a path has a filename
<code>path::has_stem()</code>	checks if a path has a stem component
<code>path::has_extension()</code>	checks if a path has an extension

We can use all of the above methods and compose an example that shows info about a given path:

```
// Chapter Filesystem/filesystem_path_info.cpp

const filesystem::path testPath{...};

if (testPath.has_root_name())
    cout << "root_name() = " << testPath.root_name() << '\n';
else
    cout << "no root-name\n";

if (testPath.has_root_directory())
    cout << "root_directory() = " << testPath.root_directory() << '\n';
else
    cout << "no root-directory\n";

if (testPath.has_root_path())
    cout << "root_path() = " << testPath.root_path() << '\n';
else
    cout << "no root-path\n";

if (testPath.has_relative_path())
    cout << "relative_path() = " << testPath.relative_path() << '\n';
else
    cout << "no relative-path\n";

if (testPath.has_parent_path())
    cout << "parent_path() = " << testPath.parent_path() << '\n';
else
    cout << "no parent-path\n";

if (testPath.has_filename())
    cout << "filename() = " << testPath.filename() << '\n';
else
```

```

    cout << "no filename\n";

    if (testPath.has_stem())
        cout << "stem() = " << testPath.stem() << '\n';
    else
        cout << "no stem\n";

    if (testPath.has_extension())
        cout << "extension() = " << testPath.extension() << '\n';
    else
        cout << "no extension\n";

```

Here's an output for a file path like "C:\Windows\system.ini":

```

root_name() = "C:"
root_directory() = "\\"
root_path() = "C:\\"
relative_path() = "Windows\system.ini"
parent_path() = "C:\Windows"
filename() = "system.ini"
stem() = "system"
extension() = ".ini"

```

Similarly on POSIX systems, for a path /usr/temp/abc.txt:

```

no root-name
root_directory() = "/"
root_path() = "/"
relative_path() = "usr/temp/abc.txt"
parent_path() = "/usr/temp"
filename() = "abc.txt"
stem() = "abc"
extension() = ".txt"

```

There's also a trick that lets you iterate over the parts of a path object. `std::filesystem::path` implements overloads for `begin()` and `end()` and that's why you can use it in a range based for loop:

```

int i = 0;
for (const auto& part : testPath)
    cout << "path part: " << i++ << " = " << part << '\n';

```

The output for C:\Windows\system.ini:

```

path part: 0 = C:
path part: 1 = \
path part: 2 = Windows
path part: 3 = system.ini

```

Path Operations

Below you can find a table with other important methods of the path class:

Operation	Description
path::append()	appends one path to the other, with a directory separator
path::concat()	concatenates the paths, without a directory separator
path::clear()	erases the elements and makes it empty
path::remove_filename()	removes the filename part from a path
path::replace_filename()	replaces a single filename component
path::replace_extension()	replaces the extension
path::swap()	swaps two paths
path::compare()	compares the lexical representations of the path and another path, returns an integer
path::empty()	checks if the path is empty

Comparison

The path class has several overloaded operators:

==, !=, <, >, <=, >=

And the path::compare() method, which returns an integer value.

All methods compare element by element, using the native format of the path.

```
fs::path p1 { "/usr/a/b/c" };  
fs::path p2 { "/usr/a/b/c" };  
assert(p1 == p2);  
assert(p1.compare(p2) == 0);  
  
p1 = "/usr/a/b/c";  
p2 = "/usr/a/b/c/d";  
assert(p1 < p2);  
assert(p1.compare(p2) < 0);
```

And on Windows we can also test the cases where we have a root element in a path:

```
p1 = "C:/test";  
p2 = "abc/xyz"; // no root path, so it's "less" than a path with a root  
assert(p1 > p2);  
assert(p1.compare(p2) > 0);
```

Or, also on Windows, a case where path formats are different:

```
fs::path p3 { "/usr/a/b/c" }; // on Windows it's converted to native format  
fs::path p4 { "\\usr/a\\b/c" };  
assert(p3 == p4);  
assert(p3.compare(p4) == 0);
```

You can play with the code in Chapter Filesystem/filesystem_compare.cpp.

Path Composition

We have two methods that let us compose a path:

- `path::append()` - adds a path with a directory separator.
- `path::concat()` - only adds the 'string' without any separator.

The functionality is also available with operators `/`, `/=` (append), `+` and `+=` (concat).

For example:

```
// append:
fs::path p1{"C:\\temp"};
p1 /= "user";
p1 /= "data";
cout << p1 << '\n';

// concat:
fs::path p2("C:\\temp\\");
p2 += "user";
p2 += "data";
cout << p2 << '\n';
```

The output:

```
C:\temp\user\data
C:\temp\userdata
```

However, appending a path has several rules that you have to be aware of.

For example, if the other path is absolute or the other path has a root-name, and the root-name is different from the current path root name. Then the append operation will replace this.

```
auto resW = fs::path{"foo"} / "D:\";    // Windows
auto resP = fs::path{"foo"} / "/bar";   // POSIX
// resW is "D:\" now
// resP is now "/bar"
```

In the above case `resW` and `resP` will contain the value from the second operand. As `D:\` and `/bar` contains root elements.

Stream Operators

The path class also implements `>>` and `<<` operators.

The operators use `std::quoted` to preserve the correct format. That's why the paths in the examples showed quotes.

On Windows, this will also cause the runtime to output `"\\\"` for paths in native format.

For example on POSIX:

```
fs::path p1 { "/usr/test/temp.xyz" };
std::cout << p1;
```

The code will print `"/usr/test/temp.xyz"`.

And on Windows:

```
fs::path p1{ "usr/test/temp.xyz" };
fs::path p2{ "usr\\test\\temp.xyz" };
std::cout << p1 << '\n' << p2;
```

The code will output:

```
"usr/test/temp.xyz"
"usr\\test\\temp.xyz"
```

Path Formats and Conversion

The filesystem library is modelled on top of POSIX (for example all Unix Based systems implements POSIX standard), but also works with other filesystems, for instance with Windows. Because of that, there are a few things to keep in mind when using paths in a portable way.

The first thing is the path format. We have two core modes:

- generic - generic format, the format as specified by the standard (based on the POSIX format)
- native - format used by the particular implementation

In POSIX systems native format is equal to generic. But On Windows it's different.

The main difference of the format is that Windows uses backslashes (\) rather than slashes (/). Another point is that Windows has a root directory - like C:, D: or other drive letters.

One more important aspect is the string type that is used to hold path elements. In POSIX it's `char` (and `std::string`), but on Windows it's `wchar_t` and `std::wstring`. The `path` type specifies `path::value_type` and `string_type` (defined as `std::basic_string<value_type>`) to expose those properties.

The `path` class has several methods that allow you to use the best matching format.

If you want to work with the native format you can use:

Operation	Description
<code>path::c_str()</code>	returns <code>value_type*</code>
<code>path::native()</code>	returns <code>string_type&</code>

And there are many methods that will convert the native format:

Operation	Description
<code>path::string()</code>	converts to <code>string</code>
<code>path::wstring()</code>	converts to <code>wstring</code>
<code>path::u8string()</code>	converts to <code>u8string</code>
<code>path::u16string()</code>	converts to <code>u16string</code>
<code>path::u32string()</code>	converts to <code>u32string</code>



Since Windows uses `wchar_t` as the underlying type for paths, then you need to be aware of “hidden” conversions to `char`.

The Directory Entry & Directory Iteration

While the `path` class represent files or paths that exist or not, we also have another object that is more concrete: it's `directory_entry` object. This object points to existing files and directories, and it's usually obtained with the aid of filesystem iterators.

What's more, implementations are encouraged to cache the additional file attributes. That way there can be fewer system calls.

Traversing a Path with Directory Iterators

You can traverse a path using two available iterators:

- `directory_iterator` - iterates in a single directory, input iterator.
- `recursive_directory_iterator` - iterates recursively, input iterator

In both approaches the order of the visited filenames is unspecified, each directory entry is visited only once.

If a file or a directory is deleted or added to the directory tree after the directory iterator has been created, it is unspecified whether the change would be observed through the iterator.

In both iterators the directories `.` and `..` are skipped.

You can iterate through a directory using the following pattern:

```
for (auto const & entry : fs::directory_iterator(pathToShow))
{
    ...
}
```

Or another way, with an algorithm, where you can also filter out paths:


```
std::filesystem::recursive_directory_iterator dirpos{ inPath };

std::copy_if(begin(dirpos), end(dirpos), std::back_inserter(paths),
             some_predicate);
```

`some_predicate` is a predicate that returns `bool` with `true` or `false` depending if a given `directory_entry` object matches our filter or not.

directory_entry Methods

Here's a list of `directory_entry` methods:

Operation	Description
<code>directory_entry::assign()</code>	replaces the path inside the entry and calls <code>refresh()</code> to update the cached attributes
<code>directory_entry::replace_filename()</code>	replaces the filename inside the entry and calls <code>refresh()</code> to update the cached attributes
<code>directory_entry::refresh()</code>	updates the cached attributes of a file
<code>directory_entry::exists()</code>	checks if a directory entry points to existing file system object
<code>directory_entry::is_block_file()</code>	returns true if the file entry is a block file
<code>directory_entry::is_character_file()</code>	returns true if the file entry is a character file
<code>directory_entry::is_directory()</code>	returns true if the file entry is a directory
<code>directory_entry::is_fifo()</code>	returns true if the file entry refers to a named pipe
<code>directory_entry::is_other()</code>	returns true if the file entry is refers to another file type
<code>directory_entry::is_regular_file()</code>	returns true if the file entry is a regular file
<code>directory_entry::is_socket()</code>	returns true if the file entry is a named IPC socket
<code>directory_entry::is_symlink()</code>	returns true if the file entry is a symbolic link
<code>directory_entry::file_size()</code>	returns the size of the file pointing by the directory entry
<code>directory_entry::hard_link_count()</code>	returns the number of hard links referring to the file
<code>directory_entry::last_write_time()</code>	gets or sets the last time write for a file
<code>directory_entry::status()</code>	returns status of the file designated by this directory entry
<code>directory_entry::symlink_status()</code>	returns the <code>symlink_status</code> of the file designated by this directory entry

Supporting Functions

So far we've covered three elements of the filesystem: the `path` class, `directory_entry` and `directory` iterators. The library also provides a set of non-member functions.

Query functions:

function	description
<code>filesystem::is_block_file()</code>	checks whether the given path refers to block device
<code>filesystem::is_character_file()</code>	checks whether the given path refers to a character device
<code>filesystem::is_directory()</code>	checks whether the given path refers to a directory
<code>filesystem::is_empty()</code>	checks whether the given path refers to an empty file or directory
<code>filesystem::is_fifo()</code>	checks whether the given path refers to a named pipe
<code>filesystem::is_other()</code>	checks whether the argument refers to another file
<code>filesystem::is_regular_file()</code>	checks whether the argument refers to a regular file
<code>filesystem::is_socket()</code>	checks whether the argument refers to a named IPC socket
<code>filesystem::is_symlink()</code>	checks whether the argument refers to a symbolic link
<code>filesystem::status_known()</code>	checks whether file status is known
<code>filesystem::exists()</code>	checks whether path refers to existing file system object
<code>filesystem::file_size()</code>	returns the size of a file
<code>filesystem::last_write_time()</code>	gets or sets the time of the last data modification

Path related:

function name	description
<code>filesystem::absolute()</code>	composes an absolute path
<code>filesystem::canonicalweakly_canonical()</code>	composes a canonical path
<code>filesystem::relativeproximate()</code>	composes a relative path
<code>filesystem::current_path()</code>	returns or sets the current working directory
<code>filesystem::equivalent()</code>	checks whether two paths refer to the same file system object

Directory and files management

function name	description
<code>filesystem::copy()</code>	copies files or directories
<code>filesystem::copy_file()</code>	copies file contents
<code>filesystem::copy_symlink()</code>	copies a symbolic link
<code>filesystem::create_directory()</code> ,	creates new directory
<code>filesystem::create_directories()</code>	
<code>filesystem::create_hard_link()</code>	creates a hard link
<code>filesystem::create_symlink()</code> ,	creates a symbolic link
<code>filesystem::create_directory_symlink()</code>	
<code>filesystem::hard_link_count()</code>	returns the number of hard links referring to the specific file
<code>filesystem::permissions()</code>	modifies file access permissions
<code>filesystem::read_symlink()</code>	obtains the target of a symbolic link
<code>filesystem::remove()</code> ,	removes a single file or whole directory recursively
<code>filesystem::remove_all()</code>	with all its content
<code>filesystem::rename()</code>	moves or renames a file or directory
<code>filesystem::resize_file()</code>	changes the size of a regular file by truncation or zero-fill
<code>filesystem::space()</code>	determines available free space on the file system
<code>filesystem::status()</code> ,	determines file attributes, determines file attributes,
<code>filesystem::symlink_status()</code>	checking the symlink target
<code>filesystem::temp_directory_path()</code>	returns a directory suitable for temporary files

Getting & Displaying the File Time

In C++17 there's one thing about `last_write_time()` values that's inconvenient.

We have one free function and a method in `directory_entry`. They both return `file_time_type` which is currently defined as:

```
using file_time_type = std::chrono::time_point</*trivial-clock*/>;
```

From the standard, 30.10.25 Header `<filesystem>` synopsis:

`trivial-clock` is an implementation-defined type that satisfies the `TrivialClock` requirements and that is capable of representing and measuring file time values. Implementations should ensure that the resolution and range of `filetime` type reflect the operating system dependent resolution and range of file time values.

In other words, it's implementation dependent.

For example in GCC/Clang STL file time is implemented on top of `chrono::system_clock`, but in MSVC it's a platform specific clock.

Here are some more details about the implementation decisions in Visual Studio: [std::filesystem::file_time_type does not allow easy conversion to time_t](https://developercommunity.visualstudio.com/content/problem/251213/stdfilesystemfile-time-type-does-not-allow-easy-co.html)¹

The situation might soon improve as in C++20 we'll get `std::chrono::file_clock` and also conversion routines between clocks. See [P0355](https://wg21.link/p0355)² (already added into C++20).

Let's have a look at some code.

```
auto filetype = fs::last_write_time(myPath);
const auto toNow = fs::file_time_type::clock::now() - filetype;
const auto elapsedSec = duration_cast<seconds>(toNow).count();
// skipped std::chrono prefix for duration_cast and seconds
```

The above code gives you a way to compute the number of seconds that have elapsed since the last update. This is however not as useful as showing a real date.

On POSIX (GCC and Clang Implementation) you can easily convert file time to `system_clock` and then obtain `std::time_t`:

¹<https://developercommunity.visualstudio.com/content/problem/251213/stdfilesystemfile-time-type-does-not-allow-easy-co.html>

²<https://wg21.link/p0355>

```

auto filetime = fs::last_write_time(myPath);
std::time_t convfiletime = std::chrono::system_clock::to_time_t(filetime);
std::cout << "Updated: " << std::ctime(&convfiletime) << '\n';

```

In MSVC the code won't compile. However there's a guarantee that `file_time_type` is usable with native OS functions that takes `FILETIME`. So we can write the following code to solve the issue:

```

auto filetime = fs::last_write_time(myPath);
FILETIME ft;
memcpy(&ft, &filetime, sizeof(FILETIME));
SYSTEMTIME stSystemTime;
if (FileTimeToSystemTime(&ft, &stSystemTime)) {
    // use stSystemTime.wYear, stSystemTime.wMonth, stSystemTime.wDay, ...
}

```

See Chapter `Filesystem/filesystem_list_files_info.cpp` to see examples.

File Permissions

In the table above you might noticed functions related to file permissions. We have two major functions:

- `std::filesystem::status()` and
- `std::filesystem::permissions()`

The first one returns `file_status` which contains information about the file type and also its permissions.

And you can use the second function to modify the file permissions. For example, to change a file to be read only.

File permissions - `std::filesystem::perms` - it's an enum class that represents the following values:

Name	Value (octal)	POSIX macro	Notes
none	0000		There are no permissions set for the file
owner_read	0400	S_IRUSR	Read permission, owner
owner_write	0200	S_IWUSR	Write permission, owner
owner_exec	0100	S_IXUSR	Execute/search permission, owner
owner_all	0700	S_IRWXU	Read, write, execute/search for owner
group_read	0040	S_IRGRP	Read permission, group
group_write	0020	S_IWGRP	Write permission, group
group_exec	0010	S_IXGRP	Execute/search permission, group
group_all	0070	S_IRWXG	Read, write, execute/search by group
others_read	0004	S_IROTH	Read permission, others
others_write	0002	S_IWOTH	Write permission, others
others_exec	0001	S_IXOTH	Execute/search permission, others
others_all	0007	S_IRWXO	Read, write, execute/search for others
all	0777		owner_all group_all others_all
set_uid	04000	S_ISUID	Set-user-ID on execution
set_gid	02000	S_ISGID	Set-group-ID on execution
sticky_bit	01000	S_ISVTX	Operating system dependent
mask	07777		all set_uid set_gid sticky_bit
unknown	0xFFFF		The permissions are not known

Here's a short code that demonstrates how to print file permissions:

```
// Chapter Filesystem/filesystem_permissions.cpp
```

```
std::ostream& operator<< (std::ostream& stream, fs::perms p)
{
    stream << "owner: "
        << ((p & fs::perms::owner_read) != fs::perms::none ? "r" : "-")
        << ((p & fs::perms::owner_write) != fs::perms::none ? "w" : "-")
        << ((p & fs::perms::owner_exec) != fs::perms::none ? "x" : "-");
    stream << " group: "
        << ((p & fs::perms::group_read) != fs::perms::none ? "r" : "-")
        << ((p & fs::perms::group_write) != fs::perms::none ? "w" : "-")
        << ((p & fs::perms::group_exec) != fs::perms::none ? "x" : "-");
    stream << " others: "
        << ((p & fs::perms::others_read) != fs::perms::none ? "r" : "-")
        << ((p & fs::perms::others_write) != fs::perms::none ? "w" : "-")
        << ((p & fs::perms::others_exec) != fs::perms::none ? "x" : "-");
    return stream;
}
```

You can use the above operator << implementation as follows:

```
std::cout << "perms: " << fs::status("myFile.txt").permissions() << '\n';
```

Setting Permissions

To change the permissions you can use the following code:

```
std::cout << "after creation: " << fs::status(sTempName).permissions() << '\n';
fs::permissions(sTempName, fs::perms::owner_read, fs::perm_options::remove);
std::cout << "after change: " << fs::status(sTempName).permissions() << '\n';
```

`std::filesystem::permissions` is a function that takes a path and then a flag and the “action” parameter.

`fs::perm_options` has three modes:

- `replace` - The permissions flag you pass will replace the existing state. It’s the default value for this parameter.
- `add` - The permission flag will be bitwise OR-ed with the existing state.
- `remove` - The permissions will be replaced by the bitwise AND of the negated argument and current permissions.
- `nofollow` - The permissions will be changed on the symlink itself, rather than on the file it resolves to

For example:

```
// remove "owner_read"
fs::permissions(myPath, fs::perms::owner_read, fs::perm_options::remove);

// add "owner_read"
fs::permissions(myPath, fs::perms::owner_read, fs::perm_options::add);

// replace and set "owner_all":
fs::permissions(myPath, fs::perms::owner_all); // replace is default param
```

Note for Windows

Windows is not a POSIX system, and it doesn’t map POSIX file permissions to its scheme. For `std::filesystem` it only supports two modes: read-only and all.

From [Microsoft Docs filesystem documentation](https://docs.microsoft.com/en-us/cpp/standard-library/filesystem-enumerations?view=vs-2017)³:

³<https://docs.microsoft.com/en-us/cpp/standard-library/filesystem-enumerations?view=vs-2017>

The supported values are essentially “readonly” and all. For a readonly file, none of the *_write bits are set. Otherwise the all bit (0777) is set.

Thus, unfortunately, you have limited options if you want to change file permissions on Windows.

Error Handling & File Races

So far the examples in this chapter have used exception handling to report errors. The filesystem API is also equipped with function and method overloads that outputs an error code. You can decide if you want exceptions or error codes.

For example we have two overloads for `file_size`:

```
uintmax_t file_size(const path& p);
uintmax_t file_size(const path& p, error_code& ec) noexcept;
```

the second one can be used in the following way:

```
const std::filesystem::path testPath("C:\\test.txt");
std::error_code ec{};
auto size = std::filesystem::file_size(testPath, ec);
if (ec == std::error_code{})
    std::cout << "size: " << size << '\n';
else
    std::cout << "error when accessing test file, size is: "
               << size << " message: " << ec.message() << '\n';
```

`file_size` takes additional output parameter - `error_code` and will set a value if something happens. If the operation is successful then `ec` will be value initialised.

File Races

It's important to point out the undefined behaviour that might happen when a file race occurs.

From *30.10.2.3 File system race behavior*:

The behavior is undefined if the calls to functions in this library introduce a file system race.

And from *30.10.9 file system race*:

The condition that occurs when multiple threads, processes, or computers interleave access and modification of the same object within a file system.

Examples

In this section, we'll analyse a few examples where `std::filesystem` is used. We'll go from a simple case - loading a file into a string, then explore creating directories and then filtering filenames using `std::regex`.

The demo sample is available in Chapter Filesystem/`filesystem_list_files.cpp` and its extended version (that shows file time and size) is located at Chapter Filesystem/`filesystem_list_files_info.cpp`

For more use cases you can also read the chapter - [How to Parallelise CSV Reader](#) - where `std::filesystem` is a crucial element for finding CSV files.

Loading a File into a String

The first example shows a compelling case where we leverage `std::filesystem`'s size related functions to build a buffer for the file contents.

Here's the code:

```
// Chapter Filesystem/filesystem_load_string.cpp

[[nodiscard]] std::string GetFileContents(const fs::path& filePath)
{
    std::ifstream inFile{ filePath, std::ios::in | std::ios::binary };
    if (!inFile)
        throw std::runtime_error("Cannot open " + filePath.string());

    const auto fsize = fs::file_size(filePath);
    if (fsize > std::numeric_limits<size_t>::max())
        throw std::runtime_error("file is too large to fit into size_t! "
                                   + filePath.string());

    std::string str(static_cast<size_t>(fsize), 0);

    inFile.read(str.data(), str.size());
    if (!inFile)
```



```

        throw std::runtime_error("Could not read the full contents from "
                                   + filePath.string());

    return str;
}

```

Before C++17 to get the file size you'd usually reposition the file pointer to the end and then read the position again. For example:

```

ifstream testFile("test.file", ios::binary);
const auto begin = myfile.tellg();
testFile.seekg (0, ios::end);
const auto end = testFile.tellg();
const auto fsize = (end-begin);

```

You could also open a file with `ios::ate` flag and then the file pointer will be positioned automatically at the end.

However, all of the above methods require to open a file but with `std::filesystem` the code is much shorter and we only have to read file properties.

What's more, the `std::filesystem` technique requires "lower" access rights as you only need parent directory read permission. There's no need to have "file read" permission.

If you use `std::filesystem::directory_entry` method, then it's possible that the file size comes from a cache.

Creating Directories

In the second example, we'll build *N* directories each with *M* files.

The core part of the `main()`:

```

// Chapter Filesystem/filesystem_build_temp.cpp

const fs::path startingPath{ argc >= 2 ? argv[1] : fs::current_path() };
const std::string strTempName{ argc >= 3 ? argv[2] : "temp" };
const int numDirectories{ argc >= 4 ? std::stoi(argv[3]) : 4 };
const int numFiles{ argc >= 5 ? std::stoi(argv[4]) : 4 };

if (numDirectories < 0 || numFiles < 0)
    throw std::runtime_error("negative input numbers...");

const fs::path tempPath = startingPath / strTempName;
CreateTempData(tempPath, numDirectories, numFiles);

```

And the `CreateTempData()` function:

```
// Chapter Filesystem/filesystem_build_temp.cpp
```

```
std::vector<fs::path> GeneratePathNames(const fs::path& tempPath,
                                       unsigned num)
{
    std::vector<fs::path> outPaths{ num, tempPath };
    for (auto& dirName : outPaths)
    {
        // use pointer value to generate unique name...
        const auto addr = reinterpret_cast<uintptr_t>(&dirName);
        dirName /= std::string("tt") + std::to_string(addr);
    }
    return outPaths;
}

void CreateTempFiles(const fs::path& dir, unsigned numFiles)
{
    auto files = GeneratePathNames(dir, numFiles);
    for (auto &oneFile : files)
    {
        std::ofstream entry(oneFile.replace_extension(".txt"));
        entry << "Hello World";
    }
}

void CreateTempData(const fs::path& tempPath, unsigned numDirectories,
                   unsigned numFiles)
{
    fs::create_directory(tempPath);
    auto dirPaths = GeneratePathNames(tempPath, numDirectories);

    for (auto& dir : dirPaths)
    {
        if (fs::create_directory(dir))
            CreateTempFiles(dir, numFiles);
    }
}
```

In `CreateTempData()` we first create the root of our folder structure. Then we generate path names in `GeneratePathNames()`. Each pathname is built from a pointer address. Such approach should give us good enough unique names. When we have a vector of unique paths, then we also generate another vector of paths for files. In this case, we use the `.txt` extension. While a directory is created using `fs::create_directory`, to create files we can use standard stream objects.

If we run the application with the following parameters: “. temp 2 4” it will create the following directory structure:

```
temp
  tt22325368
    tt22283456.txt, size 11 bytes
    tt22283484.txt, size 11 bytes
    tt22283512.txt, size 11 bytes
    tt22283540.txt, size 11 bytes
  tt22325396
    tt22283456.txt, size 11 bytes
    tt22283484.txt, size 11 bytes
    tt22283512.txt, size 11 bytes
    tt22283540.txt, size 11 bytes
```

Filtering Files Using Regex

The last example in this chapter will filter file names with the addition of `std::regex`, which has been available since C++11.

The core of the `main()`:

```
// Chapter Filesystem/filesystem_filter_files.cpp

const fs::path pathToShow{ argc >= 2 ? argv[1] : fs::current_path() };
const std::regex reg(argc >= 3 ? argv[2] : "");

auto files = CollectFiles(pathToShow);

std::sort(files.begin(), files.end());

for (auto& entry : files)
{
    const auto strFileName = entry.mPath.filename().string();
    if (std::regex_match(strFileName, reg))
        std::cout << strFileName << "\tsize: " << entry.mSize << '\n';
}
```

The application collects all the files from a given directory (recursively). Later a file entry is shown if it matches regex.

The `CollectFiles()` function uses recursive directory iterator to find all the files and it also builds basic information about each file. In `main()` we sort each file by size.

```
// Chapter Filesystem/filesystem_filter_files.cpp
struct FileEntry
{
    fs::path mPath;
    uintmax_t mSize{ 0 };

    static FileEntry Create(const fs::path& filePath) {
        return FileEntry{ filePath, fs::file_size(filePath) };
    }

    friend bool operator < (const FileEntry& a, const FileEntry& b) noexcept {
        return a.mSize < b.mSize;
    }
};

std::vector<FileEntry> CollectFiles(const fs::path& inPath)
{
    std::vector<fs::path> paths;
    if (fs::exists(inPath) && fs::is_directory(inPath))
    {
        std::filesystem::recursive_directory_iterator dirpos{ inPath };

        std::copy_if(begin(dirpos), end(dirpos), std::back_inserter(paths),
            []( const fs::directory_entry& entry) {
                return entry.is_regular_file();
            }
        );
    }
    std::vector<FileEntry> files(paths.size());
    std::transform(paths.cbegin(), paths.cend(), files.begin(), FileEntry::Create);
    return files;
}
```

In `CollectFiles` we use a recursive iterator and then `std::copy_if` to filter only regular files. Later, once the files are collected, we create the output vector of File Entries. `FileEntry::Create()` initialises objects and also fetches the size of a file.

For example, if we run the application with the following parameters “temp *.txt” we’ll be looking for all txt files in a directory.

```

.\FilterFiles.exe temp *.txt
tt22283456.txt size: 11
tt22283484.txt size: 11
tt22283512.txt size: 11
tt22283540.txt size: 11
tt22283456.txt size: 11
tt22283484.txt size: 11
tt22283512.txt size: 11
tt22283540.txt size: 11

```

Optimisation & Code Cleanup

The `CollectFiles()` function iterates over a directory and then outputs regular file's paths into a vector. Later the function creates `FileEntry` objects and returns that into a separate vector.

The thing is that we don't leverage all the possibilities of `filesystem::directory_entry` objects that we have during the scan. For example, the `filesystem::directory_entry::file_size` method will be much faster than `filesystem::file_size` because `directory_entry` usually keeps file attributes in cache.

Another element to optimise is a temporary vector of paths. We can skip it by using range-based for loop.

Here's the final code:

```

// Chapter Filesystem/filesystem_filter_files.cpp
std::vector<FileEntry> CollectFilesOpt(const fs::path& inPath)
{
    std::vector<FileEntry> files;
    if (fs::exists(inPath) && fs::is_directory(inPath))
    {
        for (const auto& entry : fs::recursive_directory_iterator{ inPath })
        {
            if (entry.is_regular_file())
                files.push_back({ entry, entry.file_size() });
        }
    }
    return files;
}

```

Chapter Summary

In this chapter, we dove into one of the biggest additions of C++17: `std::filesystem`. You saw the core elements of the library: the `path` class, `directory_entry` and iterators and lots of supporting

free functions.

Throughout the chapter, we also explored lots of examples: from simple cases like composing a path, getting file size, iterating through directories to even more complex: filtering with regex or creating temp directory structures.

You should be equipped with solid knowledge about `std::filesystem` and be prepared to explore the library on your own.



The full implementation of `std::filesystem` is described in the paper [P0218: Adopt the File System TS for C++17](https://wg21.link/p0218)⁴. There are also others updates like [P0317: Directory Entry Caching](https://wg21.link/p0317)⁵, [P0430 – File system library on non-POSIX-like operating systems](https://wg21.link/p0430)⁶, [P0492R2 - Resolution of C++17 National Body Comments](https://wg21.link/p0492)⁷, [P0392 -Adapting string_view by filesystem paths](https://wg21.link/p0392)⁸

You can find the final specification in [C++17 draft - N4687](https://ericniebler.com/2017/05/24/cplusplus17-draft-n4687/)⁹: the “filesystem” section, 30.10. Or under this online location [timsong-cpp/filesystems](https://github.com/timsong-cpp/filesystems)¹⁰.

⁴<https://wg21.link/p0218>

⁵<https://wg21.link/p0317>

⁶<https://wg21.link/p0430>

⁷<https://wg21.link/p0492>

⁸<https://wg21.link/p0392r0>

⁹<https://wg21.link/n4687>

¹⁰<https://timsong-cpp.github.io/cppwp/n4659/filesystems>

Compiler Support

GCC/libstdc++

The library was added in the version 8.0 - see [this commit - Implement C++17 Filesystem lib](#)¹¹

However, since GCC 5.3 you could play with the experimental version - the TS implementation.

In both cases (full support or experimental) you have to link with `-lstdc++fs` as the library is located outside the main standard library module.

For example to compile the sample code - Chapter Filesystem/filesystem_list_files.cpp - you should write the following command:

```
g++ -std=c++17 -O2 -Wall -Werror -pedantic filesystem_list_files.cpp -lstdc++fs
```

Clang/libc++

The support for `<filesystem>` was recently implemented in version 7.0, you can see [this commit - Implement filesystem](#)¹².

Since Clang 3.9 you could start playing with the experiential version, TS implementation.

Similarly, like GCC, you have to link to `libc++fs.a`.

Visual Studio

The full implementation of `<filesystem>` was added in Visual Studio 2017 15.7.

Before 15.7 you could play with `<experimental/filesystem>` in a much earlier version. The experimental implementation was available even in Visual Studio 2012 and later it was gradually improved with each release.

Compiler Support Summary

Feature	GCC	Clang	MSVC
Filesystem	8.0	7.0	VS 2017 15.7

¹¹<https://github.com/gcc-mirror/gcc/commit/3b90ed62fb848046ed7ddef07df7c806e7f3fadb>

¹²<https://github.com/llvm-mirror/libcxx/commit/a0866c5fb5020d15c69deda94d32a7f982b88dc9>

13. Parallel STL Algorithms

Concurrency and Parallelism are core aspects of any modern programming language. Before C++11 there was no standard support in the language for threading - you could use third-party libraries or System APIs. Modern C++ started to bring more and more necessary features: threads, atomics, locks, `std::async` and futures.

C++17 gives us a way to parallelise most of the standard library algorithms. With a powerful and yet straightforward abstraction layer, you can leverage more computing power out of a machine.

In this chapter you'll learn:

- What's on the way for C++ regarding parallelism and concurrency
- Why `std::thread` is not enough
- What the execution policies are
- How to run parallel algorithms
- Which algorithms were parallelised
- What the new algorithms are
- Expected performance of parallel execution
- Examples of parallel execution and benchmarks

Introduction

If we look at the computers that surround us, we can observe that most of them are multi-processors units. Even mobile phones have four or even eight cores. Not to mention graphics cards that are equipped with hundreds (or even thousands) of small computing cores.

The trend towards multicore machines was summarised perfectly in a famous article by Herb Sutter [The Free Lunch Is Over: A Fundamental Turn Toward Concurrency in Software](http://www.gotw.ca/publications/concurrency-ddj.htm)¹. The article appeared in 2006.

A glance around is all it takes to see that this trend is not slowing down.

While for a simple application there's probably no need to use the full computing capacity of your machine, there are applications which do require just that. Gaming, fast, responsive apps, graphics processing, video/music editing, data processing, financial, servers and many more types of systems. Spawning threads on a CPU and processing tasks concurrently is one way to achieve that.

With C++11/14 we've finally got threading into the standard library. You can now create instances of `std::thread` and not just depend on third-party libraries or a system API. What's more, there's also async processing with futures (`std::async`).

Multithreading is a significant aspect of modern C++. In the C++ Committee, there's a separate group - "SG1, Concurrency" that works on bringing more features like this to the standard.

What's on the way?

- Coroutines
- Atomic Smart pointers
- Transactional Memory
- Barriers
- Tasks blocks
- Parallelism
- Compute
- Executors
- Heterogeneous programming models support

As you can see, the plan is to expose as much of your machine's computing power as possible, directly from The Standard Library.

¹<http://www.gotw.ca/publications/concurrency-ddj.htm>

Not Only Threads

As mentioned earlier, using threads is not the only way of leveraging the power of your machine.

If your system has 8 cores in the CPU then you can use 8 threads and assuming you can split your work into separate chunks then you can hypothetically process your tasks several times faster than on a single thread.

But there's a chance to speed up things even more!

So where's the rest of the power coming from?

Vector Instructions from CPU and GPU computing.

The first element - vector instructions - allows you to compute several components of an array in a single instruction. It's also called SIMD - Single Instruction Multiple Data. Most of CPUs have 128-bit wide registers, and recent chips contain registers even 256 or 512 bits wide (AVX 256, AVX 512).

For example, using AVX-512 instructions, you can operate on 16 integer values (32-bit) at the same time!

The second element is the GPU. It might contain hundreds of smaller cores.

There are third-party APIs that allow you to access GPU/vectorisation: for example, we have CUDA, OpenCL, OpenGL, Intel TBB, OpenMP and many more. There's even a chance that your compiler will try to auto-vectorise some of the code. Still, we'd like to have such support directly from the Standard Library. That way the same code can be used on many platforms.

C++17 moves us a bit into that direction and allows us to use more computing power: it unlocks the auto vectorisation/auto parallelisation feature for algorithms in The Standard Library.

Overview

The new feature of C++17 looks surprisingly simple from a user point of view. You have a new parameter that can be passed to most of the standard algorithms: this new parameter is called **execution policy**.

```
std::algorithm_name(policy, /* normal args... */);
```

We'll go into the details later, but the general idea is that you call an algorithm and then you specify **how** it can be executed. Can it be parallel or just serial.

For example:

```
std::vector<int> v = genLargeVector();
// sort a vector using a parallel policy
std::sort(std::execution::par, v.begin(), v.end());
```

The above example will sort a vector in parallel - as specified by the first argument `std::execution::par`. The whole machinery is hidden from a user perspective. It's up to the STL implementation to choose the best approach to run tasks in parallel. Usually, they might leverage thread pools.

The hint - the execution policy parameter - is necessary because the compiler cannot deduce everything from the code. You, as the author of the code, only know if there are any side effects, possible race conditions, deadlocks, or if there's no sense in the running in parallel (such as if you have a small collection of items).



C++17's Parallelism comes from the Technical Specification that was published officially in 2015. The whole project of bringing parallel algorithms into C++ took more than five years - it started in 2012 and was merged into the standard in 2017. See the paper: [P0024 - The Parallelism TS Should be Standardised](https://ericniebler.com/2017/02/02/parallelism/)².

Execution Policies

The execution policy parameter will tell the algorithm how it should be executed.

We have the following options:

Policy Name	Description
<code>sequenced_policy</code>	It is an execution policy type used as a unique type to disambiguate parallel algorithm overloading and require that a parallel algorithm's execution not be parallelised. The corresponding global object is <code>'</code> .
<code>parallel_policy</code>	It is an execution policy type used as a unique type to disambiguate parallel algorithm overloading and indicate that a parallel algorithm's execution may be parallelised.
<code>parallel_unsequenced_policy</code>	It is an execution policy type used as a unique type to disambiguate parallel algorithm overloading and indicate that a parallel algorithm's execution may be parallelised and vectorised.

We have also three global objects corresponding to each execution policy type:

- `std::execution::seq`
- `std::execution::par`
- `std::execution::par_unseq`

Please note that execution policies are unique types, with their corresponding global objects. They are not enumerations, nor do they share the same base type.

²<https://wg21.link/P0024>

Execution policy declarations and global objects are located in the `<execution>` header.

Understanding Execution Policies

To understand the difference between execution policies let's try to build a model of how an algorithm might work.

Consider a simple vector of `float` values. In the below example, each element of the vector is multiplied by 2 and then the result is stored into an output container:

```
std::vector<float> vecX = {...}; // generate
std::vector<float> vecY(vecX.size());

std::transform(
    std::execution::seq,
    begin(vecX), end(vecX),          // input range
    begin(vecY),                    // output
    [](float x) { return x * 2.0f; }); // operation
```

Here's a pseudo code for sequential execution of the above algorithm:

```
operation
{
    load vecX[i] into RegisterX
    multiply RegisterX by 2.0f
    store RegisterX into vecY[i]
}
```

In the sequential execution, we'll access one element (from `vecX`), perform an operation and then store the result into the output (`vecY`). Execution for all elements happens on a single thread (on the calling thread).

With the `par` policy the whole operation for the *i*-th element will be executed on one thread. But there may be many threads that process different elements of the container. For example, if you have 8 free threads in the system, 8 elements of the container might be computed at the same time. The element access order is unspecified.

The Standard Library implementations might usually leverage some thread-pool to execute a parallel algorithm. The pool holds some worker threads (generally as many as system cores count), and then a scheduler will divide the input into chunks and assign them into the worker threads³. In theory, on a CPU, you could also create as many threads as elements in your container, but due to context switching that wouldn't give you good performance. On the other hand, implementations

³You might watch this great interview with Pedro Teixeira about thread pools in Windows - [Inside Windows 8: Pedro Teixeira - Thread pools](#) | [Channel 9](#). This is what MSVC implementation is using.

that use GPUs might provide hundreds of smaller “cores” so in that scenario the scheduler might work entirely differently.

The third execution policy `par_unseq` is an extension of the parallel execution policy. The operation for the i -th element will be performed on a separate thread, but also instructions of that operation might be interleaved and vectorised.

For example:

```
operation
{
    load vecX[i...i+3] into RegisterXYZW
    multiply RegisterXYZW by 2 // 4 elements at once!
    store RegisterXYZW into vecY[i...i+3]
}
```



The above pseudocode uses `RegisterXYZW` to represent a wide register that could store 4 elements of the container. For example, in SSE (Streaming SIMD Extensions) you have 128-bit registers that can handle 4 32-bit values, like 4 integers or 4 floating point numbers (or can store 8 16-bit values). However, such vectorisation might be extended to even larger registers like with AVX where you have 256 or even 512-bit registers. It's up to the implementation of the Standard Library to choose the best vectorisation scheme.

In this case, each instruction of the operation is “duplicated” and interleaved with others. That way the compiler can generate instructions that will handle several elements of the container at the same time.

In theory, if you have 8 free system threads, with 128-bit SIMD registers, and we process float values (32-bit values) - then, we can compute $8 \times 4 = 32$ values at once!



Why do you need the sequential policy?

Most of the time you'll be probably interested in using parallel policy or parallel unsequenced one. But for debugging it might be easier to use `std::execution::seq`. The parameter is also quite convenient as you might easily switch between the execution model using a template parameter. For some algorithms the sequential policy might also give better performance than the C++14 counterpart. Read more in the Benchmark section.

Limitations and Unsafe Instructions

The whole point of execution policies is to effortlessly parallelise standard algorithms in a declarative way. Nevertheless, there are some limitations you need to be aware of.

For example with `std::par` if you want to modify a shared resource you need to use some synchronisation mechanism to prevent data races and deadlocks⁴:

⁴When you only want to read a shared resource, then there's no need to synchronise.

```

std::vector<int> vec(1000);
std::iota(vec.begin(), vec.end(), 0);
std::vector<int> output;
std::mutex m;
std::for_each(std::execution::par, vec.begin(), vec.end(),
[&output, &m, &x](int& elem) {
    if (elem % 2 == 0) {
        std::lock_guard guard(m);
        output.push_back(elem);
    }
});

```

The above code filters out the input vector and then puts the elements in the output container.

If you forget about using a mutex (or another form of synchronisation), then `push_back` might cause **data races** - as multiple threads might try to add a new element to the vector at the same time.

The above example will also demonstrate weak performance, as using too many synchronisation points kills the parallel execution.



When using `par` execution policy try to access the shared resources as little as possible.

With `par_unseq` function invocations might be interleaved, so it's forbidden to use unsafe vectorised code. For example, using mutexes or memory allocation might lead to data races and deadlocks.

```

std::vector<int> vec = GenerateData();
std::mutex m;
int x = 0;
std::for_each(std::execution::par_unseq, vec.begin(), vec.end(),
[&m, &x](int& elem) {
    std::lock_guard guard(m);
    elem = x;
    x++; // increment a shared value
});

```

Since the instructions might be interleaved on one thread, you may end up with the following sequence of actions:

```
std::lock_guard guard(m) // for i-th element
std::lock_guard guard(m) // for i+1-th element
...
```

As you can see, two locks (in the same mutex) will happen on a single thread causing a deadlock!



Don't use synchronisation and memory allocation when executing with `par_unseq` policy.

Exceptions

When using execution policies, you need to be prepared for two kinds of situations.

- the scheduler, or the implementation fails to allocate resources for the invocation - then `std::bad_alloc` is thrown.
- an exception is thrown from the user code (a functor) - in that case, the exception is not re-thrown, `std::terminate()` is called.

See the example below.

```
// Chapter Parallel Algorithms/par_exceptions.cpp
try
{
    std::vector<int> v{ 1, 2, 3, 4, 5, 6, 7, 8, 9, 10 };

    std::for_each(std::execution::par, v.begin(), v.end(),
        [](int& i) {
            std::cout << i << '\n';

            if (i == 5)
                throw std::runtime_error("something wrong... !");
        });
}
catch (const std::bad_alloc& e)
{
    std::cout << "Error in execution: " << e.what() << '\n';
}
catch (const std::exception& e) // will not happen
{
    std::cout << e.what() << '\n';
}
```

```

}
catch (...)
{
    std::cout << "error!\n";
}

```

If you run the above code, the catch section will only handle `std::bad_alloc`. And if you exit a lambda because of some exception, then the `std::terminate` will be called. The exceptions are not re-thrown.



When you use parallel algorithms, for better error handling try to make your functors `noexcept`.

Algorithm Update

The execution policy parameter was added to most of the existing algorithms.

Here's the list:

<code>adjacent_difference</code>	<code>inplace_merge</code>	<code>replace_copy</code>
<code>adjacent_find</code>	<code>is_heap</code>	<code>replace_copy_if</code>
<code>all_of</code>	<code>is_heap_until</code>	<code>replace_if</code>
<code>any_of</code>	<code>is_partitioned</code>	<code>reverse</code>
<code>copy</code>	<code>is_sorted</code>	<code>reverse_copy</code>
<code>copy_if</code>	<code>is_sorted</code>	<code>rotate</code>
<code>copy_n</code>	<code>is_sorted_until</code>	<code>rotate_copy</code>
<code>count</code>	<code>lexicographical_compare</code>	<code>search</code>
<code>count_if</code>	<code>max_element</code>	<code>search_n</code>
<code>equal</code>	<code>merge</code>	<code>set_difference</code>
<code>exclusive_scan</code>	<code>min_element</code>	<code>set_intersection</code>
<code>fill</code>	<code>minmax_element</code>	<code>set_symmetric_difference</code>
<code>fill_n</code>	<code>mismatch</code>	<code>set_union</code>
<code>find</code>	<code>move</code>	<code>sort</code>
<code>find_end</code>	<code>none_of</code>	<code>stable_partition</code>
<code>find_first_of</code>	<code>nth_element</code>	<code>stable_sort</code>
<code>find_if</code>	<code>partial_sort</code>	<code>swap_ranges</code>
<code>find_if_not</code>	<code>partial_sort_copy</code>	<code>transform</code>
<code>for_each</code>	<code>partition</code>	<code>transform_exclusive_scan</code>
<code>for_each_n</code>	<code>partition_copy</code>	<code>transform_inclusive_scan</code>
<code>generate</code>	<code>remove</code>	<code>transform_reduce</code>
<code>generate_n</code>	<code>remove_copy</code>	<code>uninitialized_copy</code>
<code>includes</code>	<code>remove_copy_if</code>	<code>uninitialized_copy_n</code>
<code>inclusive_scan</code>	<code>remove_if</code>	<code>uninitialized_fill</code>

<code>inner_product</code>	<code>replace</code>	<code>uninitialized_fill_n</code>
	<code>unique</code>	<code>unique_copy</code>

New Algorithms

To fully support new parallel execution patterns The Standard Library was also equipped with a set of new algorithms:

Algorithm	Description
<code>for_each</code>	similar to <code>for_each</code> except returns <code>void</code>
<code>for_each_n</code>	applies a function object to the first <code>n</code> elements of a sequence
<code>reduce</code>	similar to <code>accumulate</code> , except out of order execution to allow parallelism
<code>transform_reduce</code>	transforms the input elements using a unary operation, then reduces the output out of order
<code>exclusive_scan</code>	parallel version of <code>partial_sum</code> , excludes the <code>i</code> -th input element from the <code>i</code> -th sum, out of order execution to allow parallelism
<code>inclusive_scan</code>	parallel version of <code>partial_sum</code> , includes the <code>i</code> -th input element in the <code>i</code> -th sum, out of order execution to allow parallelism
<code>transform_exclusive_scan</code>	applies a functor, then calculates exclusive scan
<code>transform_inclusive_scan</code>	applies a functor, then calculates inclusive scan

The new algorithms form three groups: `for_each`, `reduce` and then `scan`, plus their alternatives.

With `reduce` and `scan` you also get “fused” versions like `transform_reduce`. These compositions should give you much better performance than using two separate steps - because the cost of parallel execution setup is smaller and also you have one loop traversal less.

The new algorithms also provide overloads without the execution policy parameter so that you can use them in a standard serial version.

Below you’ll find a description of each group.

For Each Algorithm

In the serial version of `for_each`, the version that was available before C++17 you get a unary function as a return value from the algorithm.

Returning such an object is not possible in a parallel version, as the order of invocations is indeterminate.

Here’s a basic example:

```
// Chapter Parallel Algorithms/par_basic.cpp
std::vector<int> v(100);
std::iota(v.begin(), v.end(), 0);

std::for_each(std::execution::par, v.begin(), v.end(),
    [](int& i) { i += 10; });

std::for_each_n(std::execution::par, v.begin(), v.size()/2,
    [](int& i) { i += 10; });
```

The first `for_each` algorithm will update all of the elements of a vector, while the second execution will work only on the first half of the container.

Understanding Reduce Algorithms

Another core algorithm that is available with C++17 is `std::reduce`. This new algorithm provides a parallel version of `std::accumulate`. But it's important to understand the difference.

`std::accumulate` returns the sum of all the elements in a range (or a result of a binary operation that can be different than just a sum).

```
std::vector<int> v{1, 2, 3, 4, 5, 6, 7, 8, 9, 10};

auto sum = std::accumulate(v.begin(), v.end(), /*init*/0);
// sum is 55
```

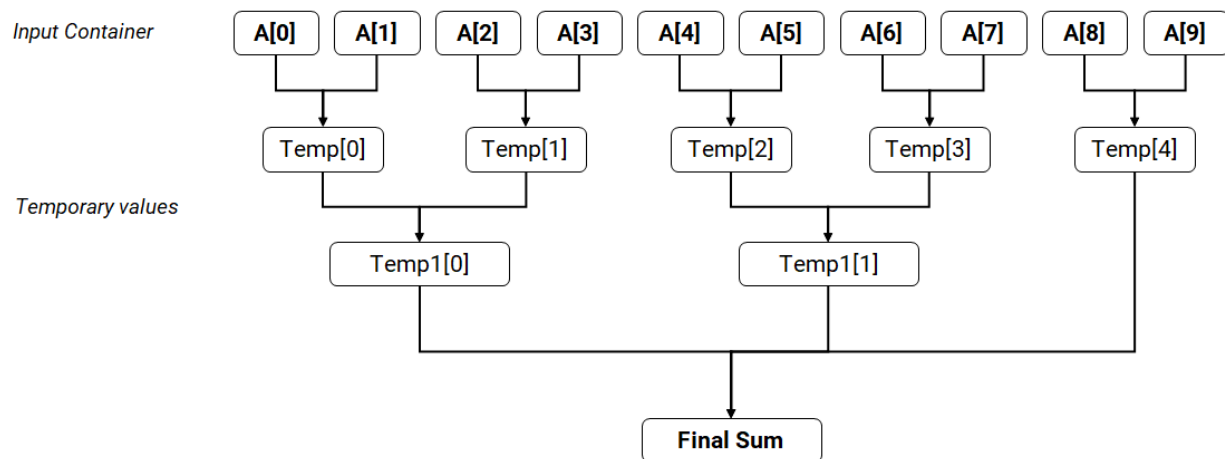
The algorithm is sequential and performs “left fold”, which means it will accumulate elements from the start to the end of a container.

The above example can be expanded into the following code:

```
sum = init +
    v[0] + v[1] + v[2] +
    v[3] + v[4] + v[5] +
    v[6] + v[7] + v[8] + v[9];
```

The parallel version - `std::reduce` - computes the final sum using a tree approach (sum sub-ranges, then merge the results, divide and conquer). This method can invoke the binary operation/sum in a nondeterministic order. Thus if `binary_op` is not associative or not commutative, the behaviour is also non-deterministic.

Here's a simplified picture that illustrates how a sum of 10 elements might work in a parallel way:



Parallel Sum Example

The above example with `accumulate` can be rewritten into `reduce`:

```
std::vector<int> v{1, 2, 3, 4, 5, 6, 7, 8, 9, 10};
```

```
auto sum = std::reduce(std::execution::par, v.begin(), v.end(), 0);
```

By default `std::plus<>{}` is used to compute the reduction steps.

A little explanation about associative and commutative operations:



A binary operation @ on a set S is **associative** if the following equation holds for all x, y, and z in S:

$$(x @ y) @ z = x @ (y @ z)$$

An operation is **commutative** if:

$$x @ y = y @ x$$

For example, we'll get the same results for `accumulate` and `reduce` for a vector of integers (when doing a sum), but we might get a slight difference for a vector of floats or doubles. That's because floating point sum operation is not associative.

An example:

```
// #include <limits> - for numeric_limits
std::cout.precision(std::numeric_limits<double>::max_digits10);
std::cout << (0.1+0.2)+0.3 << " != " << 0.1+(0.2+0.3) << '\n';
```

The output:

```
0.600000000000000009 != 0.59999999999999998
```

Another example might be the operation type: plus, for integer numbers, is associative and commutative, but minus is not associative nor commutative:

```
1+(2+3) == (1+2)+3 // sum is associative
1+8      == 8+1     // sum is commutative

1-(5-4) != (1-5)-4 // subtraction is not associative
1-7      != 7-1     // subtraction is not commutative
```

transform_reduce - Fused Algorithm

To get even more flexibility and performance, the reduce algorithm also has a version where you can apply a transform operation before doing the reduction.

```
std::vector<int> v{1, 2, 3, 4, 5, 6, 7, 8, 9, 10};
auto sumTransformed = std::transform_reduce(std::execution::par,
    v.begin(),
    v.end(),
    0,
    std::plus<int>{},
    [](const int& i) { return i * 2; }
);

// sum is 110
```

The above code will first execute the unary functor - the lambda that doubles the input value; then the results will be reduced into a single sum.

The fused version will be faster than using two algorithms: `std::reduce` firstly and then `std::transform_reduce` - because the implementation will need to perform the parallel execution setup only once.

Scan Algorithms

The third group of new algorithms is scan. They implement a version of partial sum, but out of order.

The exclusive scan does not include the i-th element in the output i-th sum, while inclusive scan does.

For example for array = {1, 2, 3, 4, 5, 6, 7, 8, 9, 10}

We'll get the following values for partials sums:

Name										
Index	0	1	2	3	4	5	6	7	8	9
Values	1	2	3	4	5	6	7	8	9	10
Exclusive partial sums	0	1	3	6	10	15	21	28	36	45
Inclusive partial sums	1	3	6	10	15	21	28	36	45	55

Similarly to `std::reduce` the order of the operations is unsequenced, so to get deterministic results the `binary_op` must be associative.

`scan` has also two fused algorithms: `transform_exclusive_scan` and `transform_inclusive_scan`. Both of the algorithms will perform a unary operation on the input container and then they will compute the prefix sums on the output.

Prefix sums have an essential role in many applications, for example for stream compaction, computing summed area tables or radix sort. Here's a link to an article that describes the algorithms in detail: [GPU Gems 3 - Chapter 39. Parallel Prefix Sum \(Scan\) with CUDA⁵](#).

Performance of Parallel Algorithms

Parallel algorithms are a powerful abstraction layer. Although they're relatively easy to use, assessing their performance is less straightforward.

The first point to note is that a parallel algorithm will generally do more work than the sequential version. That's because the algorithm needs to set up and arrange the threading subsystem to run the tasks.

For example, if you invoke `std::transform` on a vector of 100k elements, then the STL implementation needs to divide your vector into chunks and then schedule each chunk to be executed appropriately. If necessary, the implementation might even copy the elements before the execution. If you have a system with 10 free threads, the 100k element vector might be divided into chunks of 10k, and then each chunk is transformed by one thread. Due to the setup cost and other limitations of parallel code, the whole execution time won't be 10x faster than the serial version.

The **second** factor that plays an important role is synchronisation. If your operations need to synchronise on some shared objects, then the parallel execution performance decreases. A parallel algorithm performs best when you execute separate tasks.

The **third** element that has a big effect on the execution is memory throughput. If we look at a common desktop CPU, we can see that all cores share the same memory bus. That's why if your instructions wait for the data to be fetched from memory, then the performance improvement over the sequential won't be visible, as all cores will synchronise on the memory access. Algorithms like `std::copy`, `std::reverse` might be even slower than their serial versions - at least on common PC hardware⁶. It's best when your parallel tasks use CPU cycles for computation rather than waiting for memory access.

⁵https://developer.nvidia.com/gpugems/GPUGems3/gpugems3_ch39.html

⁶See section "Current Limitations of the MSVC Implementation of Parallel Algorithms" in [Using C++17 Parallel Algorithms for Better Performance | Visual C++ Team Blog](#) where parallel `std::reverse` appeared to be 1.6 times slower.

The **fourth** important thing is that the algorithms are very implementation dependent. They might use different techniques to achieve parallelism. Not to mention the device that an algorithm might be executed on - on CPU or GPU.

Right now there's only one implementation available in a popular compiler - in Visual Studio 2017. It's based on thread pools from Windows and only supports execution on the CPU and skips the vectorisation execution policy ⁷.

It will be exciting to see and compare official support from Clang/GCC in their Standard Library implementations. While Visual Studio can run only on Windows, the Clang/GCC versions would have to support more systems.

Another thing is the GPU support. With hundreds of smaller computing cores, the algorithms might perform faster than their CPU version. It's important to remember that before executing something on the GPU, you have to usually copy the data to memory visible to the GPU (unless it's a shared memory like in integrated Graphics Cards). And sometimes the cost of the data transfer might reduce the total performance.



If you decide to use parallel algorithms, it's best to measure the performance against the sequential version. Sometimes, especially for a smaller number of elements, the performance might be even slower.

Examples

So far you've seen basic code samples with parallel algorithms. This section will introduce a few more examples with more complex scenarios.

- You'll see a few benchmarks and see the performance gains over the sequential version.
- We'll discuss an example of how to process several containers in the same parallel algorithm.
- There will also be a sample of how to implement a parallel version of counting elements.



Please note that all of the measurements for benchmarks were done only using Visual Studio - as it's the only available implementation of parallel algorithms in a popular compiler. The benchmarks will be updated when implementations from GCC and Clang will be ready.

Benchmark

Finally, we can see the performance of the basic algorithms.

Let's have a look at an example where we execute separate tasks on each element - using `std::transform`. In that example, the speed up vs the sequential version should be more visible.

⁷The VS implementation will usually process elements in blocks, so there's a chance that auto-vectorisation might still optimise the code and allow vector instruction for better performance.

```
// Chapter Parallel Algorithms/par_benchmark.cpp
#include <algorithm>
#include <execution>
#include <iostream>
#include <numeric>
#include "simpleperf.h"

int main(int argc, const char* argv[])
{
    const size_t vecSize = argc > 1 ? atoi(argv[1]) : 6000000;
    std::cout << vecSize << '\n';
    std::vector<double> vec(vecSize, 0.5);
    std::vector out(vec);

    RunAndMeasure("std::transform seq", [&vec, &out] {
        std::transform(std::execution::seq, vec.begin(), vec.end(), out.begin(),
            [](double v) {
                return std::sin(v)*std::cos(v);
            }
        );
        return out.size();
    });

    RunAndMeasure("std::transform par", [&vec, &out] {
        std::transform(std::execution::par, vec.begin(), vec.end(), out.begin(),
            [](double v) {
                return std::sin(v)*std::cos(v);
            }
        );
        return out.size();
    });

    return 0;
}
```

The code calculates $\sin \cdot \cos$ ⁸ and stores the result in the output vector. Those trigonometry functions were used to use some more CPU arithmetic instructions, rather than just fetching an element from memory.

The application was run on two machines:

- i7 4720H - means Win 10 64bit, i7 4720H, 2.60 GHz base frequency, 4 Cores/8 Threads, MSVC 2017 15.8, Release mode, x86.

⁸You can also use more optimal computation of $\sin \cdot \cos$ as $\sin(2x) = 2 \sin(x) \cos(x)$.


- i7 8700 - means Win 10 64bit, i7 8700, 3,2 GHz base frequency, 6 Cores/12 Threads, MSVC 2017 15.8, Release Mode, x86.

RunAndMeasure is a helper function that runs a function and then prints the timings. The result is used later so that the compiler doesn't optimise the variable away:

```
template <typename TFunc> void RunAndMeasure(const char* title, TFunc func)
{
    const auto start = std::chrono::steady_clock::now();
    ret = func();
    const auto end = std::chrono::steady_clock::now();
    std::cout << title
        << ": " << std::chrono::duration<double, std::milli>(end - start).count()
        << " ms " << ret << '\n';
}
```

Here are the results (time in milliseconds):

algorithm	vector size	i7 4720H	i7 8700
std::transform, seq	1000000	10.9347	7.51991
std::transform, par	1000000	2.67921	1.30245
std::transform, seq	2000000	21.8466	15.028
std::transform, par	2000000	5.29644	2.34634
std::transform, seq	3000000	32.7403	22.1449
std::transform, par	3000000	7.79366	3.42295
std::transform, seq	4000000	44.2565	30.1643
std::transform, par	4000000	11.7558	4.40974



Benchmark of std::transform

The example above might be the perfect case for a parallelisation: we have an operation that requires a decent amount of instructions (trigonometry functions), and then all the tasks are separate. In this case, on a machine with 6 cores and 12 threads, the performance is almost 7X faster! On a computer with 4 cores and 8 threads, the performance is 4.2X faster.

It's also worth to notice that when the transformation instructions are simple like `return v*2.0` then the performance speed-up might be not seen. This is because all the code will be just waiting on the global memory and it might perform the same as the sequential version.

Below there's a benchmark of computing the sum of all elements in a vector:

```
// Chapter Parallel Algorithms/par_benchmark.cpp
#include <algorithm>
#include <execution>
#include <iostream>
#include <numeric>
#include "simpleperf.h"

int main(int argc, const char* argv[])
{
    const size_t vecSize = argc > 1 ? atoi(argv[1]) : 60000000;
    std::cout << vecSize << '\n';
    std::vector<double> vec(vecSize, 0.5);

    RunAndMeasure("std::accumulate", [&vec] {
```

```

        return std::accumulate(vec.begin(), vec.end(), 0.0);
    });

    RunAndMeasure("std::reduce, seq", [&vec] {
        return std::reduce(std::execution::seq,
            vec.begin(), vec.end(), 0.0);
    });


    RunAndMeasure("std::reduce, par", [&vec] {
        return std::reduce(std::execution::par,
            vec.begin(), vec.end(), 0.0);
    });

    return 0;
}

```

Here are the results:

algorithm	size	i7 4720H	i7 8700
std::accumulate	10000000	10.5814	9.62405
std::reduce seq	10000000	6.9556	4.58746
std::reduce par	10000000	4.88708	3.67831
std::accumulate	15000000	17.8769	14.9163
std::reduce seq	15000000	11.5103	5.42508
std::reduce par	15000000	9.99877	4.5679
std::accumulate	20000000	21.8888	19.6507
std::reduce seq	20000000	16.2142	6.80581
std::reduce par	20000000	10.8826	4.79214



Benchmark std::accumulate vs std::reduce

During this execution, the par version was 2x..4x faster than the standard std::accumulate!



Another reason to use sequential policy?

Sequential version of std::reduce was also faster than std::accumulate, and this might be because with reduce the order of operations is not determined, while accumulate is a left fold. So the compiler can optimise it better.

Processing Several Containers At the Same Time

When using parallel algorithms, you might sometimes want to access other containers.

For example, you might want to execute `for_each` on two containers.

The main technique is to get the index of the element currently being processed. Then you can use that index to access other containers (assuming the containers are of the same size).

We can do it in a few ways:

- by using a separate container of indices
- by using zip iterators/wrappers

Let's have a look at the techniques:

Separate Container of Indices

```
// Chapter Parallel Algorithms/par_iterating_multiple.cpp
void Process(int a, int b) { }

std::vector<int> v(100);
std::vector<int> w(100);
std::iota(v.begin(), v.end(), 0);
std::iota(w.begin(), w.end(), 0);

std::vector<size_t> indexes(v.size());
std::iota(indexes.begin(), indexes.end(), 0);

std::for_each(std::execution::par, indexes.begin(), indexes.end(),
    [&v, &w](size_t& id) {
        Process(v[id], w[id]);
    }
);
```

As you see in the above code a separate vector of indices had to be generated. Since the order of execution is not specified, then we cannot iterate using a simple `i` variable.

Zip Iterators

A more elegant technique is to use zip iterators:

```
// Chapter Parallel Algorithms/par_iterating_multiple.cpp
void Process(int a, int b) { }

std::vector<int> v(100);
std::vector<int> w(100);
std::iota(v.begin(), v.end(), 0);
std::iota(w.begin(), w.end(), 0);

vec_zipper<int, int> zipped{ v, w };
std::for_each(std::execution::seq, zipped.begin(), zipped.end(),
    [](std::pair<int&, int&>& twoElements) {
        Process(twoElements.first, twoElements.second);
    }
);
```

The example uses a custom implementation of a `ZipIterator` that works only with `std::vector`. You can make the code more general or use boost zip iterators if you have access to the library.

The full implementation of `ZipIterator` is in the example “`par_iterating_multiple.cpp`”.

Erroneous Technique

It's also important to mention one aspect:

As the standard says: in [\[algorithms.parallel.exec\]⁹](#)

Unless otherwise stated, implementations may make arbitrary copies of elements (with type `T`) from sequences where `is_trivially_copy_constructible_v<T>` and `is_trivially_destructible_v<T>` are true. [Note: This implies that user-supplied function objects should not rely on object identity of arguments for such input sequences. Users for whom the object identity of the arguments to these function objects is important should consider using a wrapping iterator that returns a non-copied implementation object such as `reference_wrapper<T>` ([\[refwrap\]^a](#)) or some equivalent solution. —end note]

^a<https://timsong-cpp.github.io/cppwp/n4659/refwrap>

Thus you cannot write:

```
vector<int> vec;
vector<int> other;
vector<int> external;
int* beg = vec.data();
std::transform(std::execution::par,
               vec.begin(), vec.end(), other.begin(),
               [&beg, &external](const int& elem) {
                   // use pointer arithmetic
                   auto index = &elem - beg;
                   return elem * externalVec[index];
               });
```

The code above uses pointer arithmetic to find the current index of the element. Then we can use this index to access other containers.

The technique, however, uses the assumption that `elem` is the exact element from the container and not its copy! Since the implementations might copy elements, the addresses might be completely unrelated! This faulty technique also assumes that the container is storing the items in a contiguous chunk of memory.

Only `for_each` and `for_each_n` has a guarantee that the elements are not being copied during the execution:

⁹<https://timsong-cpp.github.io/cppwp/n4659/algorithms.parallel#exec-2>

Implementations do not have the freedom granted under `[algorithms.parallel.exec]`^a to make arbitrary copies of elements from the input sequence.

^a<https://timsong-cpp.github.io/cppwp/n4659/algorithms.parallel.exec>

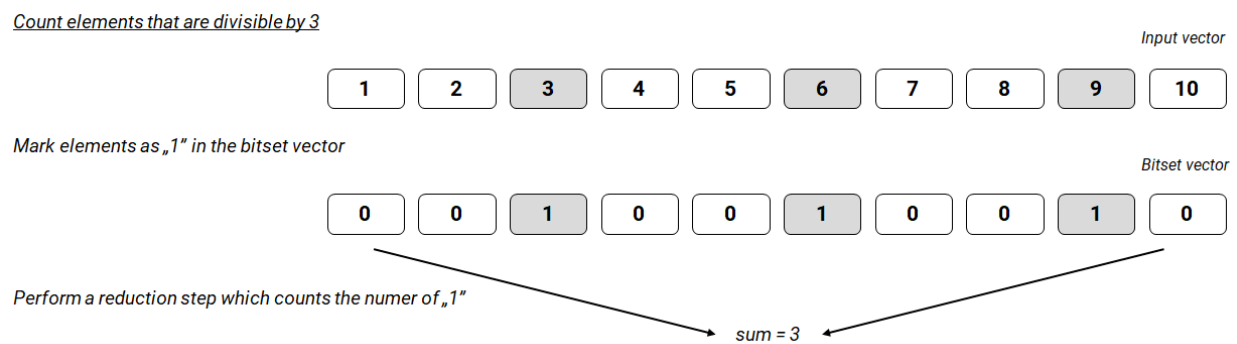
Counting Elements

To gain some practice let's build an algorithm that counts the number of elements in a container. Our algorithm will be a version of another standard algorithm `count_if`.

The main idea is to use `transform_reduce` - a new “fused” algorithm. It first applies some unary function over an element and then performs a reduce operation.

To get the count of elements that satisfy some predicate we can firstly filter each element (transform). We return 1 if the element passes the filter and 0 otherwise. Then, in the reduction step we count how many elements returned 1.

Here's a diagram that illustrates the algorithm for a simple case:



Parallel Count IF Example

- The first step is to perform the transform step in `transform_reduce` algorithm. We return 1 for matching elements and 0 otherwise.
- Then the reduce step is used to compute the sum of all 1. We have three values that satisfy the condition, so the output is 3.

And here's the code:

```
// Chapter Parallel Algorithms/par_count_if.cpp
template <typename Policy, typename Iter, typename Func>
std::size_t CountIf(Policy policy, Iter first, Iter last, Func predicate)
{
    return std::transform_reduce(policy,
        first,
        last,
        std::size_t(0),
        std::plus<std::size_t>{},
        [&predicate](const Iter::value_type& v) {
            return predicate(v) ? 1 : 0;
        }
    );
}
```

We can run it on the following test containers:

```
// Chapter Parallel Algorithms/par_count_if.cpp
std::vector<int> v(100);
std::iota(v.begin(), v.end(), 0);
auto NumEven = CountIf(std::execution::par, v.begin(), v.end(),
    [](int i) { return i % 2 == 0; }
);
std::cout << NumEven << '\n';
```

To get number of spaces in a string:

```
// Chapter Parallel Algorithms/par_count_if.cpp
std::string_view sv = "Hello Programming World";
auto NumSpaces = CountIf(std::execution::seq, sv.begin(), sv.end(),
    [](char ch) { return ch == ' '; }
);
std::cout << NumSpaces << '\n';
```

Or even on a map:

```
// Chapter Parallel Algorithms/par_count_if.cpp
std::map<std::string, int> CityAndPopulation{
    {"Cracow", 765000},
    {"Warsaw", 1745000},
    {"London", 10313307},
    {"New York", 18593220},
    {"San Diego", 3107034}
};
auto NumCitiesLargerThanMillion = CountIf(std::execution::seq,
    CityAndPopulation.begin(), CityAndPopulation.end(),
    [](const std::pair<const std::string, int>& p) {
        return p.second > 1000000;
    }
);
std::cout << CitiesLargerThanMillion << '\n';
```

The example uses simple test data and to have good performance over the sequential version the size of data would have to be significantly increased. For example, the cities and their population could be loaded from a database.

More Examples

Here's a list of a few other ideas where parallel algorithms could be beneficial:

- statistics - calculating various maths properties for a set of data
- processing CSV records line by line in parallel
- parsing files in parallel - one file per thread, or chunks of a file per thread
- computing summed area tables
- parallel matrix operations
- parallel dot product

You can find a few more examples in the following articles:

- [The Amazing Performance of C++17 Parallel Algorithms, is it Possible?](https://www.bfilipek.com/2018/11/parallel-alg-perf.html)¹⁰
- [How to Boost Performance with Intel Parallel STL and C++17 Parallel Algorithms](https://www.bfilipek.com/2018/11/pstl.html)¹¹
- [Examples of Parallel Algorithms From C++17](https://www.bfilipek.com/2018/06/parstl-tests.html)¹²
- [Parallel STL And Filesystem: Files Word Count Example](https://www.bfilipek.com/2018/07/files-word-count.html)¹³

¹⁰<https://www.bfilipek.com/2018/11/parallel-alg-perf.html>

¹¹<https://www.bfilipek.com/2018/11/pstl.html>

¹²<https://www.bfilipek.com/2018/06/parstl-tests.html>

¹³<https://www.bfilipek.com/2018/07/files-word-count.html>

Chapter Summary

After reading the chapter, you should be equipped with the core knowledge about parallel algorithms. We discussed the execution policies, how they might be executed on hardware, what are the new algorithms.

At the moment parallel algorithms show good potential. With only one extra parameter you can easily parallelise your code. Previously that would require to use some third-party library or write a custom version of some thread pooling system.

For sure we need to wait for more available implementations. Currently only Visual Studio lets you use parallel algorithms and it offers only parallel policy (not vectorised). Executing code on GPU looks especially interesting.

It's also worth quoting the TS specification [P0024](https://wg21.link/P0024)¹⁴:

The parallel algorithms and execution policies of the Parallelism TS are only a starting point. Already we anticipate opportunities for extending the Parallelism TS's functionality to increase programmer flexibility and expressivity. A fully-realised executors feature will yield new, flexible ways of creating execution, including the execution of parallel algorithms.

Things to remember:

- Parallel STL gives you set of 69 algorithms that have overloads for the execution policy parameter.
- Execution policy describes how the algorithm might be executed.
- There are three execution policies in C++17 (`<execution>` header)
 - `std::execution::seq` - sequential
 - `std::execution::par` - parallel
 - `std::execution::par_unseq` - parallel and vectorised
- In parallel execution policy functors that are passed to algorithms cannot cause deadlocks and data races
- In parallel unsequenced policy functors cannot call vectorised unsafe instructions like memory allocations or any synchronisation mechanisms
- To handle new execution patterns there are also new algorithms: like `std::reduce`, `exclusive_scan` - They work out of order so the operations must be associative to generate deterministic results
- There are “fused” algorithms: `transform_reduce`, `transform_exclusive_scan`, `transform_inclusive_scan` that combine two algorithms together.

¹⁴<https://wg21.link/P0024>

- Assuming there are no synchronisation points in the parallel execution, the parallel algorithms should be faster than the sequential version. Still, they perform more work - especially the setup and divisions into tasks.
- Implementations might usually use some thread pools to implement a parallel algorithm on CPU.

Compiler Support

As of today (October 2018) only one compiler/STL implementation supports parallel algorithms: it's Visual Studio 2017 15.7. MSVC supports most of the new algorithms, and more will be implemented in the next updates.

Visual Studio also implements `par` and `par_unseq` in the same way, so you shouldn't expect any difference between code runs.

Other major compilers Clang/GCC (and their STL implementations) are on the way.

Feature	GCC	Clang	MSVC
Parallel Algorithms	in progress	in progress	VS 2017 15.7 ¹⁵

There are also several other implementations out there:

- Codeplay - [SYCL Parallel STL](#)¹⁶
- STE||AR Group - [HPX](#)¹⁷
- Intel - [Parallel STL](#)¹⁸ - based on OpenMP 4.0 and Intel® TBB.
- [Parallel STL](#)¹⁹ - early Microsoft implementation for the Technical Specification.
- [n3554](#) - proposal implementation (started by Nvidia)²⁰
- [Thibaut Lutz Implementation @Github](#)²¹ - early implementation

¹⁵See [Announcing: MSVC Conforms to the C++ Standard](#) | Visual C++ Team Blog

¹⁶<http://github.com/KhronosGroup/SyclParallelSTL>

¹⁷<http://stellar-group.github.io/hpx/docs/html/hpx/manual/parallel.html>

¹⁸<https://software.intel.com/en-us/get-started-with-pstl>

¹⁹<https://parallelstl.codeplex.com/>

²⁰<https://github.com/n3554/n3554>

²¹<http://github.com/t-lutz/ParallelSTL>

14. Other Changes In The Library

C++17 is a significant update for the language, and it brings a lot of features in the Standard Library. So far this book has covered the most important aspects, but there are many more things to describe!

This part of the book summaries briefly other changes in the Standard Library:

- What's `std::byte`?
- What are the new functionalities of maps and sets
- New algorithms: sampling
- Special Mathematical Functions
- Shared Pointers and Arrays
- Non-member `size()`, `data()` and `empty()`
- `constexpr` additions to the Standard Library
- How to lock multiple mutexes with `scoped_lock`?
- What's a polymorphic allocator? How can it help with memory management?

std::byte

std::byte is a small type that gives you a view of bytes and bits rather than numeric/char values (like unsigned char).

In the Standard it's defined as enum:

```
enum class byte : unsigned char {} ; // in <cstdint>
```

You can initialise byte from unsigned chars with the syntax byte b{unsigned char}, and you can convert from byte into a numeric type by std::to_integer().

Let's see a basic example:

```
// Chapter STL Other/byte.cpp
#include <cstdint>

int main() {
    constexpr std::byte b{1};
    // std::byte c{3535353}; // error: narrowing conversion from int

    constexpr std::byte c{255};

    // shifts:
    constexpr auto b1 = b << 7;
    static_assert(std::to_integer<int>(b) == 0x01);
    static_assert(std::to_integer<int>(b1) == 0x80);

    // various bit operators, like &, |, &, etc
    constexpr auto c1 = b1 ^ c;
    static_assert(std::to_integer<int>(c) == 0xff);
    static_assert(std::to_integer<int>(c1) == 0x7f);

    constexpr auto c2 = ~c1;
    static_assert(std::to_integer<int>(c2) == 0x80);
    static_assert(c2 == b1);
}
```

The primary motivation behind std::byte is type-safety in the context of memory/byte access.

See the reference paper [P0298R3](https://wg21.link/P0298R3)¹

¹<https://wg21.link/P0298R3>

Improvements for Maps and Sets

In the Standard there are two notable features for maps and sets:

- Splicing Maps and Sets - [P0083](https://wg21.link/P0083)²
- New emplacement routines - [N4279](https://wg21.link/N4279)³

Splicing

You can now move nodes from one tree-based container (maps/sets) into other ones, without additional memory overhead/allocation.

For example:

```
// Chapter STL Other/set_extract_insert.cpp
#include <set>
#include <string>
#include <iostream>

struct User {
    std::string name;

    User(std::string s) : name(std::move(s)) {
        std::cout << "User::User(" << name << ")\n";
    }
    ~User() {
        std::cout << "User::~~User(" << name << ")\n";
    }
    User(const User& u) : name(u.name) {
        std::cout << "User::User(copy, " << name << ")\n";
    }

    friend bool operator<(const User& u1, const User& u2) {
        return u1.name < u2.name;
    }
};

int main() {
    std::set<User> setNames;
    setNames.emplace("John");
```

²<https://wg21.link/P0083>

³<https://wg21.link/N4279>

```

    setNames.emplace("Alex");
    setNames.emplace("Bartek");
    std::set<User> outSet;

    std::cout << "move John...\n";
    // move John to the outSet
    auto handle = setNames.extract(User("John"));
    outSet.insert(std::move(handle));

    for (auto& elem : setNames)
        std::cout << elem.name << '\n';

    std::cout << "cleanup...\n";
}

```

Output:

```

User::User(John)
User::User(Alex)
User::User(Bartek)
move John...
User::User(John)
User::~~User(John)
Alex
Bartek
cleanup...
User::~~User(John)
User::~~User(Bartek)
User::~~User(Alex)

```

In the above example, one element “John” is extracted from `setNames` into `outSet`. The `extract` method moves the found node out of the set and physically detaches it from the container. Later the extracted node can be inserted into a container of the same type.

Before C++17, if you wanted to move an object from one map (or set) and put it into another map (or set), you had to remove it from the first container and then copy/move into another. With the new functionality, you can manipulate tree nodes (by using implementation-defined type `node_type`) and leave objects unaffected. Such a technique allows you to handle non-movable elements and of course, is more efficient.

Emplace Enhancements for Maps and Unordered Maps

With C++17 you get two new methods for maps and unordered maps:

- `try_emplace()` - if the object already exists then it does nothing, otherwise it behaves like `emplace()`.
 - `emplace()` might move from the input parameter when the key is in the map, that's why it's best to use `find()` before such emplacement.
- `insert_or_assign()` - gives more information than `operator[]` - as it returns if the element was newly inserted or updated and also works with types that have no default constructor.

`try_emplace` Method

Here's an example:

```
// Chapter STL Other/try_emplace_map.cpp
#include <iostream>
#include <string>
#include <map>

int main() {
    std::map<std::string, int> m;

    m["hello"] = 1;
    m["world"] = 2;

    // C++11 way:
    if (m.find("great") == std::end(m))
        m["great"] = 3;

    // the lookup is performed twice if "great" is not in the map

    // C++17 way:
    m.try_emplace("super", 4);
    m.try_emplace("hello", 5); // won't emplace, as it's
                               // already in the map

    for (const auto& [key, value] : m)
        std::cout << key << " -> " << value << '\n';
}
```

The behaviour of `try_emplace` is important in a situation when you move elements into the map:


```
// Chapter STL Other/try_emplace_map_move.cpp
#include <iostream>
#include <string>
#include <map>

int main() {
    std::map<std::string, std::string> m;
    m["Hello"] = "World";

    std::string s = "C++";
    m.emplace(std::make_pair("Hello", std::move(s)));

    // what happens with the string 's'?
    std::cout << s << '\n';
    std::cout << m["Hello"] << '\n';

    s = "C++";
    m.try_emplace("Hello", std::move(s));
    std::cout << s << '\n';
    std::cout << m["Hello"] << '\n';
}
```

The code tries to replace ["Hello", "World"] into ["Hello", "C++"].

If you run the example the string `s` after `emplace` is empty and the value “World” is not changed into “C++”!

`try_emplace` does nothing in the case where the key is already in the container, so the `s` string is unchanged.

insert_or_assign Method

The second function `insert_or_assign`, inserts a new object in the map or assigns the new value. But as opposed to `operator[]` it also works with non-default constructible types.

For example:

```
// Chapter STL Other/insert_or_assign.cpp
#include <string>
#include <map>
#include <iostream>

struct User {
    std::string name;

    User(std::string s) : name(std::move(s)) {
        std::cout << "User::User(" << name << ")\n";
    }
    ~User() {
        std::cout << "User::~~User(" << name << ")\n";
    }
    User(const User& u) : name(u.name) {
        std::cout << "User::User(copy, " << name << ")\n";
    }

    friend bool operator<(const User& u1, const User& u2) {
        return u1.name < u2.name;
    }
};

int main() {
    std::map<std::string, User> mapNicks;
    //mapNicks["John"] = User("John Doe"); // error: no default ctor for User()

    auto [iter, inserted] = mapNicks.insert_or_assign("John", User("John Doe"));
    if (inserted)
        std::cout << iter->first << " entry was inserted\n";
    else
        std::cout << iter->first << " entry was updated\n";
}
```

Output:

```
User::User(John Doe)
User::User(copy, John Doe)
User::~~User(John Doe)
John entry was inserted
User::~~User(John Doe)
```

In the example above we cannot use `operator[]` to insert a new value into the container, as it doesn't support types with non-default constructors. We can do it with the new function.

`insert_or_assign` returns a pair of `<iterator, bool>`. If the boolean value is true, it means the element was inserted into the container. Otherwise, it was reassigned.

Return Type of Emplace Methods

Since C++11 most of the standard containers got `.emplace*` methods. With those methods, you can create a new object in place, without additional object copies.

However, most of `.emplace*` methods didn't return any value - it was `void`. Since C++17 this is changed, and they now return the reference type of the inserted object.

For example:

```
// since C++11 and until C++17 for std::vector
template< class... Args >
void emplace_back( Args&&... args );

// since C++17 for std::vector
template< class... Args >
reference emplace_back( Args&&... args );
```

This modification should shorten the code that adds something to the container and then invokes some operation on that newly added object.

For example:

```
// Chapter STL Other/emplace_return.cpp
#include <vector>
#include <string>

int main() {
    std::vector<std::string> stringVector;

    // in C++11/14:
    stringVector.emplace_back("Hello");
    // emplace doesn't return anything, so back() needed
    stringVector.back().append(" World");

    // in C++17:
    stringVector.emplace_back("Hello").append(" World");
}
```

The referenced paper with the proposal: [P0084R2](http://wg21.link/p0084r2)⁴

⁴<http://wg21.link/p0084r2>

Sampling Algorithms

New algorithm - `std::sample` - that selects `n` elements from the sequence:

```
// Chapter STL Other/sample.cpp
#include <iostream>
#include <random>
#include <iterator>
#include <algorithm>

int main() {
    std::vector<int> v { 1, 2, 3, 4, 5, 6, 7, 8, 9, 10 };
    std::vector<int> out;
    std::sample(v.begin(),                // range start
               v.end(),                  // range end
               std::back_inserter(out),  // where to put it
               3,                        // number of elements to sample
               std::mt19937{std::random_device{}}());

    std::cout << "Sampled values: ";
    for (const auto &i : out)
        std::cout << i << ", ";
}
```

Possible output:

Sampled values: 1, 4, 9,

The new sampling algorithms come from the adoption of Library Fundamentals V1 TS Components:
Sampling [P0220R1](https://wg21.link/p0220)⁵

⁵<https://wg21.link/p0220>

New Mathematical Functions

With C++17 we get lots of new mathematical functions like `std::gcd` (Greatest Common Divisor), `std::lcm` (Least Common Multiple), `std::clamp` and other special ones.

For example `std::gcd` and `std::lcm`, declared in `<numeric>` header:

```
// Chapter STL Other/numeric_gcd_lcm.cpp
#include <iostream>
#include <numeric> // for gcd, lcm

int main() {
    std::cout << std::gcd(24, 60) << '\n';
    std::cout << std::lcm(15, 50) << '\n';
}
```

Output:

```
12
150
```

Those two were introduced in [P0295R0](http://wg21.link/p0295r0)⁶

Another useful function is `std::clamp(v, min, max)`, declared in `<algorithm>` introduced in: [P0025](http://wg21.link/p0025)⁷

```
// Chapter STL Other/clamp.cpp
#include <iostream>
#include <algorithm> // clamp

int main() {
    std::cout << std::clamp(300, 0, 255) << '\n';
    std::cout << std::clamp(-10, 0, 255) << '\n';
}
```

Output:

```
255
0
```

⁶<http://wg21.link/p0295r0>

⁷<http://wg21.link/p0025>

What's more, there are newly available special functions, defined in the `<cmath>` header.

- `assoc_laguerre`
- `assoc_legendre`
- `beta`
- `comp_ellint_1`
- `comp_ellint_2`
- `comp_ellint_3`
- `conf_hyperg`
- `cyl_bessel_i`
- `cyl_bessel_j`
- `cyl_bessel_k`
- `cyl_neumann`
- `ellint_1`
- `ellint_2`
- `ellint_e`
- `erfc`
- `erf`
- `expint`
- `hermite`
- `hyperg`
- `laguerre`
- `legendre`
- `riemann_zeta`
- `sph_bessel`
- `sph_legendre`
- `sph_neumann`
- `tgamma`

The above special functions were introduced in [N1542 ver 3](http://www.open-std.org/jtc1/sc22/wg21/docs/papers/2003/n1542.pdf)⁸.

⁸<http://www.open-std.org/jtc1/sc22/wg21/docs/papers/2003/n1542.pdf>

Shared Pointers and Arrays

Before C++17, only `unique_ptr` was able to handle arrays out of the box (without the need to define a custom deleter). Now it's also possible with `shared_ptr`.

```
std::shared_ptr<int[]> ptr(new int[10]);
```

Please note that `std::make_shared` doesn't support arrays in C++17. But this will be fixed in C++20 (see [P0674](https://wg21.link/P0674)⁹ which is already merged into C++20)

Another important remark is that raw arrays should be avoided. It's usually better to use standard containers. However, sometimes, you don't have the luxury to use vectors or lists - for example in an embedded environment, or when you work with third-party API. In that situation, you might end up with a raw pointer to an array. With C++17 you'll be able to wrap those pointers into smart pointers (`std::unique_ptr` or `std::shared_ptr`) and be sure the memory is deleted correctly.

The referencing paper: [P0414R2](https://wg21.link/P0414R2)¹⁰

⁹<https://wg21.link/p0674>

¹⁰<https://wg21.link/P0414R2>

Non-member `size()`, `data()` and `empty()`

Following the approach from C++11 regarding non-member `std::begin()` and `std::end()` C++17 brings three new functions.

```
// Chapter STL Other/non_member_functions.cpp
#include <iostream>
#include <vector>

template <class Container>
void PrintBasicInfo(const Container& cont) {
    std::cout << typeid(cont).name() << '\n';
    std::cout << std::size(cont) << '\n';
    std::cout << std::empty(cont) << '\n';

    if (!std::empty(cont))
        std::cout << *std::data(cont) << '\n';
}

int main() {
    std::vector<int> iv { 1, 2, 3, 4, 5 };
    PrintBasicInfo(iv);
    float arr[4] = { 1.1f, 2.2f, 3.3f, 4.4f };
    PrintBasicInfo(arr);
}
```

Output (from GCC 8.2):

```
St6vectorIiSaIiEE
5
0
1
A4_f
4
0
1.1
```

The new functions are located in `<iterator>`, and the referencing paper is [N4280](https://wg21.link/n4280) ¹¹.

¹¹<https://wg21.link/n4280>

constexpr Additions to the Standard Library

With this enhancement you can work with iterators, accessing ranges in constexpr contexts.

The main referencing paper is [P0031 - Proposal to Add Constexpr Modifiers to reverse_iterator, move_iterator, array and Range Access](#)¹².

Have a look at a basic example, that shows a basic implementation of accumulate algorithm as constexpr (the current version of `std::accumulate` is not constexpr):

```
// Chapter STL Other/constexpr_additions.cpp
#include <array>

template<typename Range, typename Func, typename T>
constexpr T SimpleAccumulate(const Range& range, Func func, T init) {
    for (auto &&obj : range) { // begin/end are constexpr
        init += func(obj);
    }
    return init;
}

constexpr int square(int i) { return i*i; }

int main() {
    constexpr std::array arr{ 1, 2, 3 }; // clas deduction...

    // with constexpr lambda
    static_assert(SimpleAccumulate(arr, [](int i) constexpr {
        return i * i;
    }, 0) == 14);

    // with constexpr function
    static_assert(SimpleAccumulate(arr, &square, 0) == 14);

    return arr[0];
}
```

C++14 compilers would not compile the above example, but it's now possible with C++17 support.

There are several features used in this code:

- `SimpleAccumulate` is a constexpr template function and uses range access (hidden in range based for loop) to iterate over the input range.

¹²<https://wg21.link/p0031>

- `arr` is deduced as `std::array<3, int>` - class template deduction works here.
- the code uses `constexpr` lambda.
- `static_assert` without any message.

Each C++ Standard allows more and more code to be `constexpr`. In C++17 we can start using basic containers in constant expressions, and in the future, in C++20 even the standard algorithms will enable this option (if feasible).

std::scoped_lock

With C++11 and C++14 we got the threading library and many support functionalities.

For example, with `std::lock_guard` you can take ownership of a mutex and lock it in RAII style:

```
std::mutex m;  
  
std::lock_guard<std::mutex> lock_one(m);  
// unlocked when lock_one goes out of scope...
```

The above code works, however, only for a single mutex. If you wanted to lock several mutexes you had to use a different pattern, for example:

```
std::mutex first_mutex;  
std::mutex second_mutex;  
  
// ...  
  
std::lock(first_mutex, second_mutex);  
std::lock_guard<std::mutex> lock_one(first_mutex, std::adopt_lock);  
std::lock_guard<std::mutex> lock_two(second_mutex, std::adopt_lock);  
// ..
```

With C++17 things get a bit easier as with `std::scoped_lock` you can lock a variadic number of mutexes at the same time.

```
std::scoped_lock lck(first_mutex, second_mutex);
```

Due to compatibility `std::lock_guard` couldn't be extended with a variadic number of input mutexes and that's why a new type - `scoped_lock` - was needed.

You can read more information in [P0156](https://wg21.link/p0156)¹³

¹³<https://wg21.link/p0156>

std::iterator Is Deprecated

The Standard Library API requires that each iterator type has to expose five typedefs:

- `iterator_category` - the type of the iterator
- `value_type` - type stored in the iterator
- `difference_type` - the type that is the result of subtracting two iterators
- `pointer` - pointer type of the stored type
- `reference` - the reference type of the stored type

`iterator_category` must be one of `input_iterator_tag`, `forward_iterator_tag`, `bidirectional_iterator_tag` or `random_access_iterator_tag`.

Before C++17, if you wanted to define a custom iterator, you could use `std::iterator` as a base class. In C++14 it's defined as:

```
template<
    class Category,
    class T,
    class Distance = std::ptrdiff_t,
    class Pointer = T*,
    class Reference = T&
> struct iterator;
```

This helper class made defining the typedefs in a short way:

```
class ColumnIterator
: public std::iterator<std::random_access_iterator_tag, Column>
{
    // ...
};
```

In C++17 you must not derive from `std::iterator` and you need to write the trait typedefs explicitly:

```
class ColumnIterator {
public:
    using iterator_category = std::random_iterator_tag;
    using value_type = Column;
    using difference_type = std::ptrdiff_t;
    using pointer = Column*;
    using reference = Column&;

    // ...
};
```

While you have to write more code, it's much easier to read and it's less error-prone.

For example, the referencing paper mentions the following example from The Standard Library:

```
template <class T, class charT = char, class traits = char_traits<charT> >
class ostream_iterator:
    public iterator<output_iterator_tag, void, void, void, void>;
```

In the above code, it's not clear what all of those four void types mean in the definition.

Additionally `std::iterator` could lead to complicated errors in name lookup. Especially if you happen to use the same name in the derived iterator as in the base class.

You can read more information in the paper [P0174R2 - Deprecating Vestigial Library Parts in C++17¹⁴](https://wg21.link/p0174r2).

¹⁴<https://wg21.link/p0174r2>

Polymorphic Allocator, `pmr`

The polymorphic allocator is an enhancement to the standard allocator from the Standard Library.

In short, a polymorphic allocator conforms to the rules of an allocator from the Standard Library, but at its core, it uses a memory resource object to perform the memory management. Polymorphic Allocator contains a pointer to a memory resource class, and that's why it can use a virtual method dispatch. You can change the memory resource at runtime while keeping the type of the allocator.

All the types for polymorphic allocators live in a separate namespace `std::pmr` (PMR stands for Polymorphic Memory Resource), in the `<memory_resource>` header.

Core elements of `pmr`:

- `std::pmr::memory_resource` - is an abstract base class for all other implementations. It defines the following pure virtual methods: `do_allocate`, `do_deallocate` and `do_is_equal`.
- `std::pmr::polymorphic_allocator` - is an implementation of a standard allocator that uses `memory_resource` object to perform memory allocations and deallocations.
- global memory resources accessed by `new_delete_resource()` and `null_memory_resource()`
- a set of predefined memory pool resource classes: `synchronized_pool_resource`, `unsynchronized_pool_resource`, and `monotonic_buffer_resource`
- template specialisations of the standard containers with polymorphic allocator, for example `std::pmr::vector`, `std::pmr::string`, `std::pmr::map` and others. Each specialisation is defined in the same header file as the corresponding container.

Here's a short overview of the predefined memory resources:

Resource	Description
<code>new_delete_resource()</code>	a free function that returns a pointer to a global "default" memory resource. It manages memory with the global <code>new</code> and <code>delete</code>
<code>null_delete_resource()</code>	a free function that returns a pointer to a global "null" memory resource which throws <code>std::bad_alloc</code> on every allocation
<code>synchronized_pool_resource</code>	thread-safe allocator that manages pools of different sizes. Each pool is a set of chunks that are divided into blocks of uniform size.
<code>unsynchronized_pool_resource</code>	non-thread-safe <code>pool_resource</code>
<code>monotonic_buffer_resource</code>	non-thread-safe, fast, special purpose resource that gets memory from a preallocated buffer, but doesn't release it with deallocation.

It's also worth mentioning that pool resources (including `monotonic_buffer_resource`) can be chained. So that if there's no available memory in a pool, the allocator will allocate from the "upstream" resource.

Below you can find a simple example of `monotonic_buffer_resource` and `pmr::vector`:

```
// Chapter STL Other/pmr_monotonic_resource.cpp
#include <iostream>
#include <memory_resource>    // pmr core types
#include <vector>              // pmr::vector

int main() {
    char buffer[64] = {}; // a small buffer on the stack
    std::fill_n(std::begin(buffer), std::size(buffer) - 1, '_');
    std::cout << buffer << '\n';

    std::pmr::monotonic_buffer_resource pool{ std::data(buffer), std::size(buffer) };

    std::pmr::vector<char> vec{ &pool };
    for (char ch = 'a'; ch <= 'z'; ++ch)
        vec.push_back(ch);

    std::cout << buffer << '\n';
}
```

Possible output:

```
-----
aababcbcdabcbdefghabcbdefghijklmnopabcbdefghijklmnopqrstuvwxyz_____
```

In the above example, we use a monotonic buffer resource initialised with a memory chunk from the stack. By using a simple `char buffer[]` array, we can easily print the contents of the “memory”. The vector gets memory from the pool, and if there’s no more space available, it will ask for memory from the “upstream” resource (which is `new_delete_resource` by default). The example shows vector reallocations when there’s a need to insert more elements. Each time the vector gets more space, so it eventually fits all of the letters.

More Information

This section only touched upon the idea of polymorphic allocators and memory resources. If you want to learn more about this topic, there’s a long chapter in [C++17 The Complete Guide](#)¹⁵ by Nicolai Josuttis. There are also several conference talks, for example [Allocators: The Good Parts](#) by Pablo Halpern¹⁶ from CppCon 2017.

See more information in [P0220R1](#)¹⁷ and [P0337R0](#)¹⁸.

¹⁵<https://leanpub.com/cpp17>

¹⁶<https://channel9.msdn.com/Events/GoingNative/CppCon-2017/119>

¹⁷<https://wg21.link/p0220r1>

¹⁸<https://wg21.link/p0337r0>

Compiler support

Feature	GCC	Clang	MSVC
<code>std::byte</code>	7.1	5.0	VS 2017 15.3
Improvements for Maps and Sets	7.0	3.9	VS 2017 15.5
<code>insert_or_assign()/try_emplace()</code> for maps	6.1	3.7	VS 2017 15
Emplace Return Type	7.1	4.0	VS 2017 15.3
Sampling algorithms	7.1	In Progress	VS 2017 15
<code>gcd</code> and <code>lcm</code>	7.1	4.0	VS 2017 15.3
<code>clamp</code>	7.1	3.9	VS 2015.3
Special Mathematical Functions	7.1	Not yet	VS 2017 15.7
Shared Pointers and Arrays	7.1	In Progress	VS 2017 15.5
Non-member <code>size()</code> , <code>data()</code> and <code>empty()</code>	6.1	3.6	VS 2015
<code>constexpr</code> Additions to the Standard Library	7.1	4.0	VS 2017 15.3
<code>scoped_lock</code>	7.1	5.0	VS 2017 15.3
Polymorphic Allocator & Memory Resource	9.1	In Progress	VS 2017 15.6

Part 3 - More Examples and Use Cases

In the first two parts of the book you've seen Language and the Standard Library features. Most of the time they were presented isolated from each other. However, the real power of each new C++ Standard comes from the composition of the new building blocks. This part will lead you through several larger examples where multiple C++17 elements were used together.

In this part you'll learn:

- How to refactor code with `std::optional` and `std::variant`
- How `if constexpr` simplifies the complex meta-programming code.
- What are the uses for the `[[nodiscard]]` attribute.
- How to work with the filesystem and parallel algorithms and improve the performance of CSV Reader application.

15. Refactoring with `std::optional` and `std::variant`

`std::variant` and `std::optional` are called “vocabulary” types because you can leverage them to convey more design information. Your code can be much more compact and more expressive.

This chapter will show you one example of how `std::optional` and `std::variant` can help with the refactoring of one function. We’ll start with some legacy code, and through several steps, we’ll arrive with a much better solution. To give you a better understanding you’ll see the pros and cons of each step.

The Use Case

Consider a function that takes the current mouse selection for a game. The function scans the selected range and computes several outputs:

- the number of animating objects
- if there are any civil units in the selection
- if there are any combat units in the selection

The existing code looks like this:

```
class ObjSelection
{
public:
    bool IsValid() const { return true; }
    // more code...
};

bool CheckSelectionVer1(const ObjSelection &objList,
                       bool *pOutAnyCivilUnits,
                       bool *pOutAnyCombatUnits,
                       int *pOutNumAnimating);
```

As you can see above, the function uses a lot of output parameters (in the form of raw pointers), and it returns true/false to indicate success (for example the input selection might be invalid).

The implementation of the function is not relevant now, but here's an example code that calls this function:

```
ObjSelection sel;

bool anyCivilUnits { false };
bool anyCombatUnits { false };
int numAnimating { 0 };
if (CheckSelectionVer1(sel, &anyCivilUnits, &anyCombatUnits, &numAnimating))
{
    // ...
}
```

How can we improve the function?

There might be several things:

- Look at the caller's code: we have to create all the variables that will hold the outputs. It definitely generates code duplication if you call the function in many places.
- Output parameters: Core guidelines suggests not to use them.
 - [F.20: For “out” output values, prefer return values to output parameters¹](#)
- If you have raw pointers you have to check if they are valid. You might get away with the checks if you use references for the output parameters.
- What about extending the function? What if you need to add another output param?

Anything else?

How would you refactor this?

Motivated by Core Guidelines and new C++17 features, here's the plan for how we can improve this code:

1. Refactor output parameters into a tuple that will be returned.
2. Refactor tuple into a separate struct and reduce the tuple to a pair.
3. Use `std::optional` to express if the value was computed or not.
4. Use `std::variant` to convey not only the optional result but also the full error information.

The Tuple Version

The first step is to convert the output parameters into a tuple and return it from the function.

According to [F.21: To return multiple “out” values, prefer returning a tuple or struct²](#):

A return value is self-documenting as an “output-only” value. Note that C++ does have multiple return values, by the convention of using a tuple (including pair), possibly with the extra convenience of tie at the call site.

After the change the code might look like this:

¹<https://isocpp.github.io/CppCoreGuidelines/CppCoreGuidelines#f20-for-out-output-values-prefer-return-values-to-output-parameters>

²<https://isocpp.github.io/CppCoreGuidelines/CppCoreGuidelines#f21-to-return-multiple-out-values-prefer-returning-a-tuple-or-struct>

```
std::tuple<bool, bool, bool, int>
CheckSelectionVer2(const ObjSelection &objList)
{
    if (!objList.IsValid())
        return {false, false, false, 0};

    // local variables:
    int numCivilUnits = 0;
    int numCombat = 0;
    int numAnimating = 0;

    // scan...

    return {true, numCivilUnits > 0, numCombat > 0, numAnimating };
}
```

A bit better... isn't it?

- No need to check raw pointers
- Code is quite expressive

What's more on the caller site, you can use Structured Bindings to wrap the returned tuple:

```
auto [ok, anyCivil, anyCombat, numAnim] = CheckSelectionVer2(sel);
if (ok)
{
    // ...
}
```

Unfortunately, this version might not be the best one. For example, there's a risk of forgetting the order of outputs from the tuple.

The problem of function extensions is also still present. So when you'd like to add another output value, you have to extend this tuple and the caller site.

We can fix this with one further step: instead of a tuple use a structure (as also suggested by the Core Guidelines).

A Separate Structure

The outputs seem to represent related data. That's why it's probably a good idea to wrap them into a struct called `SelectionData`.

```

struct SelectionData
{
    bool anyCivilUnits { false };
    bool anyCombatUnits { false };
    int numAnimating { 0 };
};

```

And then you can rewrite the function into:

```

std::pair<bool, SelectionData> CheckSelectionVer3(const ObjSelection &objList)
{
    SelectionData out;

    if (!objList.IsValid())
        return {false, out};

    // scan...

    return {true, out};
}

```

And the caller site:

```

if (auto [ok, selData] = CheckSelectionVer3(sel); ok)
{
    // ...
}

```

The code uses `std::pair` so we still preserve the success flag, it's not the part of the new `struct`.

The main advantage that we achieved here is the improved code structure and extensibility. If you want to add a new parameter, then extend the structure. Previously - with a list of output parameters in the function declaration - you'd have to update much more code.

But isn't `std::pair<bool, MyType>` similar to the concept of `std::optional`?

With `std::optional`

From the `std::optional` chapter:

`std::optional` is a wrapper type to express “null-able” types. It either contains a value, or it's empty. It doesn't use any extra memory allocation.

That seems to be a good choice for our code. We can remove the `ok` variable and rely on the semantics of the `optional`.

The new version of the code:

```
std::optional<SelectionData> CheckSelectionVer4(const ObjSelection &objList)
{
    if (!objList.IsValid())
        return std::nullopt;

    SelectionData out;

    // scan...

    return out;
}
```

And the caller site:

```
if (auto ret = CheckSelectionVer4(sel); ret.has_value())
{
    // access via *ret or even ret->
    // ret->numAnimating
}
```

What are the advantages of the optional version?

- clean and expressive form - `optional` expresses nullable types.
- efficiency - implementations of `optional` are not permitted to use additional storage, such as dynamic memory. The contained value shall be allocated in a region of the optional storage suitably aligned for the type `T`.

With `std::variant`

The last implementation with `std::optional` omits one crucial aspect: error handling. There's no way to know the reason why a value wasn't computed. For example with the version where `std::pair` was used, we were able to return an error code to indicate the reason. What can we do about that?

If you need full information about the error that might occur in the function, you can think about an alternative approach with `std::variant`.


```
enum class [[nodiscard]] ErrorCode
{
    InvalidSelection,
    Undefined
};

variant<SelectionData, ErrorCode> CheckSelectionVer5(const ObjSelection &objList)
{
    if (!objList.IsValid())
        return ErrorCode::InvalidSelection;

    SelectionData out;
    // scan...

    return out;
}
```

As you see the code uses `std::variant` with two possible alternatives: either `SelectionData` or `ErrorCode`. It's almost like a pair, except that you'll always see one active value.

You can use the above implementation:

```
if (auto retV5 = CheckSelectionVer5(sel);
    std::holds_alternative<SelectionData>(retV5))
{
    std::cout << "ok..."
               << std::get<SelectionData>(retV5).numAnimating << '\n';
}
else
{
    switch (std::get<ErrorCode>(retV5))
    {
        case ErrorCode::InvalidSelection:
            std::cerr << "Invalid Selection!\n";
            break;
        case ErrorCode::Undefined:
            std::cerr << "Undefined Error!\n";
            break;
    }
}
```

As you can see, with `std::variant` you have even more information than when `std::optional` was used. You can return error codes and respond to possible failures.



`std::variant<ValueType, ErrorCode>` might be a possible implementation of `std::expected` - a new vocabulary type that might go into the future version of The Standard Library.

Wrap up

You can play with the code in:

Chapter Refactoring With Optional And Variant/`refactoring_optional_variant.cpp`.

In this chapter, you've seen how to refactor lots of ugly-looking output parameters to a nicer `std::optional` version. The optional wrapper clearly expresses that the computed value might be absent. Also, you've seen how to wrap several function parameters into a separate struct. Having one separate type lets you easily extend the code while keeping the logical structure untouched.

And finally, if you need the full information about errors inside a function, then you might also consider an alternative with `std::variant`. This type gives you a chance to return a full error code.

16. Enforcing Code Contracts With `[[nodiscard]]`

C++17 brought a few more standard attributes. By using those extra annotations you can make your code not only readable to other developers but also the compiler can use this knowledge. For example, it might produce more warnings about potential mistakes. Or the opposite: it might avoid a warning generation because it will notice a proper intention (for example with `[[maybe_unused]]`).

In this chapter you'll see how one attribute - `[[nodiscard]]` - can be used to provide better safety in the code.

Introduction

The `[[nodiscard]]` attribute was mentioned in the [Attributes Chapter](#), but here's a simple example to recall its properties.

The attribute is used to mark the return value of functions:

```
[[nodiscard]] int Compute();
```

When you call such function and don't assign the result:

```
void Foo() {  
    Compute();  
}
```

You should get the following (or a similar) warning:

```
warning: ignoring return value of 'int Compute()',  
declared with attribute nodiscard
```

We can go further and not just mark the return value, but a whole type:

```
[[nodiscard]] struct SuperImportantType { }  
SuperImportantType CalcSuperImportant();  
SuperImportantType OtherFoo();  
SuperImportantType Calculate();
```

and you'll get a warning whenever you call any function that returns `SuperImportantType`.

In other words, you can enforce the code contract for a function, so that the caller won't skip the returned value. Sometimes such omission might cause you a bug, so using `[[nodiscard]]` will improve code safety.



The compiler will generate a warning, but usually it's a good practice to enable "treat warnings as errors" when building the code. `/WX` in MSVC or `-Werror` in GCC. Errors stop the compilation process, so a programmer need to take some action and fix the code.

Where Can It Be Used?

Attributes are a standardised way of annotating the code. They are optional, but they might help the compiler to optimize code, detect possible errors or just express the programmer's intentions in a clear way.

Here are a few places where `[[nodiscard]]` can be potentially handy:

Errors

One crucial use case for `[[nodiscard]]` are error codes.

How many times have you forgotten to check a returned error code from a function? (Crucial if you don't rely on exceptions).

Here's some code:

```
enum class [[nodiscard]] ErrorCode {  
    OK,  
    Fatal,  
    System,  
    FileIssue  
};
```

And if we have several functions:

```
ErrorCode OpenFile(std::string_view fileName);  
ErrorCode SendEmail(std::string_view sendto,  
                    std::string_view text);  
ErrorCode SystemCall(std::string_view text);
```

Now every time you'd like to call such functions you're "forced" to check the return value.

Often you might see code where a developer checks only a few function calls, while other function invocations are left unchecked. That creates inconsistencies and can lead to some severe runtime errors.

You think your method is doing fine (because N (of M) called functions returned OK), but something still is failing. You check it with the debugger, and you notice that the Y function returns FAIL and you haven't checked it.

Should you use `[[nodiscard]]` to mark the error type or maybe some essential functions only?

For error codes that are visible through the whole application that might be the right thing to do. Of course when your function returns just `bool` then you can only mark the function, and not the type (or you can create a typedef/alias and then mark it with `[[nodiscard]]`).

Factories / Handles

Another important type of functions where `[[nodiscard]]` adds value are "factories".

Every time you call "make/create/build" you don't want to skip the returned value. Maybe that's a very obvious thing, but there's a possibility (especially when doing some refactoring), to forget, or comment out.

```
[[nodiscard]] Foo MakeFoo();
```

When Returning Non-Trivial Types?

What about such code:

```
std::vector<std::string> GenerateNames();
```

The returned type seems to be heavy, so usually, it means that you have to use it later. On the other hand, even `int` might be heavy regarding semantics of the given problem.

Code With No Side Effects?

The code in the previous section:

```
std::vector<std::string> GenerateNames(); // no side effects...
```

is also an example of a function with no side effects - no global state is changed during the call. In that case, we need to do something with the returned value. Otherwise, the function call can be removed/optimised from the code.

Everywhere?!

There's a paper that might be a "guide" [P0600R0 - `\[\[nodiscard\]\]` in the Library](https://wg21.link/p0600r0)¹. The proposal didn't make into C++17, but was voted into C++20. It suggests a few places where the attribute should be applied.

For existing API's:

- not using the return value always is a "huge mistake" (e.g. always resulting in resource leak)
- not using the return value is a source of trouble and easily can happen (not obvious that something is wrong)

For new API's (not been in the C++ standard yet):

- not using the return value is usually an error.

¹<https://wg21.link/p0600r0>

Here are a few examples where the new attribute should be added:

- `malloc()/new/allocate` - expensive call, usually not using the return value is a resource leak
- `std::async()` - not using the return value makes the call synchronous, which might be hard to detect.

On the other hand such function as `top()` is questionable, as “not very useful, but no danger and such code might exist”

It’s probably a good idea not to add `[[nodiscard]]` in all places of your code but focus on the critical places. Possibly, as mentioned before, error codes and factories are a good place to start.

How to Ignore `[[nodiscard]]`

With `[[nodiscard]]` you should use the returned value - by assigning it to a variable or by using it directly. If you forget, you’ll get an “unused variable” warning.

There are situations where you might want to suppress such a warning. To do that you can use another attribute from C++17: `[[maybe_unused]]`:

```
[[nodiscard]] int Compute() { return 42; }  
[[maybe_unused]] auto t = Compute();
```

Also, as mentioned in the Attributes Chapter, you can cast the function call to `void` and the compiler will think you “used” the value:

```
[[nodiscard]] int Compute();  
static_cast<void>(Compute()); // used
```

Another good alternative might be to write a separate function that wraps the results and pretends to use it²:

```
template <class T> inline void discard_on_purpose(T&&) {}  
  
// use case:  
discard_on_purpose(Compute());
```



Be careful with the techniques to avoid warnings with `[[nodiscard]]`. It’s better to follow the rules of the attribute rather than artificially avoid them.

²Suggested by Arne Mertz

Before C++17

Most of the attributes that went into the standardised `[[attrib]]` come from compiler extensions, same happened with `[[nodiscard]]`.

For example, in GCC/Clang, you can use:

```
__attribute__((warn_unused_result))
```

in MSVC:

`_Check_return_` - see more in [MSDN: Annotating Function Behavior](https://msdn.microsoft.com/en-us/library/jj159529.aspx)³.

Summary

To sum up: `[[nodiscard]]` is an excellent addition to all the important code: public APIs, safety-critical systems, etc. Adding this attribute will at least enforce the code contract, and a compiler will help you detect bugs - at compile time, rather than finding it in runtime.

³<https://msdn.microsoft.com/en-us/library/jj159529.aspx>

17. Replacing `enable_if` with `if constexpr` - Factory with Variable Arguments

One of the most powerful language features that we get with C++17 is the compile-time `if` in the form of `if constexpr`. It allows you to check, at compile time, a condition and depending on the result the code is rejected from the further steps of the compilation.

In this chapter, you'll see one example of how this new feature can simplify the code.

The Problem

In the item 18 of Effective Modern C++ Scott Meyers described a method called makeInvestment:

```
template<typename... Ts>
std::unique_ptr<Investment>
makeInvestment(Ts&&... params);
```

There's a factory method that creates derived classes of Investment, and the main advantage is that it supports a variable number of arguments!

For example, here are the proposed derived types:

```
// Chapter If constexpr Factory/variable_factory.cpp
// base type:
class Investment
{
public:
    virtual ~Investment() { }

    virtual void calcRisk() = 0;
};

class Stock : public Investment
{
public:
    explicit Stock(const std::string&) { }

    void calcRisk() override { }
};

class Bond : public Investment
{
public:
    explicit Bond(const std::string&, const std::string&, int) { }

    void calcRisk() override { }
};

class RealEstate : public Investment
{
public:
```

```

    explicit RealEstate(const std::string&, double, int) { }

    void calcRisk() override { }
};

```

The code from the book was too idealistic and it worked until all your classes have the same number and types of input parameters:

[Scott Meyers: Modification History and Errata List for Effective Modern C++¹](#):

The makeInvestment interface is unrealistic, because it implies that all derived object types can be created from the same types of arguments. This is especially apparent in the sample implementation code, where arguments are perfect-forwarded to all derived class constructors.

For example if you had a constructor that needed two arguments and one constructor with three arguments, then the code might not compile:

```

// pseudo code:
Bond(int, int, int) { }
Stock(double, double) { }
make(args...)
{
    if (bond)
        new Bond(args...);
    else if (stock)
        new Stock(args...)
}

```

Now, if you write make(bond, 1, 2, 3) - then the else statement won't compile - as there no Stock(1, 2, 3) available! To make it work we need a compile time if statement that rejects parts of the code that don't match a condition.

On my blog, with the help of one reader, we proposed one working solution (you can read more in [Bartek's coding blog: Nice C++ Factory Implementation 2²](#)).

Here's the code that works:

¹<http://www.aristeia.com/BookErrata/emc++-errata.html>

²<http://www.bfilipek.com/2016/03/nice-c-factory-implementation-2.html>

```

template <typename... Ts>
unique_ptr<Investment>
makeInvestment(const string &name, Ts&&... params)
{
    unique_ptr<Investment> pInv;

    if (name == "Stock")
        pInv = constructArgs<Stock, Ts...>(forward<Ts>(params)...);
    else if (name == "Bond")
        pInv = constructArgs<Bond, Ts...>(forward<Ts>(params)...);
    else if (name == "RealEstate")
        pInv = constructArgs<RealEstate, Ts...>(forward<Ts>(params)...);

    // call additional methods to init pInv...

    return pInv;
}

```

As you can see the “magic” happens inside constructArgs function.

The main idea is to return unique_ptr<Type> when Type is constructible from a given set of attributes and nullptr when it’s not.

Before C++17

In the previous solution (pre C++17) std::enable_if had to be used:

```

// before C++17
template <typename Concrete, typename... Ts>
enable_if_t<is_constructible<Concrete, Ts...>::value, unique_ptr<Concrete>>
constructArgsOld(Ts&&... params)
{
    return std::make_unique<Concrete>(forward<Ts>(params)...);
}

template <typename Concrete, typename... Ts>
enable_if_t<!is_constructible<Concrete, Ts...>::value, unique_ptr<Concrete> >
constructArgsOld(...)
{
    return nullptr;
}

```

`std::is_constructible` - allows us to test if a list of arguments could be used to create a given type.



Just a quick reminder about **`enable_if`**

`enable_if` (and `enable_if_v` since C++14). It has the following syntax:

```
template< bool B, class T = void >
struct enable_if;
```

`enable_if` will evaluate to `T` if the input condition `B` is true. Otherwise, it's SFINAE and a particular function overload is removed from the overload set.

What's more, in C++17 there's a helper:

```
is_constructible_v = is_constructible<T, Args...>::value;
```

Potentially, the code should be a bit shorter.

Still, using `enable_if` looks ugly and complicated. How about C++17 version?

With `if constexpr`

Here's the updated version:

```
template <typename Concrete, typename... Ts>
unique_ptr<Concrete> constructArgs(Ts&&... params)
{
    if constexpr (is_constructible_v<Concrete, Ts...>)
        return make_unique<Concrete>(forward<Ts>(params)...);
    else
        return nullptr;
}
```

We can even extend it with some little logging features, using fold expression:

```

template <typename Concrete, typename... Ts>
std::unique_ptr<Concrete> constructArgs(Ts&&... params)
{
    cout << __func__ << ": ";
    // fold expression:
    ((cout << params << ", "), ...);
    cout << '\n';

    if constexpr (std::is_constructible_v<Concrete, Ts...>)
        return make_unique<Concrete>(forward<Ts>(params)...);
    else
        return nullptr;
}

```

All the complicated syntax of `enable_if` went away; we don't even need a function overload for the else case. We can now wrap expressive code in just one function.

`if constexpr` evaluates the condition and only one block will be compiled. In our case, if a type is constructible from a given set of attributes, then we'll compile `make_unique` call. If not, then `nullptr` is returned (and `make_unique` is not even compiled).

You can play with the code in:

Chapter If Constexpr Factory/variable_factory.cpp

Summary

In this chapter you've seen how `if constexpr` can make code much clearer and more expressive. Before C++17 you could use `enable_if` techniques (SFINAE) or tag dispatching. Those options usually generated complicated code which might be hard to read by novice and non-meta-programming experts. `if constexpr` lowers the expertise level needed to write template code effectively.

18. How to Parallelise CSV Reader

In the [Parallel Algorithms](#) chapter we learned how to speed up code by running it automatically on multiple threads. That chapter showed a few smaller examples and benchmarks. It would also be a good idea to see more extensive applications and how they can benefit from parallelisation, so that's where we turn to now.

In the next pages, you'll see how to build a tool that works on CSV files, parses lines into sales records and then performs calculations on the data. You'll see how easy it is to add parallel execution to selected algorithms and have a performance improvement across the whole application. In the end, we'll discuss problems that we found along the way and possible future enhancements.

In this chapter you'll learn:

- How to build an application that loads CSV files
- How to efficiently use parallel algorithms
- How to use `std::filesystem` library to gather required files
- How to use other C++17 library features like `std::optional`, conversion routines - `std::from_chars` and `string_view`

Introduction and Requirements

Imagine you're working with some sales data and one task is to calculate a sum of orders for some products. Your shopping system is elementary and instead of a database, you have CSV files with the order data. There's one file per product.

Take this example of book sales:

date	coupon code	price	discount	quantity
5-12-2018		10.0	0	2
5-12-2018		10.0	0	1
6-12-2018	Santa	10.0	0.25	1
7-12-2018		10.0	0	1

Each line shows a book sale on a specific date. For example 5th Dec there were three sales, 10\$ each, and one person bought two books. On 6th Dec we had one sale with a coupon code.

The data is encoded as a CSV file: `sales/book.csv`:

```
5-12-2018;;10.0;0;2;
5-12-2018;;10.0;0;1;
6-12-2018;Santa;10.0;0.25;1;
7-12-2018;;10.0;0;1;
```

The application should read the data and then calculate the sum, so in the above case we have

```
sum = 10*2+10*1+      // 5th Dec
      10*(1-0.25)*1 + // 6th Dec with 25% coupon
      10*1;           // 7th Dec
```

For the above sales data, the final sum is 47.5\$.

Here are requirements of the application we want to build:

- The application loads all CSV files in a given folder - read from the first argument in the command line
- The files might contain thousands of records but will fit into memory. There's no need to provide extra support for huge files
- Optionally, the application reads the start and end dates from the second and the third command line argument
- Each CSV line has the following structure: `date;coupon code;unit price;quantity;discount;`
- The application sums all orders between given dates and prints the sum to the standard output

We'll implement the serial version first and then we'll try to make it parallel.

The Serial Version

For the first step, we'll cover a serial version of the application. This allows you to understand the core parts of the system and see how the tool works.

The code doesn't fit easily on a single page so you can have a look at it in the following file:

CSV Chapter/csv_reader.cpp

In the next sections, we'll explore the core parts of the application.

The Main

Let's start with the `main()` function.

```
1  int main(int argc, const char** argv) {
2      if (argc <= 1)
3          return 1;
4
5      try {
6          const auto paths = CollectPaths(argv[1]);
7
8          if (paths.empty()) {
9              std::cout << "No files to process...\n";
10             return 0;
11         }
12
13         const auto startDate = argc > 2 ? Date(argv[2]) : Date();
14         const auto endDate = argc > 3 ? Date(argv[3]) : Date();
15
16         const auto results = CalcResults(paths, startDate, endDate);
17
18         ShowResults(startDate, endDate, results);
19     }
20     catch (const std::filesystem::filesystem_error& err) {
21         std::cerr << "filesystem error! " << err.what() << '\n';
22     }
23     catch (const std::runtime_error& err) {
24         std::cerr << "runtime error! " << err.what() << '\n';
25     }
26
27     return 0;
28 }
```

Once we're sure that there are enough arguments in the command line we enter the main scope where all the processing happens:

- line 6 - gather all the files to process - in `CollectPaths()`
- line 16 - convert data from the files into record data and calculate the results - in `CalcResults()`
- line 18 - show the results on the output - in `ShowResults()`

The code relies on exceptions across the whole application.

The paths are collected using `directory_iterator` from the `std::filesystem` library:

```
bool IsCSVFile(const fs::path &p)
{
    return fs::is_regular_file(p) && p.extension() == CSV_EXTENSION;
}

[[nodiscard]] std::vector<fs::path> CollectPaths(const fs::path& startPath)
{
    std::vector<fs::path> paths;
    fs::directory_iterator dirpos{ startPath };
    std::copy_if(fs::begin(dirpos), fs::end(dirpos), std::back_inserter(paths),
                IsCSVFile);
    return paths;
}
```

As in other filesystem examples, the namespace `fs` is an alias for `std::filesystem`.

With `directory_iterator` we can easily iterate over a given directory. By using `copy_if` we can filter out unwanted files and select only those with a CSV extension. Notice how easy it is to get the elements of the path and check files' properties.

Going back to `main()`, we check if there are any files to process (line 8).

Then, in lines 13 and 14, we parse the optional dates: `startDate` and `endDate` are read from `argv[2]` and `argv[3]`.

The dates are stored in a helper class `Date` that lets you convert from strings with a simple format of `Day-Month-Year` or `Year-Month-Day`. The class also supports comparison of dates. This will help us check whether a given order fits between selected dates.

Now, all of the computations and printouts are contained in lines:

```
const auto results = CalcResults(paths, startDate, endDate);
ShowResults(results, startDate, endDate);
```

`CalcResults()` implements the core requirements of the application:

- converting data from the file into a list of records to process
- calculating a sum of records between given dates

```

struct Result
{
    std::string mFilename;
    double mSum{ 0.0 };
};

[[nodiscard]] std::vector<Result>
CalcResults(const std::vector<fs::path>& paths, Date startDate, Date endDate)
{
    std::vector<Result> results;
    for (const auto& p : paths)
    {
        const auto records = LoadRecords(p);

        const auto totalValue = CalcTotalOrder(records, startDate, endDate);
        results.push_back({ p.string(), totalValue });
    }
    return results;
}

```

The code loads records from each CSV file, then calculates the sum of those records. The results (along with the name of the file) are stored in the output vector.

We can now reveal the code behind the two essential methods `LoadRecords` and `CalcTotalOrder`.

Converting Lines into Records

`LoadRecords` is a function that takes a filename as an argument, reads the contents into `std::string` and then performs the conversion:

```
[[nodiscard]] std::vector<OrderRecord> LoadRecords(const fs::path& filename)
{
    const auto content = GetFileContents(filename);

    const auto lines = SplitLines(content);

    return LinesToRecords(lines);
}
```

We assume that the files are small enough to fit into RAM, so there's no need to process them in chunks.

The core task is to split that one large string into lines and then convert them into a collection of Records.

If you look into the code you can see that `content` is `std::string`, but `lines` is a vector of `std::string_view`. Views are used for optimisation. We guarantee to hold the large string - the file content - while we process chunks of it (views). This should give us better performance, as there's no need to copy string data.

Eventually, characters are converted into `OrderRecord` representation.

The OrderRecord Class

The main class that is used to compute results is `OrderRecord`. It's a direct representation of a line from a CSV file.

```
class OrderRecord
{
public:
    // constructors...

    double CalcRecordPrice() const noexcept;
    bool CheckDate(const Date& start, const Date& end) const noexcept;

private:
    Date mDate;
    std::string mCouponCode;
    double mUnitPrice{ 0.0 };
    double mDiscount{ 0.0 }; // 0... 1.0
    unsigned int mQuantity{ 0 };
};
```

The conversion

Once we have lines we can convert them one by one into objects:

```
[[nodiscard]] std::vector<OrderRecord>
LinesToRecords(const std::vector<std::string_view>& lines)
{
    std::vector<OrderRecord> outRecords;
    std::transform(lines.begin(), lines.end(),
                   std::back_inserter(outRecords), LineToRecord);

    return outRecords;
}
```

The code above is just a transformation, it uses `LineToRecord` to do the hard work:

```
[[nodiscard]] OrderRecord LineToRecord(std::string_view sv)
{
    const auto cols = SplitString(sv, CSV_DELIM);
    if (cols.size() == static_cast<size_t>(OrderRecord::ENUM_LENGTH))
    {
        const auto unitPrice = TryConvert<double>(cols[OrderRecord::UNIT_PRICE]);
        const auto discount = TryConvert<double>(cols[OrderRecord::DISCOUNT]);
        const auto quantity = TryConvert<unsigned int>(cols[OrderRecord::QUANTITY]);

        if (unitPrice && discount && quantity)
        {
            return { Date(cols[OrderRecord::DATE]),
                     std::string(cols[OrderRecord::COUPON]),
                     *unitPrice,
                     *discount,
                     *quantity };
        }
    }
    throw std::runtime_error("Cannot convert Record from " + std::string(sv));
}
```

Firstly, the line is split into columns, and then we can process each column.

If all elements are converted, then we can build a record.

For conversions of the elements we're using a small utility based on `std::from_chars`:

```

template<typename T>
[[nodiscard]] std::optional<T> TryConvert(std::string_view sv) noexcept
{
    T value{ };
    const auto last = sv.data() + sv.size();
    const auto res = std::from_chars(sv.data(), last, value);
    if (res.ec == std::errc{} && res.ptr == last)
        return value;

    return std::nullopt;
}

```

TryConvert uses `std::from_chars` and returns a converted value if there are no errors. As you remember, to guarantee that all characters were parsed, we also have to check `res.ptr == last`. Otherwise, the conversion might return success for input like “123xxx”.

Calculations

Once all the records are available we can compute their sum:

```

[[nodiscard]] double CalcTotalOrder(const std::vector<OrderRecord>& records,
                                   const Date& startDate, const Date& endDate)
{
    return std::accumulate(std::begin(records), std::end(records), 0.0,
        [&startDate, &endDate](double val, const OrderRecord& rec) {
            if (rec.CheckDate(startDate, endDate))
                return val + rec.CalcRecordPrice();
            else
                return val;
        }
    );
}

```

The code runs on the vector of all records and then calculates the price of each element if they fit between `startDate` and `endDate`. Then they are all summed in `std::accumulate`.

Design Enhancements

The application calculates only the sum of orders, but we could think about adding other things. For example, minimal value, maximum, average order and other statistics.

The code uses a simple approach, loading a file into a string and then creating a temporary vector of lines. We could also enhance this by using a line iterator. It would take a large string and then return a line when you iterate.

Another idea relates to error handling. For example, rather than throwing exceptions we could enhance the conversion step by storing the number of successfully processed records.

Running the Code

The application is ready to compile and we can run it on the example data shown in the introduction.

```
CSVReader.exe sales/
```

This should read a single file `sales/book.csv` and sum up all the records (as no dates were specified):

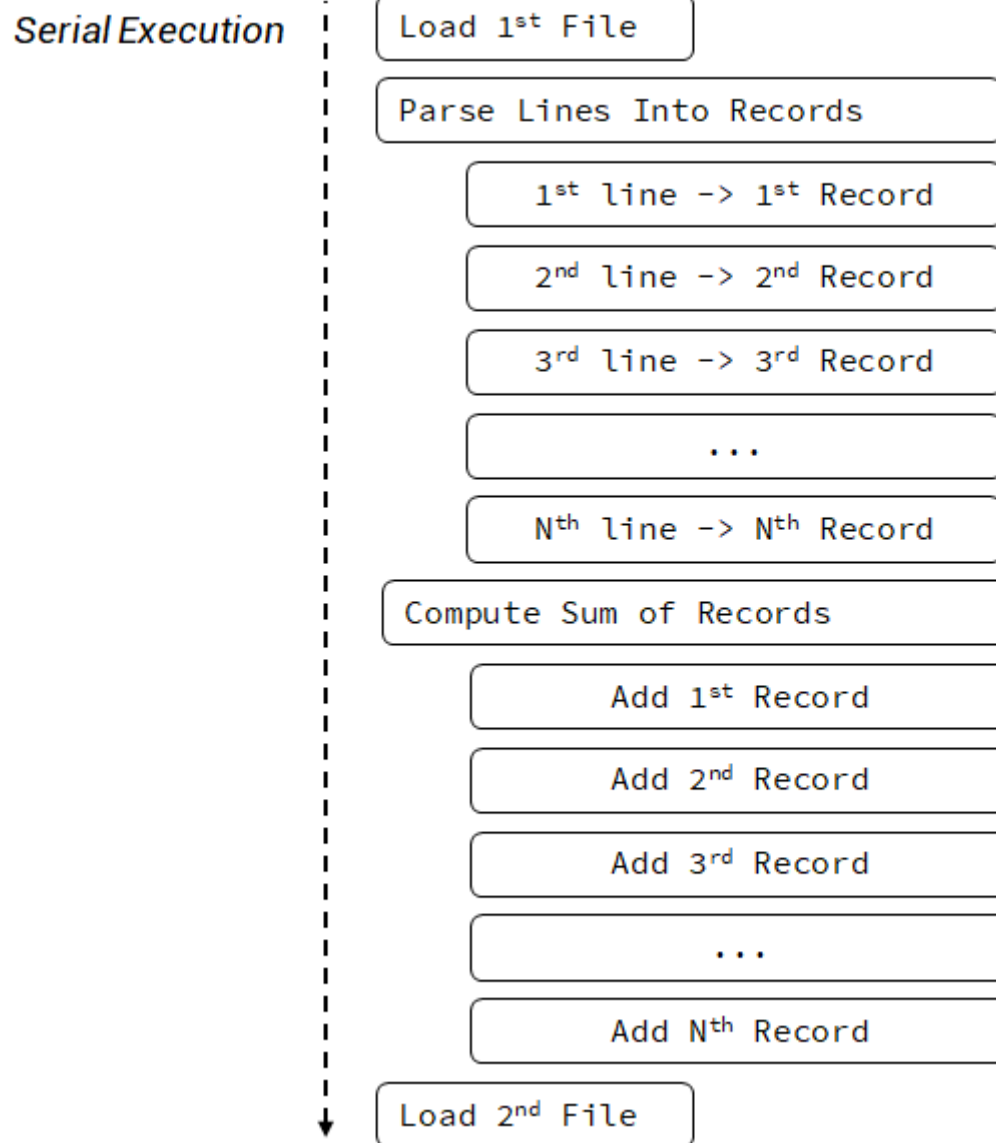
```
.\CalcOrdersSerial.exe .\sales\  
Name Of File      | Total Orders Value  
sales\book.csv   | 47.50  
CalcResults: 3.13 ms  
CalcTotalOrder: 0.01 ms  
Parsing Strings: 0.01 ms
```

The full version of the code also includes timing measurement, so that's why you can see that the operation took around 3ms to complete. The file handling took the longest; calculations and parsing were almost immediate.

In the next sections, you'll see a few simple steps you can take to apply parallel algorithms.

Using Parallel Algorithms

Previously the code was executed in a sequential way. We can illustrate it in the following diagram:

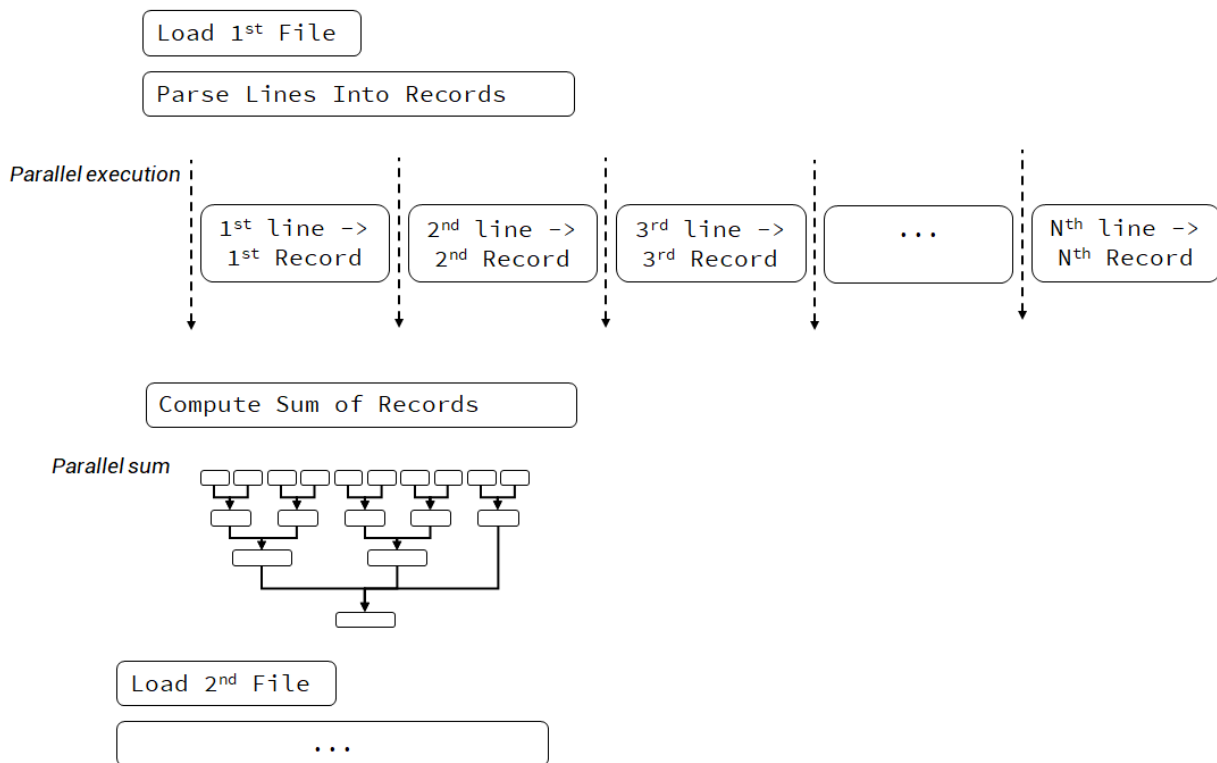


Serial Execution of CSV Reader

We open each file, process it, calculate, then we go to another file. All this happens on a single thread. However, there are several places we can consider using parallel algorithms:

- Where each file can be processed separately
- Where each line of a file can be converted independently into the Record Data
- Where calculations can be enhanced with parallel execution

If we focus on the second and the third options, we can move into the following execution model:



Parallel Execution of CSV Reader

The above diagram shows that we're still processing file one by one, but we use parallel execution while parsing the strings and making the calculations.

When doing the conversion, we have to remember that exceptions won't be re-thrown from our code. Only `std::terminate` will be called.



As of December 2018, only the MSVC compiler supports parallel execution in the Standard Library. The parallel version of the example does not work with GCC or Clang. It's possible to use a third party library like Intel Parallel STL or HPX. You can, however, compile the serial version (`csv_reader.cpp`) and play with the code. It was tested on GCC 8.2. Unfortunately, it fails to compile with Clang.

Data Size & Instruction Count Matters

How to get the best performance with parallel algorithms?

You need two things:

- a lot of data to process
- instructions to keep the CPU busy

We also have to remember one rule:

In general, parallel algorithms do more work, as they introduce the extra cost of managing the parallel execution framework as well as splitting tasks into smaller batches.

First and foremost, we have to think about the size of the data we're operating on. If we have only a few files, with a few dozen records, then we may not gain anything with parallel execution. But if we have lots of files, with hundreds of lines each, then the potential might increase.

The second thing is the instruction count. CPU cores need to compute and not just wait on memory. If your algorithms are memory-bound, then parallel execution might not give any speed-up over the sequential version. In our case, it seems that the parsing strings task is a good match here. The code performs searching on strings and does the numerical conversions, which keeps CPU busy.

Parallel Data Conversion

As previously discussed, we can add parallel execution to the place where we convert the data. We have lots of lines to parse and each parsing is independent.

```
[[nodiscard]] std::vector<OrderRecord>
LinesToRecords(const std::vector<std::string_view>& lines)
{
    std::vector<OrderRecord> outRecords(lines.size());
    std::transform(std::execution::par, std::begin(lines), std::end(lines),
                   std::begin(outRecords), LineToRecord);

    return outRecords;
}
```

Two things need to be changed to the serial version:

- we need to preallocate the vector
- we have to pass `std::execution::par` (or `par_unseq`) as the first argument

The serial code also used `std::transform`, so why cannot we just pass the execution parameter?

We can even compile it... but you should see an error like:

Parallel algorithms require forward iterators or stronger.

The reason is simple: `std::back_inserter` is very handy, but it's not a forward iterator. It inserts elements into the vector, and that causes a vector to be changed (reallocated) by multiple threads. All of the insertions would have to be guarded by some critical section, and thus the overall performance could be weak.

Since we need to preallocate the vector, we have to consider two things:

- we pay for default construction of objects inside a vector, which probably isn't a big deal when objects are relatively small, and their creation is fast.
- on the other hand, the vector is allocated once, and there's no need to grow it (copy, reallocate) as in the case of `std::back_inserter`.

Parallel Calculations

Another place where we can use parallel algorithms is `CalcTotalOrder()`.

Instead of `std::accumulate` we can use `std::transform_reduce`.



As mentioned in the [Parallel Algorithms](#) chapter, the floating point sum operation is not associative. However, in our case, the results should be stable enough to give 2 decimal places of precision. If you need better accuracy and numerical stability, you may be better off using a different method.

```
double CalcTotalOrder(const std::vector<OrderRecord>& records,
                     const Date& startDate, const Date& endDate)
{
    return std::transform_reduce(
        std::execution::par,
        std::begin(records), std::end(records),
        0.0,
        std::plus<>(),
        [&startDate, &endDate](const OrderRecord& rec) {
            if (rec.CheckDate(startDate, endDate))
                return rec.CalcRecordPrice();

            return 0.0;
        }
    );
}
```

We use the transform step of `std::transform_reduce` to “extract” values to sum. We cannot easily use `std::reduce` as it would require us to write a reduction operation that works with two `OrderRecord` objects.

Tests

We can run the two versions on a set of files and compare if the changes brought any improvements in the performance.



Our applications access files, so it's harder to make accurate benchmarks as we can quickly end up in the file system cache. Before major runs of applications, a tool called Use SysInternal's [RAMMap app](http://technet.microsoft.com/en-us/sysinternals/ff700229.aspx)¹ is executed to remove files from the cache. There are also Hard Drive hardware caches which are harder to release without a system reboot.

The application runs on a 6 core/12 thread machine - i7 8700, with a fast SSD drive, Windows 10.

Mid Size Files 1k Lines 10 Files

Let's start with 10 files, 1k lines each. Files are not in the OS cache:

Step	Serial (ms)	Parallel (ms)
All steps	74.05	68.391
CalcTotalOrder	0.02	0.22
Parsing Strings	7.85	2.82

The situation when files are in the system cache:

Step	Serial (ms)	Parallel (ms)
All steps	8.59	4.01
CalcTotalOrder	0.02	0.23
Parsing Strings	7.74	2.73

The first numbers - 74ms and 68ms - come from reading uncached files, while the next two runs were executed without clearing the system cache so you can observe how much speed-up you get by system caches.

The parallel version still reads files sequentially, so we only get a few milliseconds of improvement. Parsing strings (line split and conversion to Records) is now almost 3x faster. The sum calculations are not better as a single threaded version seem to handle sums more efficiently.

Large Set 10k Lines in 10 Files

How about larger input?

Uncached files:

Step	Serial (ms)	Parallel (ms)
All steps	239.96	178.32
CalcTotalOrder	0.2	0.74
Parsing Strings	70.46	15.39

Cached:

¹<http://technet.microsoft.com/en-us/sysinternals/ff700229.aspx>

Step	Serial (ms)	Parallel (ms)
All steps	72.43	18.51
CalcTotalOrder	0.33	0.67
Parsing Strings	70.46	15.56

The more data we process, the better our results. The cost of loading uncached files “hides” slowly behind the time it takes to process the records. In the case of 10k lines we can also see that the parsing strings step is 3.5 times faster; however, the calculations are still slower.

Largest Set 100k Lines in 10 Files

Let's do one more test with the largest files:

Uncached files:

Step	Serial (ms)	Parallel (ms)
All steps	757.07	206.85
CalcTotalOrder	3.03	2.47
Parsing Strings	699.54	143.31

Cached:

Step	Serial (ms)	Parallel (ms)
All steps	729.94	162.49
CalcTotalOrder	3.05	2.16
Parsing Strings	707.34	141.28

In a case of large files (each file is ~2MB), we can see a clear win for the parallel version.

Wrap up & Discussion

The main aim of this chapter was to show how easy it is to use parallel algorithms.

The final code is located in two files:

Chapter CSV Reader/csv_reader.cpp and Chapter CSV Reader/csv_reader_par.cpp for the parallel version.

In most of the cases, all we have to do to add parallel execution is to make sure there's no synchronisation required between the tasks and, if we can, provide forward iterators. That's why when doing the conversion we sometimes needed to preallocate `std::vector` (or other compliant collections) rather than using `std::back_inserter`. Another example is that we cannot iterate in a directory in parallel, as `std::filesystem::directory_iterator` is not a forward iterator.

The next part is to select the proper parallel algorithm. In the case of this example, we replaced `std::accumulate` with `std::transform_reduce` for the calculations. There was no need to change `std::transform` for doing the string parsing - as you only have to use the extra execution policy parameter.

Our application performed a bit better than the serial version. Here are some thoughts we might have:

- Parallel execution needs tasks that are independent. If you have jobs that depend on each other, the performance might be lower than the serial version! This happens due to extra synchronisation steps.
- Your tasks cannot be memory-bound, otherwise CPU will wait for the memory. For example, the string parsing code performed better in parallel as it has many instructions to execute: string search, string conversions.
- You need a lot of data to process to see the performance gain. In our case, each file required several thousands of lines to show any gains over the sequential version.
- Sum calculations didn't show much improvement and there was even worse performance for smaller input. This is because the `std::reduce` algorithm requires extra reduction steps, and also our calculations were elementary. It's possible that, with more statistical computations in the code, we could improve performance.
- The serial version of the code is straightforward and there are places where extra performance could be gained. For example, we might reduce additional copies and temporary vectors. It might also be good to use `std::transform_reduce` with sequential execution in the serial version, as it might be faster than `std::accumulate`. You might consider optimising the serial version first and then making it parallel.
- If you rely on exceptions then you might want to implement a handler for `std::terminate`, as exceptions are not re-thrown in code that is invoked with execution policies.

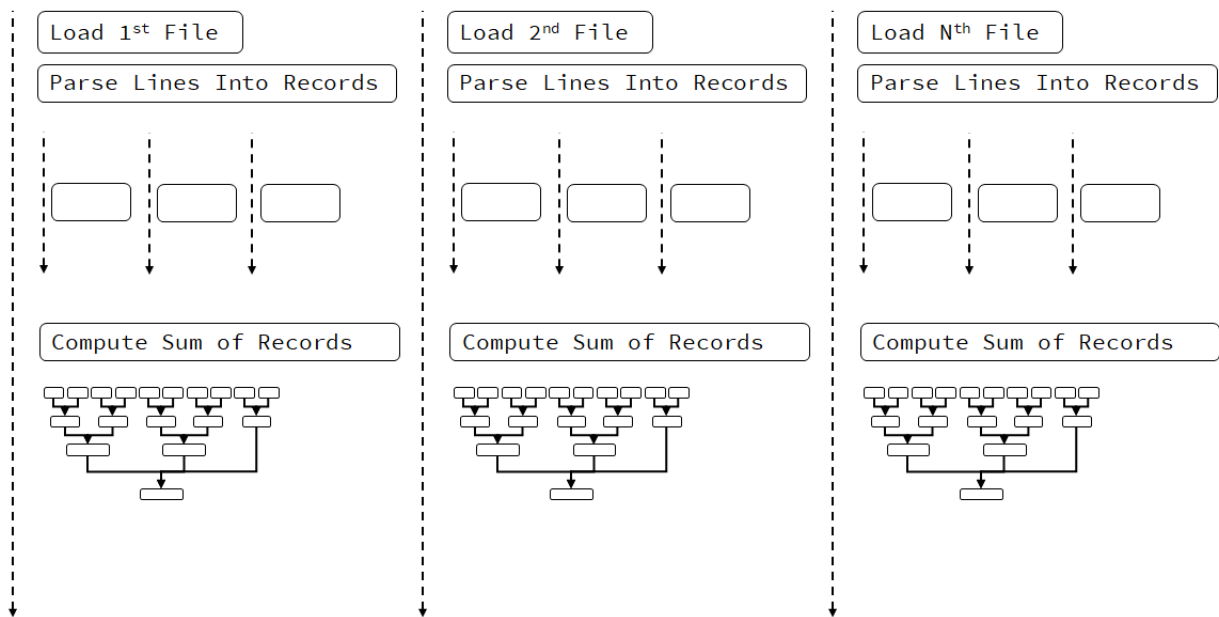
So putting it all together, the experiment with the following summary:

Parallel algorithms can bring extra speed to the application, but they have to be used wisely. They introduce an additional cost of parallel execution framework and it's essential to have lots of tasks that are independent and good for parallelisation. As always, it's important to measure the performance between the versions to be able to select the final approach with confidence.

Additional Modifications and Options

The code in the parallel version skipped one option: parallel access to files. So far we read files one by one, but how about reading separate files from separate threads?

Here's a diagram that illustrates this option:



Parallel Execution of CSV Reader, Reading files in separate threads

In the above diagram, the situation is a bit complicated. If we assume that OS cannot handle multiple file access, then threads will wait on files. But once the files are available, the processing might go in parallel.

If you want to play around with this technique, you can replace `std::execution::seq` in `CalcResults()` with `std::execution::par`. That will allow the compiler to run `LoadRecords()` and `CalcTotalOrder()` in parallel.

Is your system capable of accessing files from separate threads?

In general, the answer might be tricky, as it depends on many elements: hardware, system, and cost of computations, etc. For example on a machine with a fast SSD drive the system can handle several files reads, while on a HDD drive, the performance might be slower. Modern drives also use Native Command Queues, so even if you access from multiple threads the command to the drive will be serial and also rearranged into a more optimal way. We leave the experiments to the readers as this topic goes beyond the scope of this book.

Appendix A - Compiler Support

If you work with the latest version of a compiler like GCC, Clang or MSVC you may assume that C++17 is fully supported (with some exceptions to the STL implementations). GCC implemented the full support in the version 7.0, Clang did it in the version 6.0 and MSVC is conformant as of VS 2017 15.7 For completeness, here you have a list of features and versions of compilers where it was added.

The main resource to be up to date with the status of the features: [CppReference - Compiler Support](#)²

GCC

- [Language](#)³
- [The Library - LibSTDC++](#)⁴

Clang

- [Language](#)⁵
- [The Library - LibC++](#)⁶

VisualStudio - MSVC

- [Announcing: MSVC Conforms to the C++ Standard](#)⁷

²https://en.cppreference.com/w/cpp/compiler_support

³<https://gcc.gnu.org/projects/cxx-status.html#cxx17>

⁴<https://gcc.gnu.org/onlinedocs/libstdc++/manual/status.html#status.iso.2017>

⁵http://clang.llvm.org/cxx_status.html#cxx17

⁶http://libcxx.llvm.org/cxx1z_status.html

⁷<https://blogs.msdn.microsoft.com/vcblog/2018/05/07/announcing-msvc-conforms-to-the-c-standard/>

Compiler Support of C++17 Features

Fixes and Deprecation

Feature	GCC	Clang	MSVC
Removing <code>register</code> keyword	7.0	3.8	VS 2017 15.3
Remove Deprecated <code>operator++(bool)</code>	7.0	3.8	VS 2017 15.3
Removing Deprecated Exception Specifications	7.0	4.0	VS 2017 15.5
Removing <code>auto_ptr</code> , <code>random_shuffle</code> , old <code><functional></code>	No (kept for compatibility)	not yet	VS 2015
Removing trigraphs	5.1	3.5	VS 2010
New auto rules for direct-list-initialisation	5.0	3.8	VS 2015
<code>static_assert</code> with no message	6.0	2.5	VS 2017
Different begin and end types in range-based for	6.0	3.6	VS 2017

Clarification

Feature	GCC	Clang	MSVC
Stricter expression evaluation order	7.0	4.0	VS 2017
Guaranteed copy elision	7.0	4.0	VS 2017 15.6
Exception specifications part of the type system	7.0	4.0	VS 2017 15.5
Dynamic memory allocation for over-aligned data	7.0	4.0	VS 2017 15.5

General Language Features

Feature	GCC	Clang	MSVC
Structured Binding Declarations	7.0	4.0	VS 2017 15.3
Init-statement for <code>if/switch</code>	7.0	3.9	VS 2017 15.3
Inline variables	7.0	3.9	VS 2017 15.5
<code>constexpr</code> Lambda Expressions	7.0	5.0	VS 2017 15.3
Nested namespaces	6.0	3.6	VS 2015

Templates

Feature	GCC	Clang	MSVC
Template argument deduction for class templates	7.0/8.0 ⁸	5.0	VS 2017 15.7
Deduction Guides in the Standard Library	8.0 ⁹	7.0/in progress ¹⁰	VS 2017 15.7

⁸Additional improvements for Template Argument Deduction for Class Templates happened in GCC 8.0, [P0512R0](#).

⁹Deduction Guides are not listed in the [status pages of LibSTDC++](#), so we can assume they were implemented as part of Template argument deduction for class templates.

¹⁰The [status page for LibC++](#) mentions that `<string>`, sequence containers, container adaptors and `<regex>` portions have been implemented so far.

Feature	GCC	Clang	MSVC
Declaring non-type template parameters with auto	7.0	4.0	VS 2017 15.7
Fold expressions	6.0	3.9	VS 2017 15.5
if constexpr	7.0	3.9	VS 2017

Attributes

Feature	GCC	Clang	MSVC
<code>[[fallthrough]]</code>	7.0	3.9	VS 2017 15.0
<code>[[nodiscard]]</code>	7.0	3.9	VS 2017 15.3
<code>[[maybe_unused]]</code>	7.0	3.9	VS 2017 15.3
Attributes for namespaces and enumerators	4.9(namespaces)/6(enums)	3.4	VS 2015 14.0
Ignore unknown attributes	yes	3.9	VS 2015 14.0
Using attribute namespaces without repetition	7.0	3.9	VS 2017 15.3

The Standard Library

Feature	GCC	Clang	MSVC
<code>std::optional</code>	7.1	4.0	VS 2017 15.0
<code>std::variant</code>	7.1	4.0	VS 2017 15.0
<code>std::any</code>	7.1	4.0	VS 2017 15.0
<code>std::string_view</code>	7.1	4.0	VS 2017 15.0
String Searchers	7.1	5.0	VS 2017 15.3
String Conversions	8 (only integral types)	in progress	Partial in VS 2017 15.7
Parallel Algorithms	in progress	in progress	VS 2017 15.7
Filesystem	8.0	in progress	VS 2017 15.7

Other STL changes

Feature	GCC	Clang	MSVC
<code>std::byte</code>	7.1	5.0	VS 2017 15.3
Improvements for Maps and Sets	7.0	3.9	VS 2017 15.5
<code>insert_or_assign()/try_emplace()</code> for maps	6.1	3.7	VS 2017 15
Emplace Return Type	7.1	4.0	VS 2017 15.3
Sampling algorithms	7.1	In Progress	VS 2017 15
<code>gcd</code> and <code>lcm</code>	7.1	4.0	VS 2017 15.3
<code>clamp</code>	7.1	3.9	VS 2015.3
Special Mathematical Functions	7.1	Not yet	VS 2017 15.7
Shared Pointers and Arrays	7.1	In Progress	VS 2017 15.5
Non-member <code>size()</code> , <code>data()</code> and <code>empty()</code>	6.1	3.6	VS 2015
<code>constexpr</code> Additions to the Standard Library	7.1	4.0	VS 2017 15.3
<code>scoped_lock</code>	7.1	5.0	VS 2017 15.3
Polymorphic Allocator & Memory Resource	9.1	In Progress	VS 2017 15.6

Appendix B - Resources and References

First of all, if you want to dig into the standard on your own, you can read the latest draft here:

[N4687, 2017-07-30, Working Draft, Standard for Programming Language C++](#) - from [isocpp.org](#). It's the last draft before it got accepted, so it's almost as the final spec. Also have a look here for more information about the papers and the status: [isocpp : Standard C++](#).

For a quick overview of C++17 changes there's a paper: [P0636 - Changes between C++14 and C++17](#)

Books:

- [C++17 - The Complete Guide](#) by Nicolai Josuttis
- [C++17 STL Cookbook](#) by Jacek Galowicz
- [Modern C++ Programming Cookbook](#) by Marius Bancila
- [C++ Templates: The Complete Guide \(2nd Edition\)](#) by David Vandevorde, Nicolai M. Josuttis, Douglas Gregor
- [Professional C++, 4th Edition](#) by Marc Gregoire

General C++ Links:

- Compiler support: [C++ compiler support](#)
- The official paper with changes: [P0636r0: Changes between C++14 and C++17 DIS](#)
- Jason Turner: [C++ Weekly channel](#), where he covered most (or even all!) of C++17 features.
- [Simon Brand blog](#) - with lot's of information about C++17
- [Arne Mertz blog](#)
- [Rainer Grimm Blog](#)
- [CppCast](#)
- [FluentC++](#)

General C++17 Language Features

- Simon Brand: [Template argument deduction for class template constructors](#)
- [Class template deduction\(since C++17\) - cppreference](#).
- "Using fold expressions to simplify variadic function templates" in [Modern C++ Programming Cookbook](#).
- Simon Brand: [Exploding tuples with fold expressions](#)

- [Baptiste Wicht: C++17 Fold Expressions](#)
- [Fold Expressions - ModernesCpp.com](#)
- [Adding C++17 structured bindings support to your classes](#)
- [C++ Weekly Special Edition - Using C++17's constexpr if - YouTube](#) - real examples from Jason and his projects.
- [C++17: let's have a look at the constexpr if – FJ](#)
- [C++ 17 vs. C++ 14 — if-constexpr – LoopPerfect – Medium](#)
- [Two-phase name lookup support comes to MSVC](#)
- [What does the `carries_dependency` attribute mean? - Stack Overflow](#)
- [Value Categories in C++17 – Barry Revzin – Medium](#)
- [Rvalues redefined | Andrzej's C++ blog](#)

Expression Evaluation Order:

- [GotW #56: Exception-Safe Function Calls](#)
- [Core Guidelines: ES.43: Avoid expressions with undefined order of evaluation](#)
- [Core Guidelines: ES.44: Don't depend on order of evaluation of function arguments](#)

About `std::optional`:

- [Andrzej's C++ blog: Efficient optional values](#)
- [Andrzej's C++ blog: Ref-qualifiers](#)
- [Clearer interfaces with `optional<T>` - Fluent C++](#)
- [Optional - Performance considerations - Boost 1.67.0](#)
- [Enhanced Support for Value Semantics in C++17 - Michael Park, CppCon 2017](#)
- [std::optional: How, when, and why | Visual C++ Team Blog](#)

About `std::variant`:

- [SimplifyC++ - Overload: Build a Variant Visitor on the Fly.](#)
- [Variant Visitation by Michael Park](#)
- [Sum types and state machines in C++17](#)
- [Implementing State Machines with `std::variant`](#)
- [Pattern matching in C++17 with `std::variant`, `std::monostate` and `std::visit`](#)
- [Another polymorphism | Andrzej's C++ blog](#)
- [Inheritance vs `std::variant`, C++ Truths](#)

About `string_view`

- [CppCon 2015 `string_view` — Marshall Clow](#)
- [string_view odi et amo - Marco Arena](#)
- [C++17 `string_view` – Steve Lorimer](#)
- [Modernescpp - `string_view`](#)
- [foonathan::blog\(\) - `std::string_view` accepting temporaries: good idea or horrible pitfall?](#)
- [abseil / Tip of the Week #1: `string_view`](#)
- [std::string_view is a borrow type – Arthur O’Dwyer – Stuff mostly about C++](#)
- [C++ Russia 2018: Victor Ciura, Enough `string_view` to hang ourselves - YouTube](#)
- [StringViews, StringViews everywhere! - Marc Mutz - Meeting C++ 2017 - YouTube](#)
- [Jacek’s C++ Blog · Const References to Temporary Objects](#)
- [abseil / Tip of the Week #107: Reference Lifetime Extension](#)
- [C++17 - Avoid Copying with `std::string_view` - ModernesCpp.com](#)

String Conversions and Searchers

- [How to Convert a String to an int in C++ - Fluent C++](#)
- [How to *Efficiently* Convert a String to an int in C++ - Fluent C++](#)
- [How to encode char in 2-bits? - Stack Overflow](#)

About `filesystem`

- Chapter 7, “Working with Files and Streams” - of **Modern C++ Programming Cookbook**.
- examples like: Working with `filesystem` paths, Creating, copying, and deleting files and directories, Removing content from a file, Checking the properties of an existing file or directory, searching.
- Chapter 10 “`filesystem`” from “**C++17 STL Cookbook**”
- examples: path normalizer, Implementing a grep-like text search tool, Implementing an automatic file renamer, Implementing a disk usage counter, statistics about file types, Implementing a tool that reduces folder size by substituting duplicates with symlinks
- [C++17- `std::byte` and `std::filesystem` - ModernesCpp.com](#)
- [How similar are Boost `filesystem` and the standard C++ `filesystem` libraries? - Stack Overflow](#)

Parallel Algorithms

- Bryce Adelstein’s talk about parallel algorithms. Contains a lot of examples for map reduce (transform reduce) algorithm: [CppCon 2016: Bryce Adelstein Lelbach “The C++17 Parallel Algorithms Library and Beyond” - YouTube](#)
- And the Sean Parent talk about better concurrency in C++: [code::dive 2016 conference – Sean Parent – Better Code: Concurrency - YouTube](#)
- [Simon Brand - `std::accumulate` vs. `std::reduce`](#)
- [Using C++17 Parallel Algorithms for Better Performance | Visual C++ Team Blog](#)