

Mastering STL in C++23: New Features Updates, and Best Practices

C++23
STL

Prepared by: Ayman Alheraki

First Edition

Mastering STL in C++23: New Features, Updates, and Best Practices

Prepared By Ayman Alheraki
simplifycpp.org

January 2025

Contents

Contents	2
Author's Preface	7
1 Introduction to STL in C++	9
1.1 History of STL and Its Goals	9
1.1.1 Origins of STL	9
1.1.2 Adoption and Standardization	10
1.1.3 STL's Key Design Goals	11
1.1.4 Evolution of STL: From C++98 to C++23	13
1.2 Difference between STL and External Libraries like Boost	14
1.2.1 Core vs. Extended Functionality	14
1.2.2 Standardization: STL as the Language Standard vs. Boost as an External Library	16
1.2.3 Scope and Specialization	16
1.2.4 Performance Considerations	17
1.2.5 Community and Support	18
2 Recent Updates in C++23	21
2.1 Key Changes and Additions in STL within C++23	21

2.1.1	Ranges Library Enhancements	21
2.1.2	Enhanced <code>std::algorithm</code> with New Features	23
2.1.3	<code>std::format</code> for String Formatting	24
2.1.4	<code>std::ranges::for_each</code> for Ranges	25
2.1.5	Concepts in STL Algorithms	25
2.1.6	Performance Improvements	26
2.1.7	Deprecations and Removals	26
2.2	Explanation of Performance-Related Changes and Improvements	27
2.2.1	Optimizations in Standard Containers	28
2.2.2	Algorithmic Enhancements	29
2.2.3	Memory Management Enhancements	31
2.2.4	Multi-Threading and Concurrency Improvements	32
3	New and Modified Templates in STL	34
3.1	Newly Added Templates	34
3.1.1	<code>std::ranges::zip</code>	34
3.1.2	<code>std::ranges::cartesian_product</code>	36
3.1.3	<code>std::ranges::views::repeat</code>	38
3.1.4	<code>std::ranges::views::iota</code>	39
3.1.5	<code>std::ranges::to</code>	40
3.2	Modified or Improved Templates, such as <code>std::ranges</code> and <code>std::span</code> .	42
3.2.1	Enhancements to <code>std::ranges</code>	42
3.2.2	Improvements to <code>std::span</code>	45
3.2.3	Performance Improvements in <code>std::span</code>	47
3.2.4	Other Notable Template Improvements	48
4	New Concepts in C++23 for STL	49

4.1	<code>std::concepts</code> : How These Concepts Simplify Template Understanding and Ensure Compatibility	49
4.1.1	Overview of Concepts	50
4.1.2	How Concepts Simplify Template Understanding	50
4.1.3	Concepts and STL: Ensuring Compatibility	52
4.1.4	New Concepts in C++23	55
4.2	<code>std::ranges</code> : Expanding Template Capabilities and Enhancing Data Interactions	56
4.2.1	The Role of <code>std::ranges</code> in C++ Template Programming	57
4.2.2	Key Enhancements to <code>std::ranges</code> in C++23	57
4.2.3	Performance Improvements in <code>std::ranges</code>	61
4.2.4	Integration with Other C++23 Features	62
4.3	<code>std::span</code> : How It Can Be Used to Manage Arrays More Safely and Flexibly	63
4.3.1	Introduction to <code>std::span</code>	63
4.3.2	Benefits of Using <code>std::span</code> Over Raw Pointers	65
4.3.3	How <code>std::span</code> Works with Arrays, Vectors, and Other Containers	66
4.3.4	Slicing and Subranges with <code>std::span</code>	67
4.3.5	Improvements in C++23 for <code>std::span</code>	68
5	STL Containers in C++23	70
5.1	Modifications to Containers like <code>std::vector</code> and <code>std::list</code>	70
5.1.1	Overview of <code>std::vector</code> and <code>std::list</code>	71
5.1.2	Key Modifications in <code>std::vector</code>	71
5.1.3	Key Modifications in <code>std::list</code>	73
5.2	New or Modified Functions in Containers, like <code>insert_or_assign</code> and <code>contains</code>	75
5.2.1	The Introduction of <code>insert_or_assign</code>	76
5.2.2	The <code>contains</code> Function	78

5.2.3	Other Notable New Functions in C++23 Containers	80
6	New Libraries and Tools in STL	83
6.1	<code>std::expected</code> : Handling Errors	83
6.1.1	Introduction to <code>std::expected</code>	83
6.1.2	Structure and Usage of <code>std::expected</code>	85
6.1.3	Advantages of Using <code>std::expected</code>	88
6.1.4	Common Use Cases for <code>std::expected</code>	89
6.2	<code>std::sem</code> and <code>std::latch</code> : Advanced Concurrency Tools in C++23 . . .	90
6.2.1	Introduction to <code>std::sem</code> (Semaphore)	90
6.2.2	Introduction to <code>std::latch</code>	94
6.2.3	Comparing <code>std::sem</code> and <code>std::latch</code>	97
6.3	New Tools to Support Parallel Processing	98
6.3.1	Introduction to Parallelism in C++23	98
6.3.2	New Parallel Algorithms in C++23	98
6.3.3	Support for <code>std::execution::parallel_policy</code>	102
6.3.4	Improved Support for Parallel Sorting	103
6.3.5	Parallelism with <code>std::async</code> and <code>std::future</code>	104
7	Compatibility with Previous C++ Versions	106
7.1	How to Use Updates in Multi-Version Environments	106
7.1.1	The Importance of Backward Compatibility	107
7.1.2	Strategies for Using C++23 Updates in Multi-Version Environments . .	107
7.1.3	Keeping Dependencies Compatible with Multiple Versions	110
7.2	Enhancements in Backward Compatibility	112
7.2.1	The Importance of Backward Compatibility in C++	112
7.2.2	Enhancements in Backward Compatibility in C++23	113
7.2.3	Techniques for Maintaining Backward Compatibility in C++23	116

8	Best Practices for Using STL	119
8.1	Tips to Improve Program Performance Using STL	119
8.1.1	Understanding the Computational Complexity of STL Containers . . .	120
8.1.2	Leveraging STL Algorithms for Optimal Performance	122
8.1.3	Efficient Memory Usage with STL Containers	123
8.1.4	Optimizing for Parallelism in C++23	124
8.2	Avoiding Common Mistakes with Containers and Algorithms	126
8.2.1	Mistakes with Container Choice	127
8.2.2	Mistakes with Algorithm Usage	128
8.2.3	Mistakes with Memory Management and Efficiency	130
8.2.4	Mistakes with Iterators	132
9	Resources and Documentation	134
9.1	Links to Advanced Tutorials and Official STL Documentation	134
9.1.1	Official Documentation	134
9.1.2	Advanced Tutorials	136
9.1.3	Additional Resources	139
9.2	Books and Training Courses for Further Learning	140
9.2.1	Books for Learning STL and C++	141
9.2.2	Training Courses for C++ and STL	144

Author's Preface

The Standard Template Library (STL) has always been one of the most prominent components of the C++ language, serving as a cornerstone for building powerful and flexible applications. Since its inception, the STL has continuously evolved to reflect the changing needs of programmers and to keep up with modern programming challenges. With the updates introduced in the C++23 standard, the library has added numerous new features that open up broader possibilities for data management, performance optimization, and enhanced maintainability. This book, **"Mastering STL in C++23: New Features, Updates, and Best Practices"**, serves as your comprehensive guide to understanding and leveraging these new features. Whether you are an experienced programmer looking to enhance your efficiency with STL or a beginner seeking to grasp the fundamentals, this book provides you with the tools and practical explanations needed to develop modern and efficient applications.

In this book, we will cover:

1. **The updates and new features in C++23:** Such as enhanced ranges, improved support for dynamic memory, and other modern capabilities.
2. **Best practices for using STL:** How to write clean, efficient, and maintainable code using the Standard Template Library.
3. **Real-world case studies:** Practical examples demonstrating how STL can be employed to solve programming challenges effectively.

Additionally, the book delves into advanced tips for using STL in complex contexts such as concurrent programming, object-oriented programming, and high-performance programming. I write this book based on years of hands-on experience with C++, knowing that the power of STL lies not only in its ready-to-use functions but also in the philosophy behind its design. I hope this book serves as a guide to inspire you to explore more about the capabilities of C++ and use the STL to achieve your programming goals.

Stay Connected

For more discussions and valuable content about **C++**, I invite you to follow me on **LinkedIn**:

<https://linkedin.com/in/aymanalheraki>

You can also visit my personal website:

<https://simplifycpp.org>

I wish all **C++** enthusiasts continued success and progress on their journey with this remarkable and distinctive programming language.

Wishing you success in learning and growth,

Ayman Alheraki

Chapter 1

Introduction to STL in C++

1.1 History of STL and Its Goals

The **Standard Template Library (STL)** in C++ is a powerful set of template classes and functions that provide essential data structures and algorithms. Its primary goal is to simplify the development of efficient and reusable code. This section will delve into the historical context of STL, tracing its origins and development over time, and explaining its core goals that have shaped its evolution into an integral part of the C++ programming language.

1.1.1 Origins of STL

The history of STL can be traced back to the early 1990s, when C++ was beginning to gain traction as a widely used programming language for system and application software development. At that time, C++ was already a powerful language for object-oriented programming, but it lacked a comprehensive and efficient library for common data structures and algorithms.

The Role of Alexander Stepanov

STL's creation is most closely associated with Alexander Stepanov, a Soviet-born computer scientist who worked at Hewlett-Packard (HP). Stepanov recognized that many programming tasks could be simplified by using well-defined, efficient, and reusable abstractions. Inspired by mathematical structures, he set out to create a library that would enable developers to express common algorithms and data structures in a generic, reusable manner.

In the late 1980s and early 1990s, Stepanov, along with his colleague Meng Lee, started developing what would later become the STL. Initially, the library was designed as part of an internal project at HP to improve the usability and efficiency of their software tools. Stepanov's idea was to combine **generic programming**—a paradigm that emphasizes writing code that works with any data type—with the well-established principles of **object-oriented programming**.

1.1.2 Adoption and Standardization

In 1994, the STL was officially adopted by the C++ Standards Committee and was included in the first international C++ standard, **ISO/IEC 14882:1998**, known as **C++98**. The inclusion of STL in the C++ standard marked a key milestone in its development, as it made STL a standard library for all compliant C++ compilers. This standardization also paved the way for STL's widespread adoption in both academic and industry projects.

The STL's design was revolutionary for its time, as it introduced the concept of **generic programming** in C++. It utilized **templates** extensively, allowing developers to write algorithms and data structures that could be applied to any type of data without sacrificing performance or type safety. For example, instead of writing separate functions for sorting different types of containers (arrays, lists, etc.), STL provided a single generic algorithm (e.g., `std::sort`) that could work with any container, as long as the container's elements met certain requirements.

1.1.3 STL's Key Design Goals

When creating STL, Stepanov and his team had several key goals in mind:

1. **Generality and Reusability**

One of the main goals of STL was to create a library that could be reused across various applications and use cases. The core concept behind STL is **generic programming**, which allows developers to write algorithms and data structures that are not dependent on specific data types. This is achieved through **template classes and functions**, which enable flexibility while maintaining type safety.

For example, STL provides generic containers like `std::vector`, `std::list`, and `std::map`, which can store elements of any data type. Additionally, algorithms such as `std::sort`, `std::find`, and `std::accumulate` can be applied to these containers without modification, resulting in efficient, reusable code.

2. **Performance and Efficiency**

Performance is one of the defining characteristics of STL. From the outset, Stepanov designed STL with an emphasis on efficiency. The library was built to support both high-level abstraction and low-level control over the performance characteristics of data structures and algorithms.

STL containers, like `std::vector`, are designed to minimize memory overhead and to optimize for common use cases, such as fast access and insertion. Similarly, algorithms in STL, such as `std::sort`, are implemented to take advantage of the most efficient algorithms and techniques available for the task at hand. This dual focus on generality and performance was a key factor in STL's success.

3. **Interoperability and Modularity**

Another goal of STL was to ensure that its components could be easily combined with other parts of the C++ Standard Library and third-party libraries. STL provides a set of

iterators, which abstract the process of traversing containers. These iterators make it easy to plug STL algorithms into any C++ data structure, whether that data structure is part of STL or defined by the user.

The modular design of STL allows developers to use only the components they need for a particular application, without unnecessary overhead. For example, if a developer only requires a container to store elements, they may choose to use `std::vector` or `std::set`, depending on the specific requirements for performance and behavior. Similarly, algorithms such as `std::transform` and `std::foreach` can be used in combination with user-defined iterators or third-party data structures.

4. Abstracting Common Data Structures and Algorithms

Before STL, developers often had to implement their own versions of common data structures like stacks, queues, linked lists, and hash tables. STL's goal was to provide standardized, well-optimized, and generic implementations of these data structures. This not only reduced the need for redundant code but also allowed developers to focus on higher-level concerns.

In addition to data structures, STL also provides a wide range of algorithms for operations like searching, sorting, and modifying data. These algorithms are designed to work seamlessly with the containers, making it easier for developers to perform complex tasks with minimal code.

5. Type Safety and Compile-Time Checking

Type safety is another key aspect of STL's design. By using templates, STL ensures that the type of elements in a container is known at compile time, preventing many common runtime errors. For example, if an algorithm is passed a container of an incorrect type (such as trying to sort a container of strings with a comparison function designed for integers), the compiler will catch this error early, before the program is run.

This compile-time checking was a major breakthrough in C++ and has become a hallmark of modern C++ development. It is one of the reasons why STL remains an indispensable tool for writing efficient and reliable code.

1.1.4 Evolution of STL: From C++98 to C++23

Since its standardization, STL has evolved significantly. While C++98 provided the foundation for STL, subsequent versions of the C++ standard have introduced many improvements and new features to the library. The inclusion of features like **move semantics**, **smart pointers**, **lambda functions**, and **concurrency utilities** in later versions (C++11, C++14, C++17, and C++20) has further expanded STL's capabilities.

In particular, C++11 introduced major enhancements such as the `std::unique_ptr`, `std::shared_ptr`, and the `std::thread` class. C++14 and C++17 added improvements to performance, such as reduced overhead in `std::map` and `std::set`, and C++20 brought significant updates to ranges, concepts, and the algorithm library, further enhancing the expressiveness and efficiency of STL.

Now, with **C++23**, STL continues to evolve. Features like **calendar and timezone utilities**, **extended algorithms**, and **more powerful ranges** have expanded STL's already vast capabilities, providing modern C++ developers with even greater tools for writing clean, efficient, and reusable code.

Conclusion

The history of STL in C++ is a story of innovation and collaboration that has significantly impacted the development of the language. From its inception by Alexander Stepanov in the early 1990s to its standardization and continuous evolution through various versions of C++, STL has become a cornerstone of the C++ language. Its primary goals of reusability, performance, modularity, and type safety have shaped modern C++ programming practices and continue to drive its relevance in contemporary software development.

In the following sections, we will explore the individual components of STL—its containers, algorithms, iterators, and utilities—along with the latest features introduced in C++23, to help you understand how to master STL and leverage its power in your own projects.

1.2 Difference between STL and External Libraries like Boost

In this section, we will explore the differences between the **Standard Template Library (STL)**, which is part of the C++ Standard Library, and external libraries like **Boost**, which provides additional functionality and tools for C++ developers. While both STL and Boost offer powerful features that enhance the capabilities of C++, they differ in several key aspects such as their design goals, scope, flexibility, and community support. Understanding these differences will help you make informed decisions when choosing between STL and external libraries like Boost for your C++ projects.

1.2.1 Core vs. Extended Functionality

STL: Core Library of C++

The **Standard Template Library (STL)** is a core part of the **C++ Standard Library**. It is defined and standardized by the **International Organization for Standardization (ISO)** as part of the C++ standard (ISO/IEC 14882). STL provides essential components such as:

- **Containers** (e.g., `std::vector`, `std::list`, `std::map`, `std::unordered_map`)
- **Algorithms** (e.g., `std::sort`, `std::find`, `std::accumulate`)
- **Iterators** (which provide a standardized way to access elements of containers)
- **Functors** (objects that can be used as functions)

These components cover the fundamental needs of a C++ programmer for managing and manipulating collections of data efficiently and generically. STL focuses on providing high-performance, general-purpose data structures and algorithms that are integrated directly into the language.

Boost: Extended Functionality

On the other hand, **Boost** is an external, open-source library that extends the functionality of C++. Boost is not part of the C++ standard but is often considered one of the most comprehensive and well-regarded external libraries for C++. Boost provides a wide range of features that go beyond the scope of the STL, including:

- **Smart pointers** (e.g., `boost::shared_ptr`, `boost::scoped_ptr`)
- **Regular expressions** (`boost::regex`)
- **Graph algorithms** (`boost::graph`)
- **Multithreading utilities** (e.g., `boost::thread`, `boost::mutex`)
- **Type traits** and **metaprogramming** tools
- **File system and path manipulation** (`boost::filesystem`)

Many of the libraries in Boost focus on areas that are either not covered or only partially covered by the C++ standard. For example, while the STL provides containers and algorithms, Boost offers advanced tools for **generic programming**, **multithreading**, and **networking**. Additionally, some of the features in Boost are designed to work seamlessly with STL containers and algorithms, providing more specialized functionality.

1.2.2 Standardization: STL as the Language Standard vs. Boost as an External Library

STL: Part of the C++ Standard

The **STL** is an essential part of the C++ language itself. As a standardized component of the C++ language, STL is guaranteed to be available in any compliant C++ compiler and is automatically included in any C++ development environment. Its specifications are defined by the **C++ Standard**, meaning that all developers working with C++ have access to the same version of STL, ensuring compatibility and portability across different compilers and platforms. Since STL is a **core part of the language**, it undergoes rigorous scrutiny and updates during each version of the C++ standard. For example, significant changes to STL occurred in C++11, C++14, C++17, and most recently in C++20 and C++23, such as the addition of **ranges**, **concepts**, and **new algorithms**.

Boost: Not Part of the C++ Standard

In contrast, **Boost** is not officially part of the C++ standard, although many of its components have influenced the C++ standard library. Boost libraries are external, meaning that they are developed and maintained independently of the C++ standardization process. As a result, developers must manually download and integrate Boost into their projects, either by linking to the Boost binaries or by compiling the source code themselves.

Because Boost is an external library, it provides greater flexibility for innovation, and it can be updated and enhanced more frequently than STL. However, this also means that Boost libraries are not guaranteed to be available in every C++ development environment, and their adoption can sometimes require more configuration than STL.

1.2.3 Scope and Specialization

STL: Focused on General-Purpose Algorithms and Data Structures

STL is designed with a primary focus on **general-purpose algorithms** and **data structures**. Its goal is to provide highly optimized, reusable, and efficient containers (e.g., `std::vector`, `std::list`) and algorithms (e.g., `std::sort`, `std::find`) that can be applied across a wide range of use cases. STL's functionality is broad enough to cover most everyday programming tasks, but it does not extend into more specialized or advanced domains.

The containers and algorithms provided by STL are highly optimized for common operations, and its design emphasizes generic programming principles, ensuring that these components work with any type, provided the type meets certain requirements (e.g., being sortable for the `std::sort` algorithm).

Boost: Covers Specialized and Advanced Domains

Boost, on the other hand, is much broader in scope and often targets **specialized domains** or **advanced use cases** that may not be addressed by the C++ standard library. While STL focuses on the basics, Boost offers tools for areas like:

- **Meta-programming** (e.g., `boost::mpl`, `boost::hana`)
- **Memory management** (e.g., custom allocators and `boost::pool`)
- **Networking** (e.g., `boost::asio`)
- **Mathematics and statistics** (e.g., `boost::math`)

For example, Boost's **Boost.Asio** library allows for asynchronous I/O operations, which is far more specialized than anything available in STL. Similarly, **Boost.MPL** provides tools for metaprogramming, which allow developers to write code that manipulates types and values at compile time—something STL does not directly provide.

1.2.4 Performance Considerations

STL: Optimized for Common Use Cases

STL containers and algorithms are designed to be **highly efficient** for common programming tasks. They are tailored for a variety of standard use cases and are often highly optimized for performance. For example:

- `std::vector` is highly optimized for fast access and appending operations.
- `std::map` ensures that elements are kept in sorted order with logarithmic-time insertions and lookups.

Because STL is part of the core language, these components are designed with the assumption that they will be used frequently, so performance is a critical consideration. For most general programming tasks, STL provides optimal performance.

Boost: Specialized Performance Considerations

While **Boost** libraries are also designed with efficiency in mind, their primary goal is often to provide **extended or specialized functionality** rather than covering the most common use cases. Some Boost libraries may offer a higher level of abstraction or implement more complex algorithms, which can result in a slight performance overhead compared to STL's highly optimized core features.

However, many Boost libraries are still highly efficient, and some of them (like `boost::asio` or `boost::filesystem`) provide performance that is on par with or even superior to alternative solutions in the STL. Boost's ability to extend C++ with new features while maintaining strong performance is one of its greatest strengths.

1.2.5 Community and Support

STL: Well-Defined and Widely Supported

Because STL is part of the C++ standard, it has wide support across all **C++ compilers** and **development environments**. It is **well-documented**, and there is a wealth of educational

material, books, and tutorials available for developers to learn and master it. Given that STL is standardized, it is guaranteed to work across different platforms and compiler implementations, making it a safe and reliable choice for most C++ applications.

Boost: Open-Source and Actively Maintained

Boost is an **open-source** project that is developed and maintained by a community of developers. As a result, it is highly flexible and rapidly evolving. Boost has a large user base, and many of its libraries are widely adopted by the C++ community. Boost has become an essential tool for C++ developers who require advanced functionality that is not available in STL.

Boost also has a strong **community-driven support system**, with forums, mailing lists, and documentation available to assist developers. Many C++ developers rely on Boost for cutting-edge functionality, and several Boost libraries have been proposed for inclusion in future C++ standards (e.g., `std::filesystem`, which originated in Boost).

Conclusion

In summary, the **STL** and **Boost** serve different, though complementary, roles in the C++ ecosystem:

- **STL** is the core, standardized library that provides essential containers, algorithms, and iterators for everyday programming tasks, with an emphasis on efficiency, type safety, and portability.
- **Boost** extends the functionality of C++ with advanced libraries that target specialized domains, offering tools for metaprogramming, multithreading, networking, and more.

Choosing between STL and Boost depends on the specific needs of your project. If your requirements are met by the general-purpose features in STL, it is often the best choice due to its integration with the C++ standard and its broad support. However, if your project demands advanced features or specialized libraries, Boost is an excellent choice, offering a wealth of functionality that can complement the STL and extend C++ programming to new heights.

In the next sections, we will dive deeper into the specific components of STL and Boost, and explore how you can effectively leverage both libraries to write modern, efficient, and maintainable C++ code.

Chapter 2

Recent Updates in C++23

2.1 Key Changes and Additions in STL within C++23

The **C++23** standard, the latest version of the C++ programming language, introduces several exciting updates and additions to the **Standard Template Library (STL)**. As the C++ language evolves, the STL continues to improve by incorporating new features and refining existing ones. C++23 builds upon the work done in previous standards (C++17 and C++20) to bring more powerful, efficient, and user-friendly features to the language.

In this section, we will examine the **key changes** and **new additions** to STL in C++23, focusing on how they improve performance, usability, and versatility. These changes enhance the overall C++ programming experience, offering developers more tools for writing high-quality, modern C++ code.

2.1.1 Ranges Library Enhancements

The **ranges** library, which was introduced in C++20, saw further improvements in C++23. The concept of ranges offers a more intuitive and flexible way of working with sequences of

elements in C++. Ranges allow algorithms to operate on any container or sequence type without requiring manual iteration or the use of iterators explicitly. C++23 extends and improves the ranges library to provide even greater expressiveness and convenience.

Range Adaptors: New Additions

C++23 introduces a set of new **range adaptors**, which enhance the already existing range-based algorithms by enabling more powerful manipulations of ranges.

- **`views::transform`** now supports **in-place transformations**, where you can modify the elements of the range in place without needing to create a new container. This helps optimize memory usage and provides an efficient way to work with large datasets.
- **`views::reverse`** now also works with **non-random-access containers** (e.g., `std::list` and `std::deque`). This addition makes it more flexible, allowing the reversal of any container, even those that do not support direct element access.
- **`views::join`** has been improved to allow working with **non-ranges** in addition to ranges. This is particularly useful when working with containers of containers, as it simplifies flattening multi-level containers into a single range.

These adaptors add versatility to the **ranges** library, enabling developers to work with sequences of data more fluently and expressively.

New Range-Based Algorithms

C++23 introduces new range-based algorithms that offer more flexibility and allow operations to be performed more naturally on ranges.

- **`ranges::zip`**: The new `ranges::zip` algorithm allows you to combine multiple ranges into a single range of tuples. This is useful for scenarios where you need to iterate over multiple sequences simultaneously, such as zipping two containers together for parallel processing or transformation.

- **`ranges::clamp`**: While `std::clamp` exists in C++17, C++23 introduces a new range-based version of it. `ranges::clamp` makes it easier to clamp the elements of a range within a specified range, improving readability and conciseness when performing such operations.

These algorithms help reduce boilerplate code and enhance the expressiveness of range-based programming.

2.1.2 Enhanced `std::algorithm` with New Features

C++23 continues to enrich the set of algorithms available in the `std::algorithm` namespace, allowing developers to perform complex operations with ease.

`std::sort` with Customizable Comparison

In C++23, the `std::sort` algorithm receives an important feature that allows the use of a **customizable comparison function** directly within the algorithm. The previous versions of C++ required the use of function pointers or lambda expressions outside the algorithm call, which could sometimes be cumbersome or reduce readability.

With this new feature, developers can now pass a **custom comparator** directly into `std::sort`, simplifying syntax and improving code clarity.

```
std::sort(begin, end, std::ranges::greater<int>());
```

This improvement allows for more concise, readable, and flexible sorting logic.

`std::ranges::replace` and `std::ranges::replace_if`

C++23 introduces `std::ranges::replace` and `std::ranges::replace_if` as range-based versions of the existing `std::replace` and `std::replace_if` algorithms. These range-based algorithms provide a more intuitive and declarative syntax for performing replacements on containers and ranges.

With `std::ranges::replace`, you can replace all occurrences of a value in a range with a new value, without needing to manually specify iterators:

```
std::ranges::replace(vec, old_value, new_value);
```

Similarly, `std::ranges::replace_if` allows replacing elements based on a condition, streamlining operations that involve conditional replacements.

2.1.3 `std::format` for String Formatting

C++20 introduced **`std::format`**, which provides a modern, type-safe alternative to older string formatting mechanisms like `sprintf`. In C++23, **`std::format`** is enhanced to integrate seamlessly with STL.

Formatting Support for Containers

One of the most anticipated updates in C++23 is the ability to format **containers** and **ranges** directly using `std::format`. This new functionality allows developers to easily convert entire containers into formatted strings without manually iterating over them or creating custom serialization logic.

For example:

```
std::vector<int> vec = {1, 2, 3, 4, 5};  
std::string formatted = std::format("Container: {}", vec);
```

This enhancement significantly reduces boilerplate code and provides a clean way to format and display the contents of containers and ranges.

`std::views::cartesian_product`

Another exciting addition in C++23 is the **`std::views::cartesian_product`**, which enables you to compute the **Cartesian product** of two or more ranges. This operation allows

you to efficiently combine elements from multiple containers into a single range of all possible pairs (or tuples, in the case of more than two ranges).

```
auto product = std::views::cartesian_product(vec1, vec2);
for (auto [a, b] : product) {
    std::cout << a << ", " << b << std::endl;
}
```

The addition of Cartesian product simplifies many tasks that would otherwise require manual iteration over multiple containers, making it easier to work with combinations of data in a clean and efficient way.

2.1.4 `std::ranges::for_each` for Ranges

C++23 introduces a new version of `std::ranges::for_each`, which works similarly to the traditional `std::for_each`, but operates directly on ranges rather than requiring iterators. This makes it more natural to use in the context of modern C++ programs that rely heavily on ranges.

```
std::ranges::for_each(vec, [] (int& x) { x *= 2; });
```

This improvement simplifies code and improves readability when working with ranges.

2.1.5 Concepts in STL Algorithms

In C++20, **concepts** were introduced as a way to impose constraints on template parameters. In C++23, **concepts** are now used more extensively in STL algorithms to ensure that only valid types are passed to them, improving type safety and error detection at compile time.

For example, algorithms like `std::ranges::sort` now include **concepts** that restrict the algorithm's use to only those types that support the necessary operations (e.g., the ability to

compare or swap elements). This makes STL algorithms more robust and helps avoid misuse of templates with incompatible types.

2.1.6 Performance Improvements

Along with the new features and additions, C++23 introduces a number of **performance improvements** to STL. These improvements focus on reducing overhead, improving memory efficiency, and enhancing the performance of existing algorithms and data structures.

Some of these improvements include:

- **Optimized memory allocations** for certain container types.
- **Improved performance of `std::map` and `std::set`** by refining their internal balancing algorithms.
- **Better support for parallel and concurrent algorithms**, with more efficient implementations of parallel algorithms introduced in C++17.

2.1.7 Deprecations and Removals

C++23 also marks the deprecation or removal of certain features from STL that were either outdated or replaced by more efficient or standardized alternatives.

For example:

- **`std::pair` deprecation for some special cases:** C++23 deprecates the use of `std::pair` in specific contexts where a more appropriate and type-safe alternative exists (e.g., `std::tuple` in some cases).
- **Removal of deprecated functions and algorithms:** Functions and algorithms that were marked for deprecation in previous versions of C++ are now officially removed in C++23.

These changes ensure that STL stays modern and aligns with best practices, while also maintaining backward compatibility where possible.

Conclusion

C++23 brings significant improvements to the **Standard Template Library (STL)**, enhancing its versatility, performance, and ease of use. From new range-based algorithms and adaptors to improved string formatting support and Cartesian product functionality, these updates provide more powerful tools for C++ developers. Additionally, the expanded use of **concepts** and further optimization of existing algorithms ensure that C++23 continues to advance as a modern, high-performance language.

As C++ evolves, the STL adapts to meet the growing demands of developers, making it easier to write efficient, maintainable, and expressive code. In the next sections, we will explore in more detail how these new features can be applied in real-world applications, along with best practices for utilizing them effectively.

2.2 Explanation of Performance-Related Changes and Improvements

The **C++23** standard introduces several important **performance-related changes** and **improvements** to the Standard Template Library (STL) and the C++ language as a whole. These changes are designed to make C++ more efficient, allowing developers to write code that runs faster while using fewer resources. This section will explore the key performance enhancements that are part of C++23, focusing on the improvements to STL components, algorithms, containers, memory management, and multi-threading support.

2.2.1 Optimizations in Standard Containers

C++23 continues to optimize **standard containers** to ensure that they provide **better performance** in terms of both time complexity and memory usage. These optimizations focus on reducing overhead and improving access times, making STL containers even more efficient for common use cases.

std::vector Performance Improvements

The **std::vector** container is one of the most widely used containers in C++ due to its dynamic array-like behavior. C++23 introduces **several performance optimizations** for **std::vector**, particularly related to **memory allocations** and **element insertion**.

- **Reduced Allocation Overhead:** In C++23, **std::vector** has been optimized to reduce the overhead incurred during **dynamic memory allocations**, particularly when resizing the vector. These improvements allow **std::vector** to handle growing sizes more efficiently, without unnecessary reallocation or copying of data.
- **Improved Insertion Efficiency:** When elements are inserted into a **std::vector**, the container may need to reallocate its underlying memory to accommodate new elements. C++23 enhances the way the vector grows by improving the **allocation strategy** and reducing the number of reallocation steps. This means that vector insertion operations will now be faster and more memory-efficient, especially when handling larger datasets.
- **Optimized Move Semantics:** C++23 improves move semantics for **std::vector**, reducing the cost of transferring ownership of elements from one vector to another. This is particularly beneficial when using **std::move** to avoid unnecessary copies, leading to performance gains in scenarios involving large or complex objects.

These optimizations contribute to **faster insertion**, **better memory utilization**, and **reduced copy costs** when working with vectors, which are commonly used in performance-critical applications.

`std::map` and `std::set` Optimizations

The associative containers `std::map` and `std::set` have also received significant performance improvements in C++23. Both containers use **red-black trees** internally to maintain order and support logarithmic time complexity for insertion, deletion, and lookup operations.

- **Improved Balancing Algorithms:** The **balancing algorithms** used by `std::map` and `std::set` have been optimized in C++23 to improve performance during tree balancing operations. This leads to better performance when handling a large number of insertions or deletions, reducing the time taken to maintain the order of elements in the tree.
- **Reduced Memory Usage:** C++23 introduces improvements to the **memory footprint** of `std::map` and `std::set` by optimizing the underlying memory structures. This allows for better utilization of memory and more efficient allocation, especially when the containers store a large number of elements.

These optimizations help ensure that `std::map` and `std::set` remain efficient, even in scenarios where the size of the data set grows considerably.

2.2.2 Algorithmic Enhancements

C++23 also introduces **algorithmic optimizations** to improve the performance of standard algorithms used in STL. These optimizations aim to make algorithmic operations faster and more memory-efficient, improving the overall performance of C++ programs.

Parallel and Vectorized Algorithms

C++17 introduced the ability to execute certain algorithms in parallel through the **parallel algorithms** library (`std::for_each`, `std::sort`, `std::transform`, etc.), allowing for **multi-core optimization**. C++23 further enhances these parallel algorithms, improving their **performance and scalability**.

- **Improved Parallelism:** C++23 introduces new parallel execution policies and optimizations that allow algorithms to better utilize multi-core processors, leading to better **multi-threading performance** in computationally intensive tasks. Algorithms like `std::sort` and `std::transform` can now more effectively take advantage of hardware parallelism, leading to faster execution, particularly on multi-core systems.
- **Vectorization:** The C++23 standard includes enhancements that help algorithms better leverage **SIMD (Single Instruction, Multiple Data)** instructions and **vectorization** capabilities available on modern CPUs. By utilizing vectorized instructions, certain operations can process multiple data points simultaneously, speeding up operations like sorting, searching, and transforming data.

These improvements make it easier for developers to harness the power of modern hardware, allowing algorithms to run faster by executing on multiple cores or utilizing SIMD instructions for data parallelism.

New and Optimized Algorithms

C++23 introduces several new algorithms as well as enhancements to existing algorithms that improve performance and flexibility:

- `std::ranges::transform` and `std::ranges::copy` now come with **optimized versions** that improve their efficiency when working with large data sets. The optimizations reduce overhead and improve speed when copying or transforming ranges of elements.
- `std::ranges::find` and other searching algorithms now offer **faster performance** by optimizing the internal search mechanisms. This is especially beneficial when working with large containers, as it reduces the time required to find elements within the range.
- `std::ranges::sort` and `std::ranges::stable_sort` now include **improved implementations** that make sorting ranges more efficient. C++23 also

introduces a **parallel version** of the `ranges::sort` algorithm that provides significant speedups when sorting large ranges of data.

By improving the performance of these core algorithms, C++23 enables developers to process data more efficiently in critical sections of their code, leading to overall application performance improvements.

2.2.3 Memory Management Enhancements

Efficient **memory management** plays a vital role in optimizing the performance of C++ programs. C++23 introduces various improvements to memory allocation and deallocation strategies, which lead to better performance in programs that require frequent allocation and deallocation of memory.

Memory Resource Management

C++23 enhances the **memory resource** management capabilities of the standard library by adding features that allow more fine-grained control over memory allocations.

- **Custom Allocators:** C++23 refines the support for **custom allocators**, allowing developers to implement more efficient memory allocation strategies for containers. Custom allocators can be used to reduce the overhead of memory allocations, particularly in performance-critical applications where high throughput and low latency are required.
- **Improved `std::pmr` (Polymorphic Memory Resource):** The **polymorphic memory resource** (`std::pmr`) introduced in C++17 is further optimized in C++23. This feature allows developers to define and manage custom memory pools for specific containers, improving memory locality and cache performance. The new enhancements provide greater flexibility and efficiency in managing memory in performance-critical scenarios.

Cache-Friendly Allocations

C++23 introduces **cache-friendly memory allocation** strategies for STL containers, which focus on minimizing cache misses and improving cache locality. These strategies are particularly important in high-performance applications that process large amounts of data in a short period, such as gaming engines or scientific computing applications.

By improving how memory is allocated and accessed, C++23 reduces the time spent accessing data in memory, leading to faster execution times, especially for memory-intensive tasks.

2.2.4 Multi-Threading and Concurrency Improvements

With multi-core processors becoming the norm, **multi-threading** and **concurrency** support in C++ have been a major area of focus in recent standards. C++23 introduces further improvements to help C++ developers harness the power of modern processors in a more efficient and easy-to-use manner.

Parallel Algorithms and Thread Pools

As mentioned earlier, C++23 builds on the parallel algorithms introduced in C++17 by introducing **thread pools** and new parallel execution policies. These improvements allow algorithms to execute in parallel with more control over the number of threads being used, providing developers with more efficient ways to distribute workloads across multiple cores.

- **Thread Pools:** A thread pool allows for reusing threads instead of constantly creating and destroying them. This reduces the overhead of thread creation and management, leading to improved performance in multi-threaded applications.
- **Refined Parallel Execution Policies:** The parallel execution policies (`std::execution::par` and `std::execution::par_unseq`) have been enhanced to provide better control over the execution of parallel algorithms. This helps improve the efficiency and scalability of parallel algorithms, especially when running on systems with many CPU cores.

Synchronization Primitives

C++23 introduces several **synchronization primitives** that help improve performance when dealing with multi-threaded environments:

- **`std::atomic` optimizations:** New optimizations to `std::atomic` operations in C++23 reduce the latency of atomic operations, leading to more efficient multi-threading. These improvements are especially useful for high-concurrency applications where atomic operations are frequently used.
- **Improved mutexes:** New types of **mutexes** and locks are introduced in C++23, providing more efficient and lower-latency synchronization mechanisms.

Conclusion

The **performance-related improvements** introduced in **C++23** significantly enhance the efficiency, scalability, and memory utilization of C++ programs. Key updates include:

- **Optimizations for standard containers** like `std::vector` and `std::map`, which make memory allocation and element insertion more efficient.
- **Enhancements to standard algorithms** with better parallelism, vectorization, and execution policies, leading to faster execution on modern hardware.
- **Memory management improvements** that provide better control over memory allocation, leading to reduced overhead and improved cache performance.
- **Concurrency and multi-threading advancements** that allow more efficient parallel execution and synchronization, improving performance in multi-threaded applications.

These improvements ensure that C++23 continues to meet the growing demands for performance in modern software development. By taking advantage of these enhancements, developers can write faster, more efficient code that makes better use of system resources.

Chapter 3

New and Modified Templates in STL

3.1 Newly Added Templates

C++23 introduces several new templates to the **Standard Template Library (STL)**, significantly expanding its capabilities. These new templates enhance the expressiveness of the STL, making it easier for developers to write high-quality, efficient, and type-safe code. In this section, we will focus on the **newly added templates** in C++23, discussing their purposes, usage, and how they fit into the broader STL ecosystem.

3.1.1 `std::ranges::zip`

Introduced in C++23, `std::ranges::zip` is a new template function designed to combine multiple ranges into a single range of tuples. This template provides an elegant way to work with multiple sequences of data simultaneously without needing to manually manage iterators or indices. It is particularly useful when you need to process multiple containers in parallel or perform operations involving multiple ranges.

Usage

The `std::ranges::zip` template allows you to "zip" two or more ranges together, creating a range of tuples. Each tuple contains one element from each of the input ranges at the same position.

For example, if you have two containers, you can use `zip` to create a single range of pairs:

```
#include <iostream>
#include <ranges>
#include <vector>

int main() {
    std::vector<int> vec1 = {1, 2, 3};
    std::vector<char> vec2 = {'a', 'b', 'c'};

    for (auto [i, c] : std::ranges::zip(vec1, vec2)) {
        std::cout << i << " " << c << "\n";
    }

    return 0;
}
```

This code will output:

```
1 a
2 b
3 c
```

The `std::ranges::zip` template provides an easy and efficient way to iterate over multiple containers concurrently, without manually managing iterators or indices.

Key Features

- **Tuple-like structure:** The result of `zip` is a range of tuples, making it easy to access elements from each range together.

- **No need for manual indexing:** With `zip`, you can avoid manually managing the indices of multiple containers, simplifying your code and reducing the risk of errors.
- **Works with ranges of different types:** `zip` can handle ranges of different types and sizes, allowing for more flexible combinations of sequences.

This template is especially useful in cases like parallel iteration, transforming multiple containers together, or combining data from different sources.

3.1.2 `std::ranges::cartesian_product`

Another exciting addition to C++23 is `std::ranges::cartesian_product`, a template that calculates the **Cartesian product** of multiple ranges. The Cartesian product of two or more sets is the set of all possible pairs (or tuples in the case of more than two sets) that can be formed by combining one element from each set.

Usage

The `std::ranges::cartesian_product` template generates a range of tuples representing the Cartesian product of the input ranges. This template makes it easy to iterate over all possible combinations of elements from multiple ranges.

For example, if you want to combine two vectors into all possible pairs, you can use `cartesian_product` as follows:

```
#include <iostream>
#include <ranges>
#include <vector>

int main() {
    std::vector<int> vec1 = {1, 2};
    std::vector<char> vec2 = {'a', 'b'};
```

```
for (auto [i, c] : std::ranges::cartesian_product(vec1, vec2)) {  
    std::cout << i << ", " << c << "\n";  
}  
  
return 0;  
}
```

This code will output:

```
1, a  
1, b  
2, a  
2, b
```

Key Features

- **Generates all combinations:** The Cartesian product template efficiently generates all possible combinations of elements from the input ranges.
- **Works with any number of ranges:** You can calculate the Cartesian product of any number of input ranges, not just two.
- **Tuple-based results:** The result of the Cartesian product is a range of tuples, allowing easy access to each element in the combination.

The `std::ranges::cartesian_product` template is useful for cases like generating combinations or permutations of data, solving optimization problems, or working with multidimensional data structures.

3.1.3 `std::ranges::views::repeat`

`std::ranges::views::repeat` is a new range adaptor introduced in C++23, enabling the creation of an infinite sequence of repeated values. This template is particularly helpful when you need to generate repetitive sequences or perform operations where an infinite repetition of a value is required.

Usage

The `repeat` adaptor generates an infinite range that repeats the specified value, and it can be used in combination with other range algorithms and adaptors to create interesting and flexible sequences.

Example usage:

```
#include <iostream>
#include <ranges>

int main() {
    auto repeated_range = std::ranges::views::repeat(42) |
        ↪ std::ranges::view::take(5);

    for (auto val : repeated_range) {
        std::cout << val << " ";
    }

    return 0;
}
```

This code will output:

```
42 42 42 42 42
```

Key Features

- **Infinite sequence:** The primary purpose of `std::ranges::views::repeat` is to generate an infinite sequence of repeated values. You can combine this with other range adaptors like `take` to limit the number of elements or `transform` to apply a function.
- **Highly flexible:** You can use `repeat` for various applications, such as filling containers, generating default values, or simulating repeated patterns.

This new template enhances the expressiveness of range-based programming by providing an easy-to-use tool for creating repetitive sequences.

3.1.4 `std::ranges::views::iota`

The `std::ranges::views::iota` template is a new range adaptor introduced in C++23, which generates a sequence of increasing values starting from a given number. This is similar to `std::iota`, but it is specifically designed to work with ranges.

Usage

You can use `std::ranges::views::iota` to generate a sequence of numbers that increment from a given starting value. By default, it increments by 1, but you can specify any step value.

Example usage:

```
#include <iostream>
#include <ranges>

int main() {
    auto range = std::ranges::views::iota(1, 6); // Generates numbers
    ↪ from 1 to 5
```

```
for (auto num : range) {  
    std::cout << num << " ";  
}  
  
return 0;  
}
```

This code will output:

```
1 2 3 4 5
```

You can also specify a custom step value:

```
auto range_with_step = std::ranges::views::iota(1, 10, 2); // Generates 1,  
↪ 3, 5, 7, 9
```

Key Features

- **Customizable ranges:** `std::ranges::views::iota` allows you to define the start, end, and step size for the sequence. This flexibility makes it useful for generating various types of ranges based on numerical patterns.
- **Range-based iteration:** This template works seamlessly with the range-based iteration features in C++20 and later, making it a natural fit for modern C++ codebases.

`std::ranges::views::iota` simplifies the process of generating sequences, particularly when working with algorithms that require arithmetic progressions.

3.1.5 `std::ranges::to`

C++23 introduces the `std::ranges::to` template, which allows converting a range directly to a container. This is particularly useful when you want to collect the elements of a range into a

specific container type (e.g., `std::vector`, `std::list`, etc.) without manually iterating over the range.

Usage

The `std::ranges::to` template simplifies the process of converting ranges to containers.

```
#include <iostream>
#include <ranges>
#include <vector>

int main() {
    std::vector<int> vec = {1, 2, 3, 4, 5};

    auto result = vec | std::ranges::views::transform([](int x) { return x
        ↪ * x; })
                  | std::ranges::to<std::vector>();

    for (auto val : result) {
        std::cout << val << " ";
    }

    return 0;
}
```

This code will output:

```
1 4 9 16 25
```

Key Features

- **Direct range-to-container conversion:** `std::ranges::to` makes it easier to convert a range to a container type in one line, reducing boilerplate code.

- **Works with any container type:** You can convert a range to any container type that supports the range's element type.

This addition simplifies code that involves transforming ranges and collecting the results in containers, enhancing both readability and efficiency.

Conclusion

The newly added templates in C++23 significantly enhance the flexibility, expressiveness, and convenience of working with ranges and containers in C++. Key additions like `std::ranges::zip`, `std::ranges::cartesian_product`, `std::ranges::views::repeat`, and others provide powerful tools for modern C++ developers, making it easier to work with complex data structures and sequences. By incorporating these new templates, developers can write more efficient, cleaner, and more maintainable code, leveraging the full potential of the STL in C++23.

3.2 Modified or Improved Templates, such as `std::ranges` and `std::span`

While C++23 introduces new templates to the Standard Template Library (STL), it also enhances and modifies existing templates to improve usability, performance, and functionality. In this section, we will focus on some of the most significant **modified or improved templates** in C++23, including `std::ranges` and `std::span`. These improvements are aimed at streamlining the way developers interact with ranges, views, and data structures, enhancing the flexibility and expressiveness of C++ code.

3.2.1 Enhancements to `std::ranges`

Introduced in C++20, the `std::ranges` library brought the power of **range-based programming** to C++, significantly improving the way developers could interact with

sequences of data. C++23 builds on the foundation laid by `std::ranges` by adding new functionalities and optimizations to the library. These modifications make it even easier to manipulate, filter, and transform sequences in a more expressive and efficient manner.

Expanded Support for Range Adaptors

One of the key enhancements to `std::ranges` in C++23 is the introduction of new **range adaptors** and the optimization of existing ones. Range adaptors allow developers to transform, filter, and combine ranges without the need for explicit loops or temporary containers.

Some of the notable new or improved range adaptors in C++23 include:

- **`std::ranges::view::transform`**: While the `transform` adaptor was already available in C++20, C++23 introduces optimizations that improve its performance when applied to large ranges. The `transform` adaptor now makes better use of **move semantics** and **parallel execution**, making it more efficient in certain scenarios.
- **`std::ranges::view::filter`**: The **`filter`** adaptor has received updates to allow **better memory management** and **faster execution** when filtering large ranges. In particular, filtering operations that involve complex conditions or large datasets benefit from the optimizations introduced in C++23.
- **`std::ranges::view::take` and `std::ranges::view::drop`**: These adaptors, which allow developers to **take** or **drop** elements from the start of a range, have been enhanced to support more complex slicing operations. They now allow for **non-zero starting indices** and can efficiently work with ranges that are not necessarily contiguous in memory.

These adaptors are part of the broader effort to provide more intuitive, efficient, and powerful tools for working with ranges in modern C++.

`std::ranges::to` Enhancements

C++23 introduces a significant improvement to the `std::ranges::to` adaptor, which allows for easy conversion from a range to a container. This improvement provides greater flexibility and control over the types of containers that can be created from ranges, enhancing the developer experience.

- **Custom Containers:** The `std::ranges::to` adaptor now supports **custom container types** more easily. Developers can specify their own container types (such as `std::deque`, `std::set`, or even user-defined containers) and use `to` to transform ranges into these containers with minimal boilerplate.
- **Improved Performance:** C++23 optimizes the implementation of `std::ranges::to` to reduce overhead and improve performance, especially when dealing with large data sets. This allows for faster conversions between ranges and containers, improving overall efficiency in C++ applications.

These improvements make `std::ranges::to` more versatile, allowing developers to seamlessly work with a broader range of container types and reducing unnecessary overhead when converting ranges.

New Algorithms in `std::ranges`

C++23 introduces **new algorithms** in the `std::ranges` library, extending its utility and performance. Some of these new algorithms help improve the flexibility and efficiency of working with ranges:

- **`std::ranges::is_sorted`:** This algorithm checks whether a range is sorted. In C++23, this algorithm is optimized for **short-circuiting**; that is, it will stop as soon as it detects that the range is not sorted. This optimization improves the algorithm's efficiency in cases where the range is already sorted early on.

- **`std::ranges::starts_with`**: This algorithm checks if a range starts with a specific subsequence. This new addition makes it easier to handle cases where you need to check if a sequence begins with a given prefix, such as string or array processing.

These new algorithms further enrich the `std::ranges` library, making it easier to perform complex range-based operations while improving performance.

3.2.2 Improvements to `std::span`

Introduced in C++20, **`std::span`** is a **non-owning view** over a sequence of elements, providing an efficient way to work with contiguous data without the overhead of copying or managing memory. C++23 enhances `std::span` with several modifications, making it even more powerful and flexible.

Support for Dynamic Extent `std::span`

In C++20, **`std::span`** required a **fixed size** for the number of elements it points to. This limitation has been relaxed in C++23 with the addition of **dynamic extent support**. A **dynamic extent span** allows the creation of spans where the size is **not known at compile time**, making it easier to work with runtime data sizes.

```
#include <iostream>
#include <span>

void print_span(std::span<int> sp) {
    for (auto val : sp) {
        std::cout << val << " ";
    }
    std::cout << "\n";
}

int main() {
```

```
int arr[] = {1, 2, 3, 4, 5};
std::span<int> s(arr); // span with dynamic extent

print_span(s);
}
```

In this example, the size of the span is determined at runtime, providing more flexibility when working with data.

Enhanced Subspan Capabilities

C++23 introduces further improvements to the **subspan** function, which allows extracting a subrange from an existing span. These enhancements include:

- **Better Handling of Views:** C++23 optimizes **subspan operations** for better performance, particularly when subspans are being created from large spans. This allows for more efficient handling of slices, especially when working with large data structures.
- **Improved Bound Checking:** The bounds-checking mechanisms have been improved in C++23, making it easier to work safely with spans. When slicing or accessing elements beyond the current extent, C++23 ensures that better error messages and checks are provided, improving both usability and safety.

std::span and Container Compatibility

In C++23, **std::span** has been enhanced to support better interaction with **standard containers** such as **std::vector**, **std::array**, and others. The most important change is the **std::span constructor** that can now accept containers directly, making it easier to create spans from existing containers without manually extracting pointers or sizes.

For example:

```
#include <iostream>
#include <vector>
#include <span>

int main() {
    std::vector<int> vec = {10, 20, 30, 40, 50};
    std::span<int> sp(vec); // Directly from a container

    for (auto val : sp) {
        std::cout << val << " ";
    }

    return 0;
}
```

This change simplifies working with spans in real-world code, reducing the need for manual pointer management.

3.2.3 Performance Improvements in `std::span`

C++23 introduces **performance optimizations** in `std::span` to make it even more efficient for performance-critical applications. These improvements mainly focus on reducing the overhead when accessing elements within spans, especially for large datasets. The goal is to ensure that `std::span` provides the best possible performance when used with high-performance applications such as scientific computing, gaming, and data processing. Some of the key performance improvements include:

- **Reduced memory access latency:** The memory access patterns for spans are optimized for modern hardware architectures, reducing cache misses and improving data locality.
- **Increased compiler optimization:** The implementation of `std::span` is now more

amenable to aggressive compiler optimizations, such as loop unrolling and vectorization.

3.2.4 Other Notable Template Improvements

Beyond `std::ranges` and `std::span`, C++23 brings several minor modifications and optimizations to other STL templates:

- **`std::array`**: While not a new template, `std::array` has received several **performance improvements** in C++23, particularly in how it interacts with modern compiler optimizations and reduces overhead when used with small fixed-size arrays.
- **`std::optional`**: The `std::optional` template now offers **better performance** for cases where it is frequently constructed and destructed, improving scenarios like storing optional values in containers or working with result types in algorithms.

Conclusion

C++23 brings significant enhancements to existing STL templates, such as **`std::ranges`** and **`std::span`**, improving both **performance** and **usability**. The new features and optimizations introduced in this version make it easier for developers to work with sequences and data structures efficiently, while offering greater flexibility in modern C++ code.

- **`std::ranges`** benefits from improved adaptors, more powerful algorithms, and enhanced support for container manipulation, allowing developers to write cleaner, more efficient code.
- **`std::span`** is enhanced with better runtime support, improved subspan capabilities, and performance optimizations, making it an even more valuable tool for working with contiguous data.

With these improvements, C++23 sets a new standard for working with data in C++, giving developers the tools they need to write more expressive, high-performance applications.

Chapter 4

New Concepts in C++23 for STL

4.1 `std::concepts`: How These Concepts Simplify Template Understanding and Ensure Compatibility

In C++20, the **concepts** feature was introduced to provide a more expressive and declarative way of specifying requirements for template parameters. Concepts allow developers to specify constraints on types that are passed to templates, enhancing the readability, clarity, and safety of template-based code. In **C++23**, this feature has been significantly expanded and refined, allowing for even more powerful and flexible constraints, particularly in the context of the **Standard Template Library (STL)**.

In this section, we will delve into **`std::concepts`** in C++23, explaining how they simplify template programming, improve template understanding, and help ensure compatibility between different components. We will also examine new and improved concepts in C++23, their practical applications, and how they fit into the broader design philosophy of the STL.

4.1.1 Overview of Concepts

A **concept** in C++ is a set of **requirements** that a type must satisfy in order to be used with a specific template. Prior to C++20, constraints on template types were typically enforced using **SFINAE (Substitution Failure Is Not An Error)** or **enable_if**. While these methods were powerful, they were often verbose, hard to read, and error-prone. Concepts, by contrast, provide a much clearer and more intuitive way of specifying these requirements.

Concepts allow you to specify exactly what kinds of types are valid for a particular template, ensuring that only types that meet certain criteria can be passed to the template. This enhances the type safety of your code by catching errors earlier in the compilation process and provides better error messages when constraints are violated.

4.1.2 How Concepts Simplify Template Understanding

Concepts **simplify template programming** by making it more explicit what kinds of types are acceptable for templates. This eliminates much of the **implicit complexity** that comes with traditional template metaprogramming techniques such as SFINAE.

Clear and Readable Template Constraints

With **concepts**, you can define template constraints in a way that is **self-documenting**. Rather than having to read through complex **enable_if** conditions or meta-functions to understand what types are expected, a **concept** directly expresses the intended behavior.

For example, consider a function template that accepts only **iterable types**:

```
#include <concepts>
#include <vector>
#include <iostream>

template <std::input_iterator It>
void print_elements(It first, It last) {
```

```
    for (auto it = first; it != last; ++it) {
        std::cout << *it << ' ';
    }
}

int main() {
    std::vector<int> vec = {1, 2, 3, 4, 5};
    print_elements(vec.begin(), vec.end());
    return 0;
}
```

Here, the `std::input_iterator` concept makes it clear that the function **print_elements** requires its template argument to be an input iterator, and it is enforced by the compiler. This makes the function's constraint **explicit** and **easy to understand** at the point of use.

Better Error Messages

One of the most powerful features of **concepts** is the **improved error messages** that come with them. Prior to C++20, when a template failed to compile due to an invalid type, the error message was often long, cryptic, and difficult to decipher. Concepts provide more **meaningful error messages** that clearly explain what the problem is.

For example, if we try to pass a non-iterable type to the `print_elements` function:

```
#include <concepts>
#include <iostream>

template <std::input_iterator It>
void print_elements(It first, It last) {
    for (auto it = first; it != last; ++it) {
        std::cout << *it << ' ';
    }
}
```

```
}

int main() {
    int x = 10;
    print_elements(x, x); // Error: 'int' is not an input iterator
}
```

The compiler will generate an error like:

```
error: no matching function for call to 'print_elements'
note: template parameter 'It' has requirements:
      'It' is an input iterator
```

This error message is **much clearer** and helps the developer quickly identify that `int` does not meet the `std::input_iterator` requirement, making debugging simpler.

4.1.3 Concepts and STL: Ensuring Compatibility

One of the primary goals of **concepts** is to ensure **compatibility** between template components by enforcing constraints that ensure the proper type is passed to a template function or class. In C++23, the STL embraces **concepts** as part of its design, ensuring that only compatible types are used with STL algorithms and containers.

Enforcing Type Compatibility in Algorithms

C++23 brings additional **concepts** to improve the expressiveness of STL algorithms. Many STL algorithms now use concepts to enforce more specific type requirements on their parameters, which helps ensure compatibility and improve performance.

For instance, let's consider the `std::ranges::sort` algorithm. In C++23, the algorithm now ensures that the type passed to it meets certain constraints, such as being **swappable** and having a proper **ordering relation**.

```
#include <iostream>
#include <ranges>
#include <vector>
#include <algorithm>

template <std::sortable Range>
void sort_range(Range& r) {
    std::ranges::sort(r);
}

int main() {
    std::vector<int> vec = {4, 2, 5, 1, 3};
    sort_range(vec); // works, since std::vector<int> is sortable
    sort_range("hello"); // Error: char* is not sortable
}
```

In this example, the **std::sortable** concept ensures that the range passed to **std::ranges::sort** is compatible with the algorithm. If the type doesn't meet the **sortable** constraint, the compiler will immediately produce a meaningful error message.

std::ranges::viewable_range and std::input_iterator

With the expansion of **ranges** in C++23, **std::ranges::viewable_range** and **std::input_iterator** concepts have become more crucial in the implementation of various algorithms. These concepts ensure that algorithms work only on ranges that are **viewable** (i.e., valid for range-based operations) or on types that provide **input iteration** capabilities. By using these concepts, we can write algorithms that work seamlessly across different container types, reducing the risk of type mismatches and ensuring that the correct operations are performed.

For example, **std::ranges::viewable_range** ensures that only objects that can be used with a range-based for loop or other range operations are passed to range algorithms:

```
#include <ranges>
#include <vector>

template <std::ranges::viewable_range R>
void print_range(const R& r) {
    for (auto&& elem : r) {
        std::cout << elem << " ";
    }
}

int main() {
    std::vector<int> vec = {1, 2, 3};
    print_range(vec); // works fine, as std::vector is viewable
}
```

This guarantees that only ranges that support the necessary operations (e.g., `begin()`, `end()`, etc.) are passed to algorithms like `print_range`, ensuring **type compatibility** and **correct behavior**.

Specialized Concepts for Container Types

In C++23, specialized concepts for container types ensure that only the correct container types are passed to container-specific functions. For example, containers that support random access iterators can be used with algorithms that require such iterators, and containers that provide a specific comparison function can be passed to sorting algorithms.

```
#include <concepts>
#include <vector>
#include <algorithm>
#include <iostream>

template <std::random_access_iterator It>
```

```

void reverse_range(It first, It last) {
    std::reverse(first, last);
}

int main() {
    std::vector<int> vec = {5, 4, 3, 2, 1};
    reverse_range(vec.begin(), vec.end()); // Works with std::vector
    // reverse_range("hello", "world"); // Error: char* is not a random
    ↪ access iterator
}

```

Here, **`std::random_access_iterator`** ensures that the `reverse_range` function only accepts iterators that support random access, preventing errors caused by passing incompatible types like `char*`, which only supports pointer arithmetic, not random access.

4.1.4 New Concepts in C++23

C++23 introduces several new concepts that improve type safety, enforce algorithm constraints, and simplify STL usage. Some of these new concepts include:

- **`std::movable`**: A concept that requires types to be **move-constructible** and **move-assignable**. It is used to ensure that algorithms that expect to move elements only work on types that support efficient move operations.
- **`std::swappable`**: This concept guarantees that types can be swapped with each other. It is crucial for algorithms like `std::ranges::sort`, where the ability to swap elements is essential for the sorting operation to work.
- **`std::copyable`**: A concept that ensures the type supports **copying**, i.e., both copy constructor and copy assignment. This is useful for algorithms that rely on copying elements from one container to another.

Conclusion

The introduction and refinement of **concepts** in C++23 represent a **major advancement** in the way developers express template requirements. Concepts provide a clear, readable, and efficient way to specify type constraints, helping to **simplify template programming** and improve **template understanding**. They also play a crucial role in **ensuring compatibility** between different components of the STL, making algorithms and data structures more predictable and easier to work with.

With new concepts like `std::movable`, `std::swappable`, and `std::copyable`, C++23 continues to refine and expand the expressiveness of its template system, providing developers with powerful tools to write cleaner, more maintainable, and type-safe code. Concepts are rapidly becoming an integral part of modern C++ programming, enhancing the **expressiveness**, **safety**, and **compatibility** of STL algorithms and beyond.

4.2 `std::ranges`: Expanding Template Capabilities and Enhancing Data Interactions

The `std::ranges` library, first introduced in C++20, brought a revolutionary shift in how C++ developers interact with sequences and containers. By providing an abstraction layer over traditional algorithms and iterators, it allowed for more expressive, flexible, and type-safe operations on data ranges. With C++23, the capabilities of `std::ranges` have been significantly expanded, introducing new concepts, optimizations, and features that further enhance its integration with the Standard Template Library (STL).

In this section, we will explore how C++23 expands the power of `std::ranges`, not only improving the way we interact with data, but also enhancing template capabilities and offering new methods for working with data ranges more efficiently. We will discuss the changes made to `std::ranges` in C++23, how they improve the overall development experience, and the benefits they bring to template-based programming.

4.2.1 The Role of `std::ranges` in C++ Template Programming

Before C++20, manipulating containers and data sequences typically required a combination of **iterators**, **loops**, and **standard algorithms**. While powerful, this approach could be verbose, error-prone, and less flexible, especially when it came to handling various types of containers or data structures. The introduction of `std::ranges` aimed to address these concerns by simplifying interactions with data ranges and making them more declarative and intuitive. With C++23, the `std::ranges` library evolves further to bring additional features and improvements to its functionality. Key to this evolution are the new **concepts** and **template capabilities** that make it even easier to compose, filter, transform, and manipulate ranges, while also ensuring type-safety and better integration with modern C++ idioms.

4.2.2 Key Enhancements to `std::ranges` in C++23

C++23 focuses on enhancing the functionality, efficiency, and expressiveness of `std::ranges` through several new features, concepts, and optimizations. Here are some of the key updates and their impact on how we interact with data ranges:

New Range Adaptors for More Expressive Data Transformations

Range adaptors in `std::ranges` provide a flexible and composable way to transform, filter, or modify data within a range. C++23 introduces several new adaptors, expanding the toolkit available to developers for manipulating sequences of data. These new adaptors aim to make it easier to work with more complex data transformations and operations, especially when working with larger data sets or when performance is a critical concern.

- **`std::ranges::view::reverse`**: While **`reverse`** was available in previous versions of C++, C++23 introduces an improved version that is **more efficient** and **faster**, particularly for large ranges. This is particularly useful when reversing large data sets, as the new version reduces unnecessary allocations and memory overhead.

- **`std::ranges::view::slide`**: This new adaptor allows for **sliding windows** over ranges, providing an elegant way to handle fixed-size subsets of ranges. It allows the developer to specify a window size and move it over the range, useful for applications like signal processing, time-series analysis, or sliding window algorithms.

```
#include <ranges>
#include <vector>
#include <iostream>

int main() {
    std::vector<int> vec = {1, 2, 3, 4, 5, 6};
    auto slides = vec | std::ranges::view::slide(3); // Sliding
    ↪ window of size 3

    for (auto&& slide : slides) {
        for (auto&& val : slide) {
            std::cout << val << ' ';
        }
        std::cout << '\n';
    }

    return 0;
}
```

In this example, the `slide` adaptor creates sliding windows of the specified size (3 in this case), making it easier to iterate over fixed-length subsets of the range.

- **`std::ranges::view::chunk`**: This new adaptor divides a range into **chunks** of specified sizes. It can be particularly useful for parallel processing, batching operations, or breaking up large data sets into smaller, more manageable sections.

```
auto chunks = vec | std::ranges::view::chunk(2);
```

These new range adaptors make working with data more **flexible** and **expressive**, allowing developers to perform more complex data manipulations in a clean, composable manner.

Enhanced Range-based Algorithms

C++23 builds on the existing algorithms in the `std::ranges` library by introducing several new range-based algorithms. These algorithms are designed to operate directly on ranges, rather than requiring explicit iterators, improving the readability and efficiency of the code.

- **`std::ranges::find_if_not`**: This algorithm is similar to `find_if`, but instead of finding the first element that satisfies a condition, it finds the first element that **does not** satisfy the condition. This allows for more natural expressions of certain search patterns.
- **`std::ranges::is_sorted_until`**: This algorithm finds the point where a range stops being sorted, making it easier to check whether a range is sorted, and if not, find the point of failure.
- **`std::ranges::is_partitioned`**: This algorithm checks whether the elements in a range are partitioned according to a predicate. This is particularly useful when dealing with ranges that have already been partially sorted or partitioned by another algorithm.

By providing direct, range-based alternatives to traditional iterator-based algorithms, C++23 makes it easier and more intuitive to work with ranges while keeping the code clean and readable.

Range-based Views for More Efficient Data Manipulation

In C++23, there are **optimizations** to range views, which allow data to be processed lazily (i.e., evaluated only when needed) rather than eagerly constructing intermediate containers. This

approach is particularly useful for improving **memory efficiency** when working with large datasets, as it avoids unnecessary copies or memory allocations.

- **Lazy evaluation:** `std::ranges::view` enables **lazy evaluation** of transformations and filtering operations. For example, when a range is transformed using the `transform` adaptor, the result is not immediately computed; instead, it is computed only when the transformed range is iterated over, improving performance and reducing overhead in cases where transformations are not immediately required.

```
auto view = vec | std::ranges::view::transform([](int n) { return n *
↪ 2; });
```

Here, the transformation of multiplying each element by two happens lazily, only when the range is actually iterated.

`std::ranges` and Concepts: Enforcing Type Safety

C++23 introduces new **concepts** and refines existing ones, further strengthening the ability to specify the exact type requirements for ranges and algorithms. These concepts improve the **type safety** of operations and ensure that the proper types are used in conjunction with range-based algorithms.

For example, concepts like **`std::ranges::input_range`** and

`std::ranges::viewable_range` are now more **strictly enforced** in C++23, ensuring that algorithms only work with ranges that support the necessary operations (e.g., iterating, indexing, and so on).

```
template <std::ranges::input_range R>
void process_range(R&& range) {
    for (auto&& elem : range) {
        // Process each element
    }
}
```

```
}  
}
```

In this case, the function `process_range` can only accept types that are **input ranges**, enforcing that the passed range can be **iterated over** safely.

By leveraging these concepts, C++23 ensures that only valid ranges and iterators are passed to range-based algorithms, preventing **type errors** at compile time and improving overall code safety.

4.2.3 Performance Improvements in `std::ranges`

With **C++23**, **performance optimizations** have been implemented throughout the `std::ranges` library. These improvements are particularly beneficial for applications that need to process large amounts of data or require high performance in data manipulation operations.

Optimized Range Adaptors

Many of the range adaptors introduced in C++23, such as `slide`, `chunk`, and `reverse`, have been optimized for better **memory locality** and **cache performance**. This means that operations on large ranges are more efficient, with reduced overhead compared to previous versions of C++.

Parallel Execution Support

C++23 enhances the ability to perform **parallel executions** over ranges, enabling developers to leverage **multi-core processors** more effectively. Range-based algorithms can now more easily support parallel execution, allowing for faster processing of large datasets in cases where such optimizations are beneficial.

- For example, the `std::ranges::sort` algorithm can now take advantage of multi-threaded parallel execution for sorting large ranges, dramatically improving performance in certain scenarios.

4.2.4 Integration with Other C++23 Features

The `std::ranges` library in C++23 integrates seamlessly with other **new C++23 features**, such as `std::span`, `std::format`, and `std::atomic`, allowing for even greater flexibility in data manipulation.

- **`std::span`**: A `std::span` provides a **view** of an array or part of an array and integrates well with `std::ranges`. This means you can treat `std::span` objects as ranges and apply `std::ranges` algorithms to them without needing to worry about pointer arithmetic.
- **`std::format`**: The `std::ranges` library can be used alongside the new `std::format` functionality to format the elements in a range before outputting them. This allows for more powerful and concise ways of printing or processing data ranges.

Conclusion

C++23 takes the already powerful `std::ranges` library introduced in C++20 and significantly expands its capabilities, making it even more expressive, efficient, and integrated with the rest of the STL. With new range adaptors, enhanced algorithms, and optimizations for memory and performance, `std::ranges` in C++23 makes it easier to manipulate and interact with data in a clear, efficient, and type-safe manner.

The combination of **range-based views**, **concepts**, and **performance improvements** ensures that `std::ranges` remains a key feature in modern C++ programming, enabling developers to write cleaner, more expressive code that handles complex data interactions with ease. With these improvements, C++23 further cements `std::ranges` as a powerful tool in the developer's toolkit for working with data.

4.3 `std::span`: How It Can Be Used to Manage Arrays More Safely and Flexibly

In modern C++, managing raw arrays has often been seen as error-prone and cumbersome. While raw pointers provide flexibility, they often require manual memory management and careful bounds checking, making the code harder to maintain and more susceptible to bugs such as buffer overflows or memory leaks. To address these concerns, the `std::span` class was introduced in C++20 as a lightweight, non-owning view of a sequence of elements, providing a more **type-safe**, **flexible**, and **convenient** way to work with arrays, buffers, and other contiguous sequences of data.

With C++23, `std::span` has been further enhanced, improving its integration with modern C++ features and expanding its use cases. This section will explore how `std::span` improves array management by providing safety, flexibility, and performance while making working with arrays and contiguous data structures much more efficient.

4.3.1 Introduction to `std::span`

`std::span` is part of the `<span>` header in the C++ standard library, and it represents a **view** or **window** over a contiguous sequence of elements, such as an array, a vector, or a dynamically allocated memory buffer. The primary purpose of `std::span` is to provide a **non-owning** view of an array or sequence of elements without needing to pass around raw pointers or manually track the size of the array.

Key properties of `std::span` include:

- It provides **bounds checking**, preventing access to out-of-bounds elements and avoiding common mistakes associated with pointer arithmetic.
- It can represent both **static arrays** (arrays with a fixed size known at compile time) and **dynamic arrays** (e.g., vectors, arrays with size determined at runtime).

- It allows efficient **passing of arrays** or **buffers** to functions without needing to copy the data.

A `std::span` object is constructed with a pointer to the beginning of the sequence and the size of the sequence. This allows for safe and efficient access to a range of data, without taking ownership or requiring explicit memory management.

Basic Example of `std::span`

```
#include <span>
#include <iostream>
#include <vector>

void print_span(std::span<int> span) {
    for (auto elem : span) {
        std::cout << elem << ' ';
    }
    std::cout << std::endl;
}

int main() {
    int arr[] = {1, 2, 3, 4, 5};
    std::span<int> span(arr, 5); // span over the array

    print_span(span); // prints: 1 2 3 4 5

    // Can also work with vectors
    std::vector<int> vec = {10, 20, 30};
    print_span(vec); // prints: 10 20 30
}
```

In this example, `std::span` is used to wrap both a raw array (`arr`) and a `std::vector<int>`. By passing a `std::span` to the function `print_span`, we can

safely access the data from both container types without having to worry about their underlying implementation.

4.3.2 Benefits of Using `std::span` Over Raw Pointers

Before the introduction of `std::span`, developers relied on raw pointers and manual size management to handle arrays. This could lead to various issues, including:

- **Lack of bounds checking:** A raw pointer does not inherently know the size of the array it points to, making it easy to access memory outside the bounds of the array, causing undefined behavior and hard-to-diagnose bugs.
- **Memory management overhead:** Raw pointers often require explicit memory management, leading to potential memory leaks or double frees if not handled carefully.
- **Inconsistent array handling:** Passing arrays to functions typically required passing both the pointer and the size, often leading to confusion and errors.

`std::span` solves these issues by providing several key benefits:

- **Bounds checking:** By wrapping a contiguous range of data, `std::span` ensures that any access to the data is within bounds. If a function attempts to access an element outside the range, the compiler can raise a warning or error.

```
std::span<int> span(arr, 5);  
std::cout << span[5]; // Out of bounds, undefined behavior
```

- **Automatic size tracking:** `std::span` knows its size, so functions accepting a `span` do not require the array size to be passed explicitly. This reduces the potential for errors and makes the code more maintainable.

- **Safety and clarity:** `std::span` expresses the intent more clearly than raw pointers. It is immediately obvious to readers of the code that a `std::span` refers to a sequence of elements, whereas raw pointers may be more ambiguous.
- **Non-owning nature:** `std::span` does not take ownership of the data it refers to. It only provides a **view** over the data, meaning that it does not modify the ownership or lifetime of the sequence it points to.

4.3.3 How `std::span` Works with Arrays, Vectors, and Other Containers

One of the most powerful features of `std::span` is its ability to work with a wide range of data types. In addition to supporting raw arrays, `std::span` can also be used with **`std::vector`**, **`std::array`**, and even **C-style strings**. This flexibility makes `std::span` an ideal abstraction for handling a wide variety of data structures in a consistent manner.

`std::span` with Raw Arrays

When used with a raw array, `std::span` provides a safer, more flexible way to represent and manipulate the array. For example, it provides a consistent interface for **iterating**, **slicing**, and **viewing** portions of the array.

```
int arr[] = {1, 2, 3, 4, 5};
std::span<int> span(arr);
std::cout << span.size(); // Outputs: 5
```

Here, `std::span` automatically determines the size of the array and provides a range over it.

`std::span` with `std::vector`

`std::span` can also be constructed from `std::vector` (and other standard container types that expose a contiguous memory layout, like `std::deque`). This enables functions to operate on ranges of data without knowing the underlying container type, making the code more **generic** and **flexible**.

```
std::vector<int> vec = {10, 20, 30, 40};  
std::span<int> span(vec);
```

This allows `std::span` to treat the vector as a sequence of elements, allowing for safer operations on the vector.

`std::span` with `std::array`

Another commonly used container that works seamlessly with `std::span` is `std::array`. Since `std::array` is a **fixed-size array**, `std::span` can take advantage of its size at compile time, providing even more type safety.

```
std::array<int, 5> arr = {1, 2, 3, 4, 5};  
std::span<int> span(arr); // No need to pass the size separately
```

`std::span` with C-style Strings

A very practical use of `std::span` is with **C-style strings**. Since C-style strings are often passed as pointers to `char` with an implicit size, `std::span` makes it easy to manage them safely.

```
const char* str = "Hello, world!";  
std::span<char> span(str, std::strlen(str));
```

With `std::span`, working with C-style strings becomes more manageable, eliminating the need for manually tracking lengths or worrying about off-by-one errors.

4.3.4 Slicing and Subranges with `std::span`

One of the great features of `std::span` is its ability to **slice** ranges. This allows for **subranges** of the original sequence to be created without the need for copying the data. By using

`std::span::subspan`, you can create a new span that represents a subset of the original data.

Example of Slicing a Span

```
std::array<int, 5> arr = {1, 2, 3, 4, 5};
std::span<int> span(arr);

std::span<int> subrange = span.subspan(1, 3); // Elements from index 1 to
↪ 3
for (auto elem : subrange) {
    std::cout << elem << ' '; // Output: 2 3 4
}
```

In this example, the `subspan` function creates a **view** of the array starting at index 1 and containing 3 elements.

4.3.5 Improvements in C++23 for `std::span`

While `std::span` was introduced in C++20, C++23 has introduced several small improvements to its usability and behavior, including:

- Better integration with **`std::ranges`**: `std::span` can now be used seamlessly with range-based algorithms, making it even easier to use in modern C++ code.
- Enhanced **bounds-checking** and **error reporting**: C++23 improves the compiler's ability to report errors when accessing out-of-bounds elements, making `std::span` even safer.
- Improved performance optimizations for subspans and slicing, making these operations more efficient for large datasets.

Conclusion

`std::span` is a powerful addition to the C++ Standard Library that simplifies working with contiguous data structures. It offers a safer, more flexible alternative to raw pointers, with added benefits like **bounds checking**, **type safety**, and **performance optimizations**. In **C++23**, these capabilities have been further enhanced, making `std::span` an even more essential tool for modern C++ programming. By abstracting away the complexities of manual memory management and offering a consistent interface for a variety of data types, `std::span` makes handling arrays and buffers easier and more robust than ever before.

Chapter 5

STL Containers in C++23

5.1 Modifications to Containers like `std::vector` and `std::list`

In **C++23**, the **Standard Template Library (STL)** has seen important updates across various components, including the **container types**. Two of the most commonly used containers, `std::vector` and `std::list`, have undergone modifications that improve their performance, usability, and flexibility. These changes enhance the overall experience of working with the STL in modern C++ development, ensuring that developers can take advantage of the latest language features to write more efficient and maintainable code.

This section delves into the modifications introduced in **C++23** for the `std::vector` and `std::list` containers, explaining their key updates, practical benefits, and how they improve the handling of dynamic arrays and linked lists.

5.1.1 Overview of `std::vector` and `std::list`

Before diving into the C++23 updates, let's briefly refresh the functionality of these two containers:

- **`std::vector`** is a dynamic array that provides random access to elements and can grow or shrink in size automatically. It is highly efficient when it comes to **random access** and appending elements to the end of the container. However, insertion or deletion of elements in the middle or at the beginning of a vector can be expensive due to the need to shift elements.
- **`std::list`** is a **doubly linked list**, which allows fast insertions and deletions anywhere in the container, as no shifting of elements is required. However, it sacrifices **random access**, meaning accessing an element by index takes linear time. It is primarily used when frequent insertions and deletions are needed.

Both containers have been critical in C++ programming due to their simplicity and general-purpose nature, but C++23 introduces specific updates that refine their behavior and functionality.

5.1.2 Key Modifications in `std::vector`

Support for Move-Only Types

One of the important improvements to `std::vector` in C++23 is its enhanced support for **move-only types**. In earlier versions of C++, `std::vector` could hold elements of move-only types like `std::unique_ptr`, but it often required additional care when pushing, popping, or resizing the vector. This resulted in either copying or reallocation that could prevent move semantics from being fully utilized.

In **C++23**, **move-only types** are now **handled more efficiently** when working with `std::vector`. Specifically, the internal allocation strategy has been updated to better support

moving elements during reallocation, ensuring that move-only types, such as `std::unique_ptr` and other types with deleted copy constructors, can be used in a vector without performance penalties or needing additional workarounds.

Faster Insertion at the End

Another key improvement in C++23 is the **faster insertion** at the end of `std::vector` through the **`std::vector::emplace_back`** and **`std::vector::push_back`** operations. C++23 introduces optimizations that make these operations **more efficient** by reducing the internal memory reallocations and improving how the vector's capacity is managed.

- In previous versions of C++, when a vector's size exceeded its capacity, it would need to reallocate memory and copy elements to the new location. This could result in unnecessary copying or moving of elements, which was especially costly with large objects.
- With the C++23 update, `std::vector` uses **smarter growth strategies** and internal buffering techniques to reduce the need for reallocations and improve overall insertion time.

This means that inserting elements at the end of the vector (especially when the container is near its capacity) is now **significantly faster**, especially in cases where move semantics are involved.

New **reserve** and **shrink_to_fit** Behavior

In C++23, `std::vector`'s **reserve** and **shrink_to_fit** functions have been improved to better manage memory.

- **reserve:** The `reserve` function now ensures that additional space is allocated in a way that minimizes reallocations in the future. It introduces a more **efficient memory allocation scheme**, which dynamically calculates the amount of space to allocate without causing excessive memory fragmentation.

- **shrink_to_fit**: The `shrink_to_fit` function is now optimized to release excess capacity more effectively. This ensures that the vector shrinks to the exact size required without leaving unnecessary empty slots. This improvement not only saves memory but also makes memory usage more predictable, especially when a vector's size is reduced after many elements are removed.

Memory Allocation Enhancements

The **allocation strategy** for `std::vector` has been optimized in C++23 to improve both **performance** and **predictability** in memory usage. Specifically, **custom allocators** are now more efficiently integrated into the vector's allocation process. This improvement makes it easier to use `std::vector` with custom memory management strategies, which can be particularly useful in performance-critical applications.

5.1.3 Key Modifications in `std::list`

Although `std::list` is not as commonly used as `std::vector` due to its lack of random access, it has still undergone several improvements in C++23, making it a more effective and flexible container for certain use cases.

Optimized Node Management

One of the more notable changes in C++23 is the **optimized management of the internal nodes** in `std::list`. In earlier versions of C++, `std::list` used a straightforward doubly linked list implementation where each element had pointers to its previous and next neighbors. However, the overhead of maintaining these pointers, particularly during operations like insertion and removal, could be inefficient in some cases.

In C++23, `std::list` has been optimized to reduce **memory fragmentation** and improve **node allocation**. The new improvements in **node pooling** and **allocator management** result in **better memory locality** and **lower allocation costs**, especially when working with large lists or performing many insertions and deletions.

Improved `splice` Operations

The **`splice`** operation in `std::list`, which is used to transfer elements from one list to another, has been enhanced for **better performance**. In C++23, the `splice` operation has become **more efficient** in terms of both time and memory usage, particularly for large lists. The optimizations allow elements to be transferred between lists without needing to allocate new nodes or copy elements unnecessarily.

This improvement is especially useful when working with lists that require frequent merging, rearranging, or transferring of elements, making `std::list` a more attractive option for these tasks compared to other containers.

New Algorithms for `std::list`

C++23 has added new **algorithm support** for `std::list`, particularly around **range-based algorithms**. This includes algorithms such as **`std::ranges::reverse`**, **`std::ranges::unique`**, and **`std::ranges::sort`**, which can now be used directly with `std::list`. Previously, using algorithms on `std::list` required explicit conversion to other container types or manual iteration.

With these new range-based algorithm enhancements, developers can now apply standard algorithms directly to `std::list` containers, making the code cleaner and more expressive while maintaining the performance benefits of the list structure.

Conclusion

The **C++23 updates** to `std::vector` and `std::list` bring several important improvements to these already powerful containers, increasing their efficiency, flexibility, and usability in modern C++ development.

- **`std::vector`** benefits from enhancements to its move semantics, memory management, and optimized insertion at the end of the container. These improvements make it faster and more predictable, especially when working with large amounts of data or move-only types.

- **`std::list`** sees improvements in internal node management, better `splice` operations, and integration with new range-based algorithms. These changes make it a more effective choice for use cases where frequent insertions and deletions are required, and it now integrates better with the overall C++23 algorithm ecosystem.

These updates to `std::vector` and `std::list` in **C++23** help developers write more efficient, maintainable, and expressive code, while also offering significant performance improvements for certain use cases. With these modifications, both containers continue to play a crucial role in modern C++ programming, providing the flexibility and power needed to manage data in a wide range of applications.

5.2 New or Modified Functions in Containers, like `insert_or_assign` and `contains`

In **C++23**, several **new** and **modified** functions have been introduced to improve the usability, flexibility, and performance of the **Standard Template Library (STL) containers**. These changes aim to simplify common operations on containers, enhance their functionality, and make them more intuitive to use. Key functions like `insert_or_assign` and `contains` have been added or enhanced in containers such as `std::map`, `std::unordered_map`, and other container types, providing better control over data insertion, search, and assignment.

This section delves into the **new** and **modified functions** in STL containers in **C++23**, focusing on the addition of `insert_or_assign`, `contains`, and other useful utilities. Understanding these new functions will enable developers to write more concise, readable, and efficient code when working with STL containers.

5.2.1 The Introduction of `insert_or_assign`

The `insert_or_assign` function is a new addition to the **associative containers** like `std::map`, `std::unordered_map`, `std::set`, and `std::unordered_set`. This function combines two operations that were previously separate: **insertion** and **assignment**.

Functionality of `insert_or_assign`

In earlier versions of C++, if you wanted to either **insert** a new element or **update** an existing one in an associative container, you typically had to perform two separate operations:

1. **Check** if the element exists using `find` or `count`.
2. **Insert** or **update** the element accordingly.

The `insert_or_assign` function simplifies this process by handling both the **insertion** and **assignment** in a single call. It works as follows:

- If the key already exists in the container, it updates the corresponding value.
- If the key does not exist, it inserts a new key-value pair.

This eliminates the need for explicit checks before performing these operations, streamlining the code and reducing boilerplate.

Syntax and Example Usage

The syntax for `insert_or_assign` is as follows:

```
std::map<Key, T>::insert_or_assign(const Key& key, T&& value);  
std::unordered_map<Key, T>::insert_or_assign(const Key& key, T&& value);
```

Here, `Key` is the key type, and `T` is the type of the value associated with the key.

- If the key exists, the associated value is updated with the provided value.

- If the key does not exist, a new key-value pair is inserted into the container.

Example:

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

int main() {
    std::map<int, std::string> my_map;

    // Insert a key-value pair
    my_map.insert_or_assign(1, "One");
    my_map.insert_or_assign(2, "Two");

    // Update an existing key
    my_map.insert_or_assign(1, "Uno");

    // Print the map
    for (const auto& pair : my_map) {
        std::cout << pair.first << " => " << pair.second << std::endl;
    }

    return 0;
}
```

Output:

```
1 => Uno
2 => Two
```

In this example, the first insertion adds the key-value pairs (1, "One") and (2, "Two"). When we call `insert_or_assign(1, "Uno")`, the value for key 1 is updated to "Uno".

Benefits of `insert_or_assign`

- **Cleaner code:** It eliminates the need for manual checking of whether an element exists before deciding whether to insert or update.
- **Efficiency:** Internally, the function performs the required operation with a minimal number of lookups, improving performance for containers with many elements.
- **Simplicity:** It simplifies many common patterns where you need to insert or update elements in associative containers.

5.2.2 The `contains` Function

The `contains` function is a new addition introduced in **C++23** for **associative containers** like `std::map`, `std::unordered_map`, `std::set`, and `std::unordered_set`. Before the introduction of `contains`, checking whether a key existed in these containers required using functions like `find` or `count`. However, these functions often led to more verbose code or required additional steps.

Functionality of `contains`

The `contains` function simplifies checking whether a key exists in an associative container by directly returning a **boolean value**. It checks if a given **key** is present in the container and returns `true` if the key is found, or `false` if it is not.

The syntax for `contains` is as follows:

```
std::map<Key, T>::contains(const Key& key) const;  
std::unordered_map<Key, T>::contains(const Key& key) const;  
std::set<Key>::contains(const Key& key) const;  
std::unordered_set<Key>::contains(const Key& key) const;
```

Here, `Key` is the type of the key, and `T` is the type of the value (for `map` and `unordered_map`).

Example Usage of `contains`

Example with `std::unordered_map`:

```
#include <iostream>
#include <unordered_map>

int main() {
    std::unordered_map<int, std::string> my_map = {{1, "One"}, {2, "Two"},
    ↪ {3, "Three"}};

    // Use contains to check if a key exists
    if (my_map.contains(2)) {
        std::cout << "Key 2 exists with value: " << my_map[2] <<
        ↪ std::endl;
    }

    if (!my_map.contains(4)) {
        std::cout << "Key 4 does not exist." << std::endl;
    }

    return 0;
}
```

Output:

```
Key 2 exists with value: Two
Key 4 does not exist.
```

In this example, `contains` is used to check whether keys 2 and 4 exist in the map. This is much more intuitive and concise than using `find` or `count`.

Benefits of `contains`

- **Simplicity:** It provides a simple and intuitive way to check for the existence of a key in a

container, reducing the need for more complex code like `find` or `count`.

- **Readability:** The code becomes more expressive, making it easier to understand the intent of checking for the presence of a key.
- **Efficiency:** Just like `find`, `contains` has a constant-time complexity ($O(1)$) for unordered containers and logarithmic complexity ($O(\log n)$) for ordered containers.

5.2.3 Other Notable New Functions in C++23 Containers

In addition to `insert_or_assign` and `contains`, several other **new** or **modified** functions have been introduced or improved across various STL containers in **C++23**. These include:

`erase_if` Function (for `std::map`, `std::unordered_map`, and others)

C++23 introduces the **`erase_if`** function, which provides a way to remove elements from a container based on a predicate.

- For associative containers like `std::map` or `std::unordered_map`, `erase_if` allows you to erase elements for which the predicate returns `true`.
- This function eliminates the need for manually iterating over the container to find and erase elements based on a condition.

Example with `std::unordered_map`:

```
#include <iostream>
#include <unordered_map>

int main() {
    std::unordered_map<int, std::string> my_map = {{1, "One"}, {2, "Two"},
    ↪ {3, "Three"}};
```

```
// Erase elements where the key is greater than 1
my_map.erase_if([](const auto& pair) {
    return pair.first > 1;
});

// Print the remaining elements
for (const auto& pair : my_map) {
    std::cout << pair.first << " => " << pair.second << std::endl;
}

return 0;
}
```

Output:

```
1 => One
```

Here, the `erase_if` function removes all elements where the key is greater than 1.

`insert_or_erase` (for `std::map` and `std::unordered_map`)

This function combines the functionality of inserting a key-value pair if the key does not exist, or erasing the key-value pair if it does. It's useful for situations where you want to **update** or **remove** a key-value pair based on a condition.

Conclusion

The introduction of functions like **`insert_or_assign`** and **`contains`** in **C++23** marks a significant improvement in the flexibility, efficiency, and ease of use of STL containers. These new functions make working with **associative containers** like `std::map` and `std::unordered_map` more concise, readable, and maintainable, reducing the need for manual checks and improving code clarity.

Moreover, the addition of **`erase_if`**, **`insert_or_erase`**, and other new utilities enrich the

toolkit available to C++ developers, empowering them to write cleaner and more efficient code when managing containers.

These changes make STL containers in **C++23** more powerful, intuitive, and easier to use in a wide range of applications, from simple data structures to complex algorithms.

Chapter 6

New Libraries and Tools in STL

6.1 `std::expected`: Handling Errors

One of the most notable and significant additions to the **Standard Template Library (STL)** in **C++23** is the introduction of **`std::expected`**, a new class template designed to handle errors and exceptional situations more effectively and efficiently. The **`std::expected`** type addresses several long-standing issues in error handling in C++ by providing a robust, modern approach to managing errors without the overhead of exceptions. It is part of the broader effort in the C++23 standard to make error handling safer and more flexible.

In this section, we will delve into the **`std::expected`** type, explore its features, advantages, and provide guidance on how to use it effectively in your C++ programs.

6.1.1 Introduction to `std::expected`

In traditional C++ error handling, there are two primary methods for signaling errors: **exceptions** and **error codes**. Both have their strengths and weaknesses:

- **Exceptions** provide a clean, centralized way to handle errors, but they can introduce

performance overhead due to stack unwinding and the need for try-catch blocks.

- **Error codes**, on the other hand, are a simpler, low-cost mechanism but they can lead to scattered error-handling logic and harder-to-maintain code.

std::expected aims to combine the best of both worlds by offering a type that allows errors to be handled explicitly without the need for exceptions. It provides a type-safe alternative that enables developers to return either a **successful result** or an **error state** while keeping error-handling logic clear and concise.

The class is typically used to return a value of type **T** (on success) or an **error type** (on failure).

std::expected was introduced to bridge the gap between exceptions and error codes, giving developers a more expressive, flexible way to handle errors in C++ programs.

The Basic Concept of **std::expected**

At its core, **std::expected** is a template class that holds either:

- A valid result (of type **T**), or
- An **error value** (of type **E**).

The syntax for **std::expected** looks like this:

```
std::expected<T, E>
```

- **T** is the type of the expected value when the operation is successful.
- **E** is the type of the error when the operation fails.

The **expected** class is designed to hold one of these two states and expose methods to interact with either the value or the error. This gives developers clear, deterministic control over error propagation.

6.1.2 Structure and Usage of `std::expected`

Creating and Accessing a Value

To work with `std::expected`, developers can use the constructor to create either a valid result or an error state.

- If the operation is successful, you create an instance of `std::expected` with the expected value.
- If the operation fails, you create an instance with an error value.

For example, here is how you can create a successful result or an error state:

```
#include <expected>
#include <iostream>
#include <string>

// A simple function that either returns a string result or an error
std::expected<std::string, std::string> get_user_name(bool success) {
    if (success) {
        return "John Doe"; // Return a successful result
    } else {
        return std::unexpected("Failed to get user name"); // Return an
        ↪ error
    }
}

int main() {
    auto result = get_user_name(true); // Change to 'false' to simulate
    ↪ failure

    if (result) {
```

```
        std::cout << "Success: " << *result << std::endl;    // Dereference
        ↪ to get the result
    } else {
        std::cout << "Error: " << result.error() << std::endl; // Access
        ↪ the error state
    }

    return 0;
}
```

Output (Success Case):

```
Success: John Doe
```

Output (Failure Case):

```
Error: Failed to get user name
```

In this example:

- The function `get_user_name` returns a `std::expected<std::string, std::string>` type.
- If the operation succeeds, it returns a valid string. If the operation fails, it returns an error message encapsulated in `std::unexpected`.

Key Methods in `std::expected`

There are several important methods in `std::expected` that help manage the result or error:

- **operator bool():** This method allows the object to be used in a boolean context. It returns `true` if the result is valid, and `false` if an error occurred.

```
if (result) { /* success */ } // Works similarly to checking for  
↳ success
```

- **operator*()** and **operator->()**: These operators allow direct access to the result when it is valid. They are similar to dereferencing a pointer.

```
auto value = *result; // Access the successful result
```

- **error()**: If the operation failed, **error()** gives access to the error value.

```
auto err = result.error(); // Access the error state
```

- **has_value()**: Returns **true** if the expected value is present, **false** otherwise.

```
if (result.has_value()) { /* Handle success */ }
```

- **value()**: Returns the valid value. If there is no value (i.e., an error state), it throws **std::bad_expected_access**.

```
auto value = result.value(); // Throws if no value is available
```

std::unexpected

std::unexpected is a special object that is used to represent the error part of the **std::expected** type. You can create an error value using **std::unexpected**, which allows the error state to be propagated.

```
return std::unexpected("Error message");
```

This is similar to the `std::exception` mechanism in exception-based error handling, but it works in the **expected** type, thus allowing developers to handle errors explicitly without using exceptions.

6.1.3 Advantages of Using `std::expected`

Explicit Error Handling

Unlike exceptions, which often require try-catch blocks scattered across the codebase, **std::expected** provides an explicit, predictable way to manage errors. You always know whether a function has succeeded or failed based on whether the object contains a value or an error. This makes error handling both **deterministic** and **visible**, as there's no need to rely on the exceptions mechanism or check global states.

Performance Efficiency

`std::expected` avoids the overhead associated with throwing and catching exceptions. It uses a **lightweight mechanism** for indicating errors, which can be more efficient in performance-critical code where exceptions are considered too costly.

Cleaner Code

When using `std::expected`, you no longer need to write boilerplate error-checking code or rely on implicit return codes. The code remains cleaner and more expressive:

- There is no need to check error codes manually.
- The success and failure paths are clearly distinct, improving both maintainability and readability.

Composability

`std::expected` allows for better **composability** when combining multiple functions that may return errors. You can chain operations and propagate errors through them without manually checking after each step.

```
auto result = first_function()
    .and_then(second_function)
    .and_then(third_function);
```

Each function in the chain can return an expected value or propagate an error, making the logic concise and expressive.

6.1.4 Common Use Cases for `std::expected`

- **Error propagation:** When an operation might fail but does not require exception handling, `std::expected` provides a type-safe way to propagate errors.
- **Returning multiple outcomes:** Functions that may have two distinct outcomes (success or failure) but need to return complex results are great candidates for `std::expected`.
- **Error reporting:** Instead of throwing exceptions or returning raw error codes, `std::expected` allows you to encapsulate the error information in a more structured way.

Conclusion

`std::expected` is a powerful and flexible tool introduced in **C++23** for more explicit and predictable error handling. It combines the clarity of error codes with the safety and efficiency of modern C++ features, allowing developers to handle errors without resorting to exceptions. By offering explicit success and error states, along with useful utility methods, `std::expected` simplifies error management while avoiding unnecessary overhead, making it a compelling alternative to traditional error handling techniques.

This new addition to the C++ standard library helps achieve cleaner, more maintainable, and performant code in applications where error handling is crucial but exceptions are not ideal. By adopting `std::expected`, developers can enjoy better error handling in modern C++ programming, while ensuring their code remains both efficient and clear.

6.2 `std::sem` and `std::latch`: Advanced Concurrency Tools in C++23

The C++23 standard introduces two new concurrency primitives, `std::sem` (semaphore) and `std::latch`, which provide enhanced tools for synchronization in concurrent programming. These new constructs make it easier for developers to write efficient, scalable, and easy-to-understand multi-threaded programs. As modern applications increasingly rely on multi-core processors and parallelism to improve performance, the demand for more flexible and powerful concurrency mechanisms has grown. In this section, we will explore the design, use, and benefits of `std::sem` and `std::latch`, and explain how they can simplify the task of managing synchronization in C++ programs.

6.2.1 Introduction to `std::sem` (Semaphore)

A **semaphore** is a well-known synchronization primitive that has been part of many programming languages and libraries. It controls access to a resource by maintaining a counter that represents the number of available units of the resource. Semaphores are typically used in situations where a limited number of threads can access a resource concurrently, making them essential in many real-time and system programming applications.

In C++23, the `std::sem` class template was introduced as a standardized semaphore implementation. This makes semaphores part of the C++ standard library, eliminating the need for custom implementations or reliance on platform-specific APIs.

The Structure of `std::sem`

The `std::sem` class is part of the `<semaphore>` header, and its primary purpose is to synchronize access to a shared resource among multiple threads. The semaphore manages a counter, which is initially set to a specified value representing the number of available resources or "permits". Threads attempting to acquire the semaphore will either:

- **Wait** if the counter is zero (i.e., no resources are available), or
- **Acquire** the semaphore, decrementing the counter when resources are available.

The `std::sem` class allows two main operations:

1. **acquire()**: A thread attempts to decrement the semaphore count (i.e., acquire a resource). If the count is zero, the thread blocks until the semaphore count is greater than zero.
2. **release()**: A thread increments the semaphore count (i.e., releases a resource). This may unblock one or more threads waiting for the semaphore.

Example of Using `std::sem`

Here is a simple example demonstrating how to use a `std::sem` to manage access to a shared resource by multiple threads:

```
#include <iostream>
#include <semaphore>
#include <thread>
#include <vector>

std::binary_semaphore sem(3); // Semaphore with 3 resources

void access_shared_resource(int thread_id) {
```

```
sem.acquire(); // Wait for an available resource

std::cout << "Thread " << thread_id << " is accessing the resource."
↳ << std::endl;

// Simulate resource usage
std::this_thread::sleep_for(std::chrono::seconds(1));

std::cout << "Thread " << thread_id << " has finished using the
↳ resource." << std::endl;

sem.release(); // Release the resource
}

int main() {
    const int num_threads = 5;
    std::vector<std::thread> threads;

    // Create threads
    for (int i = 0; i < num_threads; ++i) {
        threads.emplace_back(access_shared_resource, i);
    }

    // Join all threads
    for (auto& t : threads) {
        t.join();
    }

    return 0;
}
```

Output (execution order might vary):

```
Thread 0 is accessing the resource.  
Thread 1 is accessing the resource.  
Thread 2 is accessing the resource.  
Thread 0 has finished using the resource.  
Thread 3 is accessing the resource.  
Thread 1 has finished using the resource.  
Thread 2 has finished using the resource.  
Thread 4 is accessing the resource.  
Thread 3 has finished using the resource.  
Thread 4 has finished using the resource.
```

In this example:

- We create a **`std::binary_semaphore`** with an initial count of 3, allowing three threads to access the shared resource concurrently.
- Each thread waits for an available permit, uses the resource, and then releases the permit.

Types of Semaphores

`std::sem` provides two primary types of semaphores:

1. **`std::binary_semaphore`**: A binary semaphore where the counter is either 0 or 1. It is used for mutual exclusion, where only one thread can access a resource at a time.
2. **`std::counting_semaphore`**: A counting semaphore where the counter can be greater than 1. It is useful when there are multiple instances of a resource that threads can concurrently access.

Use Cases for `std::sem`

- **Managing Resource Limits:** When a system has a fixed number of resources (e.g., threads, connections, buffers), semaphores allow threads to acquire a resource, work with it, and then release it.
- **Rate Limiting:** Semaphores can control how often a thread can perform a certain operation within a time period.
- **Worker Pools:** Semaphores can be used to limit the number of worker threads operating concurrently on tasks.

6.2.2 Introduction to `std::latch`

While semaphores are used to control access to resources and limit concurrency, `std::latch` is a synchronization primitive used to block one or more threads until a certain number of events have occurred. It is designed for situations where threads need to wait for other threads to reach a particular point in execution, often referred to as a "barrier".

The `std::latch` class template was introduced in **C++23** to make it easier to synchronize threads based on a fixed number of signals.

The Structure of `std::latch`

A `std::latch` is initialized with a specified count, which represents the number of threads or events that need to occur before the latch "opens." Once the count reaches zero, the waiting threads are unblocked. Threads can decrement the latch counter using the `count_down()` method, and waiting threads can block until the count reaches zero using the `wait()` method.

Example of Using `std::latch`

Here is an example demonstrating how to use a `std::latch` to synchronize threads:

```
#include <iostream>
#include <latch>
```

```
#include <thread>
#include <vector>

std::latch latch(3); // Latch that waits for 3 threads

void thread_task(int thread_id) {
    std::cout << "Thread " << thread_id << " is starting." << std::endl;

    // Simulate some work
    std::this_thread::sleep_for(std::chrono::seconds(1));

    std::cout << "Thread " << thread_id << " is done." << std::endl;

    latch.count_down(); // Signal that this thread is done
}

int main() {
    std::vector<std::thread> threads;

    // Create and start 3 threads
    for (int i = 0; i < 3; ++i) {
        threads.emplace_back(thread_task, i);
    }

    // Wait for all threads to reach the count_down() call
    latch.wait();
    std::cout << "All threads have finished. Continuing main thread." <<
        ↵ std::endl;

    // Join threads
    for (auto& t : threads) {
        t.join();
    }
}
```

```
    }  
  
    return 0;  
}
```

Output (execution order might vary):

```
Thread 0 is starting.  
Thread 1 is starting.  
Thread 2 is starting.  
Thread 0 is done.  
Thread 1 is done.  
Thread 2 is done.  
All threads have finished. Continuing main thread.
```

In this example:

- The **`std::latch`** is initialized with a count of 3, indicating that three threads need to call **`count_down()`** before the main thread can proceed.
- Each worker thread simulates some work and then calls **`count_down()`** to signal its completion.
- The main thread waits for all worker threads to finish by calling **`wait()`**, at which point the latch counter reaches zero, and the main thread can continue.

Use Cases for **`std::latch`**

- **Thread Synchronization:** A common use case for **`std::latch`** is synchronizing a set of worker threads before starting a collective task.

- **Barrier Synchronization:** Ensuring that all threads reach a certain point of execution before proceeding to the next phase.
- **Coordinating Task Completion:** When you need to ensure that multiple threads have completed their work before performing further operations.

6.2.3 Comparing `std::sem` and `std::latch`

While both `std::sem` and `std::latch` are synchronization primitives, they are designed to solve different problems:

- **`std::sem` (Semaphore):** Useful for controlling access to resources (like limiting concurrent access). It is a general-purpose primitive that can be used to allow a fixed number of threads to work concurrently.
- **`std::latch`:** Useful for synchronization between threads, ensuring that all threads or events are complete before proceeding. It does not allow resetting and is typically used to block threads until all required threads have reached a certain point.

Conclusion

The introduction of `std::sem` and `std::latch` in C++23 enhances the concurrency capabilities of the C++ language, providing developers with more flexible and efficient tools for synchronizing threads. `std::sem` offers control over resource access, while `std::latch` enables synchronization between threads, ensuring that specific events occur before continuing execution. These new primitives contribute to more readable, maintainable, and high-performance multi-threaded programs. As multi-threaded applications continue to grow in complexity, the inclusion of these tools in the standard library marks a significant step forward in making C++ a more powerful language for concurrent programming.

6.3 New Tools to Support Parallel Processing

Parallel processing has become a cornerstone for performance optimization in modern software development. As hardware evolves with more cores and processing units, the need for efficient parallel execution grows. **C++23** continues to improve its capabilities for managing parallelism, providing new tools to help developers take full advantage of multi-core systems. In this section, we will explore the latest additions and updates in the C++23 Standard Library that enhance support for parallel processing, making it easier and more efficient to implement concurrent and parallel algorithms.

6.3.1 Introduction to Parallelism in C++23

Parallel processing allows tasks to be executed simultaneously across multiple processing units (such as CPU cores), reducing the time required to process large datasets, perform computations, or handle multiple tasks concurrently. However, parallel programming is complex and requires careful management of thread synchronization, load balancing, and task decomposition to avoid issues such as data races, deadlocks, and inefficient CPU usage.

In **C++23**, several new tools have been added to simplify and enhance parallel processing, making it more accessible to developers while improving performance. These tools focus on fine-grained control over parallel execution, better integration with existing algorithms, and improved support for multi-core processors.

6.3.2 New Parallel Algorithms in C++23

Parallel algorithms have been a part of the C++ Standard Library since **C++17**, where the execution policies (`std::execution::par`, `std::execution::par_unseq`, etc.) were introduced to indicate how algorithms should be executed. **C++23** introduces new parallel algorithm features and improvements, making it easier to apply parallelism to a broader set of

use cases. These improvements extend the versatility and efficiency of parallel operations on data structures like vectors, arrays, and other containers.

`std::transform_reduce` and `std::transform_exclusive_scan`

Two powerful new algorithms introduced in **C++23** are **`std::transform_reduce`** and **`std::transform_exclusive_scan`**. These algorithms provide more flexible ways to apply parallelism to common data reduction and scanning operations.

- **`std::transform_reduce`**: This is an algorithm that combines both transformation and reduction in a single step. It allows you to apply a function to each element of a range, followed by a reduction (e.g., summing, multiplying) of the results. The main advantage is that it can be executed in parallel efficiently, without the need to manually split and combine results.

Example:

```
#include <iostream>
#include <vector>
#include <numeric>
#include <execution>

int main() {
    std::vector<int> data = {1, 2, 3, 4, 5};

    // Perform a parallel transform-reduce operation
    int sum_of_squares = std::transform_reduce(
        std::execution::par,           // Parallel execution
        ↪ policy
        data.begin(), data.end(),      // Input range
        0,                             // Initial value
        std::plus<int>(),              // Reduction function
        ↪ (sum)
```

```

    [](int x) { return x * x; } // Transformation function
    ↪ (square)

);

std::cout << "Sum of squares: " << sum_of_squares << std::endl;
return 0;
}

```

- **std::transform_exclusive_scan:** This algorithm applies a transformation to each element of a range, followed by an exclusive scan (or prefix sum), which computes the cumulative sum up to each element. The benefit of this algorithm is its ability to run in parallel while maintaining the sequential semantics of exclusive scans.

Example:

```
#include <iostream>
#include <vector>
#include <numeric>
#include <execution>

int main() {
    std::vector<int> data = {1, 2, 3, 4, 5};
    std::vector<int> result(data.size());

    // Parallel exclusive scan with transformation
    std::transform_exclusive_scan(
        std::execution::par,           // Parallel execution
        ↪ policy
        data.begin(), data.end(),      // Input range
        result.begin(),               // Output range
        0,                            // Initial value
```

```

        std::plus<int>(), // Scan function (sum)
        [](int x) { return x * x; } // Transformation function
        ↪ (square)
    );

    for (const auto& val : result) {
        std::cout << val << " ";
    }
    std::cout << std::endl;
    return 0;
}

```

std::for_each with Parallel Execution

The **std::for_each** algorithm in **C++23** allows for parallel execution when provided with an appropriate execution policy. This allows you to apply a function to each element in a range concurrently. For example, you can use **std::execution::par** to apply a function to a collection of data in parallel, speeding up the computation for large datasets.

Example:

```

#include <iostream>
#include <vector>
#include <execution>

int main() {
    std::vector<int> data = {1, 2, 3, 4, 5};

    // Apply a function to each element in parallel
    std::for_each(std::execution::par, data.begin(), data.end(), [](int& x)
    ↪ {
        x *= 2; // Doubling each element in the vector
    }

```

```
});  
  
for (const auto& val : data) {  
    std::cout << val << " ";  
}  
std::cout << std::endl;  
return 0;  
}
```

6.3.3 Support for `std::execution::parallel_policy`

In **C++23**, the parallel execution policies have been further refined to allow for better control over parallelism, particularly in combination with the new parallel algorithms. The **`std::execution::par`** policy has been improved to ensure that parallel execution can scale better with the number of available hardware threads.

- **`std::execution::par_unseq`** (parallel and unsequenced execution) allows algorithms to be executed in parallel with the possibility of vectorization, where operations on adjacent elements are processed simultaneously using SIMD (Single Instruction, Multiple Data) instructions.
- **`std::execution::par`** (parallel execution) is used when parallel execution is desired, but vectorization isn't guaranteed.

These execution policies are particularly useful in scenarios where you need to maximize performance on multi-core processors without having to manage the complexity of threads manually.

6.3.4 Improved Support for Parallel Sorting

Sorting is one of the most common and performance-critical operations in many applications.

C++23 continues to enhance parallel sorting algorithms. One key update is the

`std::ranges::sort` algorithm, which has been made compatible with parallel execution policies, enabling more efficient sorting for large datasets.

By applying **`std::execution::par`** to the **`std::ranges::sort`** algorithm, C++23 allows developers to parallelize sorting operations with minimal code changes, making it easier to leverage multi-core processors for large collections.

Example:

```
#include <iostream>
#include <vector>
#include <ranges>
#include <execution>

int main() {
    std::vector<int> data = {5, 2, 9, 1, 3, 8};

    // Parallel sorting
    std::ranges::sort(std::execution::par, data);

    for (const auto& val : data) {
        std::cout << val << " ";
    }
    std::cout << std::endl;

    return 0;
}
```

6.3.5 Parallelism with `std::async` and `std::future`

While `std::async` and `std::future` were introduced in C++11 to simplify asynchronous programming, C++23 brings improvements in how they interact with the execution policies and thread management. This makes it easier to combine parallel algorithms with asynchronous tasks, ensuring that you can achieve better load balancing and task coordination.

For example, you can use `std::async` to launch parallel tasks and `std::future` to collect their results, enabling a parallel flow of computations across multiple threads.

Example:

```
#include <iostream>
#include <vector>
#include <future>
#include <numeric>

int compute_sum(const std::vector<int>& data) {
    return std::accumulate(data.begin(), data.end(), 0);
}

int main() {
    std::vector<int> data = {1, 2, 3, 4, 5};

    // Parallel asynchronous task
    std::future<int> result = std::async(std::launch::async, compute_sum,
    ↪ std::cref(data));

    // Retrieve the result
    std::cout << "Sum: " << result.get() << std::endl;

    return 0;
}
```

Conclusion

The **C++23** standard significantly enhances support for parallel processing, making it easier and more efficient to take advantage of modern multi-core CPUs. The introduction of new algorithms such as **std::transform_reduce** and

std::transform_exclusive_scan, combined with enhanced execution policies and parallel sorting, gives developers more power to optimize their applications. The ability to seamlessly incorporate parallelism into existing code using simple execution policies—such as **std::execution::par**—increases both performance and maintainability.

These new tools will help developers optimize computational-heavy tasks, improve the performance of data processing pipelines, and make multi-core parallelism accessible without complex manual thread management. By leveraging these features, developers can write high-performance, scalable applications that meet the demands of modern hardware and software environments.

Chapter 7

Compatibility with Previous C++ Versions

7.1 How to Use Updates in Multi-Version Environments

One of the ongoing challenges when working with C++ is ensuring that your code remains compatible across multiple versions of the language, especially as new features and updates are introduced in newer standards such as **C++17**, **C++20**, and **C++23**. C++ developers often face the dilemma of maintaining backward compatibility with older compilers and toolchains while taking advantage of the latest features for performance improvements, better readability, and enhanced functionality.

In this section, we will focus on how you can effectively use the updates introduced in **C++23** while maintaining compatibility with previous versions of C++, such as **C++17** and **C++20**, in multi-version environments. This is essential for organizations with legacy codebases or when working in environments where upgrading to the latest standard is not feasible immediately.

7.1.1 The Importance of Backward Compatibility

C++ is a complex language with numerous features, and each new version introduces both improvements and new functionalities. However, not all developers are able to upgrade their environments to the latest standard right away. Some projects require maintaining compatibility with legacy systems or external libraries that have not yet been updated to newer versions of the language.

Additionally, different compilers may support different versions of C++ or may only provide partial support for certain features in newer versions. For example, some compilers might implement C++17 features without fully supporting C++20 or C++23. In these scenarios, being able to selectively enable or disable newer language features is crucial for maintaining the portability and functionality of the code.

7.1.2 Strategies for Using C++23 Updates in Multi-Version Environments

When working in environments where you need to support multiple versions of C++, you can follow several strategies to use **C++23** features while keeping your code compatible with previous C++ versions like **C++17** or **C++20**. These strategies include using conditional compilation, compiler flags, and feature tests to ensure that newer features are only used in compatible environments.

Conditional Compilation

Conditional compilation allows you to selectively compile sections of code based on the version of the compiler or the C++ standard it supports. This is one of the most straightforward ways to ensure that your code works across different versions of C++.

You can use the `__cplusplus` macro, which is defined by the compiler, to check the version of C++ being used. This macro holds the value corresponding to the standard version, such as `201703L` for C++17, `202002L` for C++20, and `202302L` for C++23.

Example:

```
#include <iostream>

int main() {
    // Check for C++17 or newer
    #if __cplusplus >= 201703L
        std::cout << "C++17 or later features available." << std::endl;
    #endif

    // Check for C++20 or newer
    #if __cplusplus >= 202002L
        std::cout << "C++20 or later features available." << std::endl;
    #endif

    // Check for C++23 or newer
    #if __cplusplus >= 202302L
        std::cout << "C++23 features available!" << std::endl;
    #else
        std::cout << "Older version of C++ detected." << std::endl;
    #endif

    return 0;
}
```

In this example:

- The `__cplusplus` macro is used to check if the C++ standard version is greater than or equal to C++17, C++20, or C++23.
- This allows you to use **C++23**-specific features only when the compiler supports it, while still maintaining backward compatibility with older versions of C++.

Compiler Flags and Preprocessor Directives

When working with a multi-version environment, it's also important to configure your compiler settings to target specific versions of C++ through compiler flags. For example, when using **GCC** or **Clang**, you can specify the C++ standard to use via the `-std=c++XX` flag, where `XX` is the version number (e.g., `c++17`, `c++20`, `c++23`). Similarly, with **MSVC (Microsoft Visual C++)**, you can configure the C++ version through project settings or the `/std:c++XX` option.

You can combine these flags with preprocessor directives in your source code to enable or disable specific features based on the compiler version being used.

Example using GCC/Clang compiler flag:

```
# Compile for C++17
g++ -std=c++17 my_program.cpp -o my_program

# Compile for C++20
g++ -std=c++20 my_program.cpp -o my_program

# Compile for C++23
g++ -std=c++23 my_program.cpp -o my_program
```

By using compiler flags in combination with preprocessor directives, you can ensure that the correct code paths are taken depending on the targeted C++ version.

Feature Test Macros

In addition to the version-based conditional compilation, **C++23** introduces a number of feature test macros. These macros enable you to check for specific language features, making it easier to write code that is compatible across different versions of C++ while still using the latest features when available.

Feature test macros allow you to test for the presence of specific features such as **`std::span`**, **`std::ranges`**, or **`std::expected`** introduced in C++23. These macros allow your code to gracefully handle missing features or use fallback implementations if certain features are not

available in the compiler being used.

For example:

- `__cpp_lib_span` is a macro defined if the compiler supports `std::span`.
- `__cpp_lib_ranges` is a macro defined if the compiler supports **`std::ranges`**.

Example:

```
#include <iostream>

#ifdef __cpp_lib_span
    #include <span>
#endif

int main() {
    #ifdef __cpp_lib_span
        std::span<int> sp({1, 2, 3, 4});
        std::cout << "std::span is supported." << std::endl;
    #else
        std::cout << "std::span is not available." << std::endl;
    #endif

    return 0;
}
```

This code checks whether **`std::span`** is available before using it, ensuring compatibility with older compilers that do not yet support this feature.

7.1.3 Keeping Dependencies Compatible with Multiple Versions

When working with libraries or dependencies, it's essential to ensure that the libraries you use are compatible with different C++ versions. Many libraries and frameworks are gradually

updated to support newer versions of the language, but you may still need to work with older versions in certain cases.

Use of Conditional Includes

One way to manage dependencies is by conditionally including headers based on the C++ version being used. For example, libraries that support multiple versions of C++ may have separate headers for C++17, C++20, and C++23.

Example:

```
#if __cplusplus >= 202002L
    #include <some_library/cpp20_feature.hpp>
#else
    #include <some_library/legacy_feature.hpp>
#endif
```

Managing Compatibility in Build Systems

For large projects with multiple dependencies, it's important to configure your build system to handle different C++ versions properly. **CMake**, **Makefiles**, and other build systems can be configured to compile different modules of the project with different C++ standards.

Example CMake snippet:

```
if(CMAKE_CXX_STANDARD LESS 202002L)
    target_compile_definitions(my_target PRIVATE -DOLD_CPP_VERSION)
endif()
```

This enables the inclusion of version-specific code paths or libraries, making it easier to maintain compatibility with different versions of C++.

Conclusion

In multi-version C++ environments, the ability to selectively use new features in **C++23** while

maintaining compatibility with **C++17** or **C++20** is essential for ensuring portability and functionality across various compilers and toolchains. By leveraging conditional compilation, compiler flags, feature test macros, and careful dependency management, developers can write modern C++ code that is compatible with a wide range of environments and compilers. Using these techniques allows you to gradually adopt new features of C++23 without breaking compatibility with older versions, enabling smoother transitions and more maintainable code across versions. This approach is critical for large projects with legacy systems, ensuring that new optimizations and language features can be adopted at a pace that suits the development lifecycle and organizational needs.

7.2 Enhancements in Backward Compatibility

C++ is known for its commitment to backward compatibility, ensuring that code written for older versions of the language continues to work seamlessly with newer compilers and toolchains. This focus on backward compatibility is crucial in C++ because it allows long-running projects and large codebases to benefit from improvements in newer versions of the language without requiring a complete rewrite of existing code.

In **C++23**, significant enhancements have been made to ensure backward compatibility, particularly for legacy code that might rely on older features or compiler behaviors. These improvements provide smoother transitions from older versions, such as **C++17** and **C++20**, while introducing new features in **C++23**. This section will explore how C++23 maintains backward compatibility and what changes have been made to improve it for developers working in multi-version environments.

7.2.1 The Importance of Backward Compatibility in C++

Backward compatibility in C++ is vital for several reasons:

- **Legacy Systems:** Many systems in production today are based on older versions of C++

(e.g., C++98, C++03, C++11, C++14, or C++17). These systems often have large, complex codebases that cannot be easily refactored to use newer versions of the language.

- **Third-Party Libraries:** Developers frequently use libraries that may not be updated immediately to support the latest C++ standard. Ensuring backward compatibility with older versions of the language allows these libraries to continue working in newer compilers.
- **Gradual Adoption:** Some teams prefer to migrate incrementally to newer standards to avoid the risk of introducing bugs and breaking changes into a working system. Backward compatibility helps ensure that the transition process is smooth and that new language features can be adopted without breaking existing code.

By maintaining backward compatibility, C++ allows developers to upgrade their toolchains, take advantage of performance improvements and new features, and still maintain the stability of existing applications.

7.2.2 Enhancements in Backward Compatibility in C++23

While C++23 introduces several new features, enhancements, and changes to the language, the standard body has made significant efforts to ensure that older codebases remain functional with minimal adjustments. Below are the major backward compatibility enhancements introduced in C++23.

Removal of Deprecated Features

One of the key aspects of maintaining backward compatibility is ensuring that deprecated features are removed gradually. C++23 continues to clean up older, less useful features, but it does so with care to minimize disruption. In many cases, deprecated features are marked with warnings in earlier versions (such as C++17 or C++20), giving developers time to migrate their code before removal.

For example, `std::auto_ptr`, which was deprecated in **C++11**, is fully removed in **C++23**, but compilers will emit warnings about its usage long before this happens. This provides a transition period where developers can replace `std::auto_ptr` with the safer `std::unique_ptr` or `std::shared_ptr`, which are now standard practice for managing ownership of dynamically allocated memory.

This gradual deprecation and removal process ensures that legacy code continues to work for the most part, while providing developers with clear guidance and sufficient time to adopt new practices.

Compatibility with Older Compiler Implementations

As compiler implementations evolve, C++ standards need to maintain compatibility with older versions of the language. C++23 ensures that older features, such as the **C++98** standard, can still be compiled by modern compilers. This is important for applications that still depend on legacy toolchains or for specific platforms that do not support the latest versions of the language. Modern compilers like GCC, Clang, and MSVC provide extensive backward compatibility by supporting flags to specify the desired C++ standard. By maintaining this compatibility, C++23 allows users to choose which version of the language they want to work with, depending on the project's needs.

For example, compilers typically support the following flags:

- **GCC/Clang:** `-std=c++98`, `-std=c++11`, `-std=c++14`, `-std=c++17`,
`-std=c++20`, `-std=c++23`
- **MSVC:** `/std:c++98`, `/std:c++11`, `/std:c++14`, `/std:c++17`, `/std:c++20`,
`/std:c++23`

These flags let developers compile code with older versions of C++ while still having the option to leverage new features when required, ensuring compatibility across multiple versions of C++.

Improved Deprecation Warnings for Legacy Features

C++23 continues the trend set by previous standards of providing **deprecation warnings** for features that are likely to be removed in future versions of the language. These warnings inform developers of code that might become incompatible with future versions of the language, giving them ample time to update their code before those features are eventually removed.

For example, features like `std::auto_ptr` or certain overloaded functions might trigger deprecation warnings when compiled with a modern compiler in C++23, even if the code was written in an older version of the language. These warnings help maintain compatibility between versions by alerting developers to potential problems well in advance of removal.

Library Backward Compatibility

In **C++23**, the Standard Library has also made efforts to ensure that existing library code continues to work seamlessly with new versions. For example:

- **Iterators** and **ranges** have been extended in C++23, but legacy iterators and algorithms (from C++17 and C++20) still work, ensuring that older code that uses conventional iterators and algorithms does not break.
- `std::span`, introduced in **C++20**, is fully compatible with existing array-based code in **C++23**, ensuring that legacy array management code will continue to function without modification.
- The `std::expected` type, introduced in **C++23**, has been designed to integrate smoothly with existing exception-handling patterns in C++ code, allowing for optional error handling mechanisms without breaking existing code.

This ensures that when developers update their toolchains to C++23, they don't have to go back and rework significant portions of their codebase simply to accommodate changes in the library.

Compiler and Standard Library Evolution

C++23 emphasizes the importance of evolving compilers and libraries in such a way that code written for older versions of C++ (e.g., C++17 or C++20) can continue to compile and run with minor adjustments. This is achieved by:

- **Backward-compatible additions:** New language features and library improvements in C++23 have been designed to extend existing functionality without breaking the behavior of older code. For example, `std::ranges` and `std::span`, although introduced in newer standards, maintain compatibility with existing iterators, arrays, and other legacy constructs.
- **Backporting features:** Many of the features in C++23, like certain algorithms and optimizations, can be backported to older versions of C++ through external libraries. This means that developers using older compilers can still take advantage of some of the performance improvements and new functionality by incorporating libraries like **Boost** or specialized C++23 feature libraries.

By ensuring that compilers and libraries evolve in a backward-compatible manner, C++23 improves the ease with which developers can transition their projects to the latest standards.

7.2.3 Techniques for Maintaining Backward Compatibility in C++23

To take full advantage of C++23's backward compatibility features, developers can employ the following techniques:

Use of Feature Test Macros

Feature test macros in C++ allow developers to detect the availability of specific features (e.g., new library components or language features). This enables developers to write code that adapts to different versions of the language by using the features available in the target environment. For example, the following feature test macros check for the presence of new **C++23** features:

- `__cpp_lib_expected`: Checks if `std::expected` is available.

- `__cpp_lib_span`: Checks if `std::span` is available.
- `__cpp_lib_ranges`: Checks if `std::ranges` is available.

These macros allow developers to use newer features while maintaining compatibility with older compilers that do not support them.

Conditional Compilation with `__cplusplus`

The `__cplusplus` macro is one of the most useful tools for ensuring backward compatibility in multi-version environments. It allows code to be conditionally compiled depending on the version of the C++ standard supported by the compiler. By checking the value of `__cplusplus`, developers can ensure that code written for newer versions is only compiled if the target environment supports the appropriate C++ version.

```
#if __cplusplus >= 202002L
    // C++20 or later specific code
#else
    // Fallback to older standard compatibility
#endif
```

Testing with Multiple Compilers and Versions

To ensure that code works across multiple versions of C++, developers should test their code with different compilers and standard versions. Many continuous integration (CI) systems allow for testing across different environments. By setting up CI pipelines that compile and run tests across multiple C++ standards (C++17, C++20, C++23), developers can ensure their code remains compatible and functional with different versions.

Conclusion

C++23 provides significant enhancements to backward compatibility, ensuring that existing C++ codebases continue to function across multiple versions of the language. These improvements

benefit developers working with legacy systems, libraries, and compilers, and they allow for a smoother transition to the newer language features introduced in C++23.

By leveraging techniques like conditional compilation, feature test macros, and multi-version testing, developers can ensure their code remains portable and maintainable, whether working with legacy code or taking advantage of the latest features. This focus on backward compatibility in C++23 allows developers to continue building on their existing investments while adopting new standards and practices incrementally.

Chapter 8

Best Practices for Using STL

8.1 Tips to Improve Program Performance Using STL

The Standard Template Library (STL) is an essential part of C++ programming, providing a set of efficient, reusable, and well-tested data structures and algorithms. However, to fully realize its performance potential, developers must make informed decisions about how they use STL containers, algorithms, and iterators. Even though STL provides high-quality implementations of common data structures and algorithms, the performance of a program can still vary based on how effectively these tools are used.

This section provides a detailed exploration of tips and best practices to optimize program performance when using the STL in C++23. By understanding how STL components interact with each other and how the choice of data structures, algorithms, and their operations can impact performance, developers can create more efficient, scalable, and maintainable code.

8.1.1 Understanding the Computational Complexity of STL Containers

A fundamental step in optimizing performance with STL is understanding the time complexities of the operations performed on different containers. Each container in the STL has its own set of strengths and weaknesses based on the type of data and the operations you need to perform.

Choosing the Right Container

Different STL containers offer different time complexities for common operations like insertion, deletion, and access. Below is a breakdown of common STL containers and their associated time complexities for some of the most important operations:

- **`std::vector`:**
 - Best for random access and sequential operations.
 - Insertion at the **end** is $O(1)$, but insertion in the **middle** or **beginning** is $O(n)$.
 - Access to an element is $O(1)$.
 - Resizing can be $O(n)$ if reallocation occurs.
- **`std::deque`:**
 - Fast insertions and deletions at both the **front** and the **back**.
 - Access to an element is $O(1)$, but random access is slower than `std::vector` due to internal data structures.
- **`std::list`:**
 - Insertion and deletion are $O(1)$ when the position is known, but access is $O(n)$.
 - Ideal for scenarios where you frequently insert and remove elements but do not require fast random access.

- **`std::unordered_map`** and **`std::unordered_set`**:
 - These containers use hash tables and offer average constant time $O(1)$ for insertion, deletion, and search.
 - Worst-case time complexity is $O(n)$, which happens when all elements hash to the same bucket, though this is rare.
- **`std::map`** and **`std::set`**:
 - These containers are implemented as balanced binary search trees, with $O(\log n)$ time complexity for insertion, deletion, and search.

Minimizing Reallocation with **`std::vector`**

A **`std::vector`** dynamically resizes as elements are added, but reallocating memory is expensive because it may involve copying all elements to a new memory location. To improve performance, you can avoid unnecessary reallocations by **reserving capacity** ahead of time when you know the number of elements you will store in the vector.

```
std::vector<int> vec;  
vec.reserve(1000); // Avoid reallocating during insertions up to 1000  
↪ elements
```

This reduces the cost of multiple reallocations as the vector grows and ensures that the memory management overhead is minimized.

Avoiding Excessive Memory Usage with **`std::list`**

A **`std::list`** uses more memory per element than a **`std::vector`** or **`std::deque`** because each element requires additional memory for the pointers to the previous and next elements. In scenarios where random access and fast iteration are not crucial, **`std::deque`** or **`std::vector`** may offer better performance and memory efficiency than **`std::list`**.

8.1.2 Leveraging STL Algorithms for Optimal Performance

STL algorithms, such as `std::sort`, `std::find`, and `std::accumulate`, are implemented efficiently, making them preferable over manually writing loops for common operations. By using the STL algorithms, you benefit from highly optimized, often parallelized, and well-tested implementations. However, the performance gains come from knowing when and how to apply these algorithms.

Using `std::sort` Instead of `std::stable_sort`

`std::sort` is generally more efficient than `std::stable_sort`, which guarantees the relative order of equal elements. If the stability of sorting is not required, using `std::sort` avoids the overhead of maintaining order for equal elements and provides better performance in practice.

```
std::vector<int> vec = {5, 1, 3, 4, 2};  
std::sort(vec.begin(), vec.end()); // Faster, unless stability is needed
```

Prefer `std::find_if` and `std::count_if` Over Manual Loops

The `std::find_if` and `std::count_if` algorithms are optimized for searching and counting elements that meet specific criteria. Using these algorithms ensures that your code takes advantage of internal optimizations, such as early exits from loops when a match is found or an element meets the required condition.

```
// Example using std::find_if  
auto result = std::find_if(vec.begin(), vec.end(), [](int x) { return x >  
    < 10; });
```

This approach is typically more efficient and concise than writing a manual loop to achieve the same behavior.

Using `std::for_each` to Apply Functions in Bulk

When performing an operation across all elements in a container, `std::for_each` can be an efficient way to apply a function to every element. It can be especially useful when used in combination with lambda expressions to simplify your code and improve performance.

```
std::vector<int> vec = {1, 2, 3, 4, 5};  
std::for_each(vec.begin(), vec.end(), [](int& x) { x *= 2; });
```

This avoids manually writing a loop and ensures that the operation is optimized by the underlying implementation of `std::for_each`.

8.1.3 Efficient Memory Usage with STL Containers

Efficient memory usage is crucial for high-performance applications, especially when handling large datasets. STL containers, such as `std::vector` and `std::unordered_map`, can be optimized by controlling their capacity, avoiding unnecessary copying, and selecting the right container type based on the use case.

Avoiding Unnecessary Copies

Passing containers by reference, or even by pointer in certain cases, is a best practice when working with STL. Copying large containers can be expensive, particularly when they hold large amounts of data. Always prefer **pass-by-reference** or **move semantics** to avoid copying the container when possible.

```
void process_data(std::vector<int>& data) {  
    // Pass-by-reference avoids copying  
}  
  
std::vector<int> data = {1, 2, 3};  
process_data(data);
```

Additionally, you can use **std::move** to transfer ownership of a container's data instead of copying it, which can greatly improve performance in certain cases.

```
std::vector<int> data1 = {1, 2, 3};  
std::vector<int> data2 = std::move(data1); // Avoid copy, transfer  
→ ownership
```

Avoiding Contiguous Memory Copying with **std::deque**

While **std::vector** provides fast random access and efficient memory usage for contiguous blocks of memory, **std::deque** is more efficient when insertions and deletions occur at both ends of the container. However, it may be less efficient for large data copying because its internal structure is not a single contiguous block of memory.

If your program requires a mix of both fast insertions and random access, **std::deque** might be the better choice than **std::vector**.

Memory Pools and Allocators

For scenarios where memory allocation overhead is a bottleneck (e.g., frequent allocations and deallocations), using custom **allocators** or memory **pools** can help reduce the costs associated with allocating memory dynamically. STL containers allow the use of custom allocators, which can be implemented to pool memory or use specialized allocation strategies.

```
std::vector<int, MyAllocator<int>> vec; // Using a custom allocator
```

Allocators can help reduce fragmentation and improve memory access patterns, especially when working with large-scale systems or real-time applications.

8.1.4 Optimizing for Parallelism in C++23

C++23 has introduced several updates to STL, specifically in the realm of parallelism. Many STL algorithms now support parallel execution, allowing them to leverage multi-core processors

for faster execution.

Parallel Algorithms

The STL in C++23 introduces parallel versions of commonly used algorithms (e.g., `std::for_each`, `std::sort`, and `std::transform`) that can automatically utilize multiple cores for better performance.

You can enable parallel execution by passing the `std::execution::par` execution policy to these algorithms:

```
std::vector<int> vec = {1, 2, 3, 4, 5};
std::for_each(std::execution::par, vec.begin(), vec.end(), [](int& x) { x
↪  *= 2; });
```

Parallelizing algorithms in this way helps achieve significant performance improvements in data-heavy applications, especially when dealing with large datasets.

`std::ranges` and Parallel Execution

With the introduction of `std::ranges` in C++20 and its enhanced capabilities in C++23, developers can perform operations on ranges of data while also enabling parallel execution. By combining the power of ranges and parallelism, you can simplify your code and improve its efficiency:

```
#include <ranges>
#include <execution>

std::vector<int> vec = {1, 2, 3, 4, 5};
auto result = vec | std::ranges::view::transform([](int x) { return x * 2;
↪  })
                | std::ranges::to<std::vector>();
```

By using `std::ranges`, developers gain a more declarative approach to parallelism and

optimization.

Conclusion

Improving program performance with STL requires a deep understanding of the time complexities, container types, algorithms, and memory management techniques available. By carefully selecting the right STL containers, algorithms, and leveraging parallelism, developers can optimize their applications for speed and efficiency.

Key strategies include:

- Choosing the right container for your data access patterns.
- Avoiding unnecessary copies of containers and utilizing move semantics.
- Leveraging optimized STL algorithms and parallel execution policies in C++23.
- Efficient memory management with custom allocators or memory pools.

By following these best practices, you can unlock the full potential of STL in C++23, creating high-performance, maintainable, and scalable software systems.

8.2 Avoiding Common Mistakes with Containers and Algorithms

In this section, we will discuss common pitfalls that developers encounter when working with STL containers and algorithms in C++ and how to avoid them. While STL is designed to be efficient and flexible, it is essential to understand the nuances of container behavior, algorithm usage, and memory management to prevent inefficiencies, bugs, or performance problems. This section focuses on the most frequent mistakes developers make when using STL, as well as the best practices for steering clear of these issues.

8.2.1 Mistakes with Container Choice

One of the most critical decisions in optimizing performance is choosing the right container for the job. Picking a container that does not fit your use case can lead to unnecessary overhead and slow performance.

Using `std::vector` when `std::list` or `std::deque` is More Suitable

`std::vector` is often the default container choice due to its simplicity and performance for random access. However, for scenarios where frequent insertions and deletions occur at the **front** or **middle** of the container, `std::list` or `std::deque` may be more appropriate. Misusing `std::vector` for such cases can lead to inefficient algorithms.

- **Mistake:** Using `std::vector` for a container that requires frequent insertions/deletions at both ends or the middle.
- **Fix:** Use `std::deque` for fast insertions and deletions at both ends, or use `std::list` for constant-time insertions/deletions at any point.

```
std::deque<int> dq; // Best for frequent insertions and deletions at both
↪ ends
dq.push_front(1); // O(1) operation for both front and back insertions
```

In contrast, operations on `std::vector`—like inserting or removing elements in the middle or front—are **O(n)** due to the need for shifting elements.

Using `std::map` or `std::set` When Hash Maps Are a Better Fit

If you need a fast lookup table with **unique keys** and don't need to maintain sorted order, `std::unordered_map` or `std::unordered_set` (which use hash tables) may provide a better solution than `std::map` or `std::set` (which are typically implemented as balanced trees).

- **Mistake:** Using `std::map` when an unordered container would suffice, leading to unnecessary logarithmic time complexity.
- **Fix:** Opt for `std::unordered_map` or `std::unordered_set` for faster $O(1)$ average-time lookups.

```
std::unordered_map<int, std::string> umap;  
umap[10] = "apple"; // Fast O(1) lookup and insertion
```

`std::map` and `std::set`, while sorted, incur an $O(\log n)$ time complexity for operations such as insertion, deletion, and lookup.

Choosing the Right Container for Memory Efficiency

Containers like `std::list` store additional pointers (to the previous and next elements), which can lead to overhead, especially when the data size is large. On the other hand, `std::vector` is often the most memory-efficient option because it uses a contiguous block of memory.

- **Mistake:** Using `std::list` for small, simple datasets when `std::vector` would be more memory-efficient.
- **Fix:** Use `std::vector` for scenarios where fast random access and lower memory overhead are crucial.

8.2.2 Mistakes with Algorithm Usage

STL algorithms are highly optimized, but developers often misuse them or fail to leverage their full power, resulting in inefficient or incorrect code.

Not Using the Right Algorithms for the Task

One of the most common mistakes in using STL is trying to implement logic that can be easily solved using STL algorithms, resulting in redundant or inefficient code. Algorithms like `std::find`, `std::transform`, and `std::sort` are highly optimized and should be used over custom loops or manual implementations of common tasks.

- **Mistake:** Writing a custom loop to find an element when `std::find` can achieve the same result more efficiently.
- **Fix:** Use `std::find` or other STL algorithms to simplify and optimize your code.

```
auto it = std::find(vec.begin(), vec.end(), target); // Efficient O(n)
↪ lookup
if (it != vec.end()) {
    // Element found, process it
}
```

Similarly, **`std::sort`** should be used over writing your sorting logic unless you have specific requirements that cannot be addressed by the standard sorting algorithm.

Misusing **`std::copy`** and **`std::move`**

One of the more subtle mistakes that can affect both performance and correctness is misunderstanding the difference between copying and moving elements. When moving objects, **`std::move`** transfers ownership without copying, while **`std::copy`** performs a deep copy.

- **Mistake:** Using `std::copy` instead of `std::move` when transferring ownership of objects.
- **Fix:** When you no longer need the original object and want to transfer ownership, use `std::move`.

```
std::vector<int> original = {1, 2, 3};  
std::vector<int> moved = std::move(original); // Efficient move operation
```

If you're dealing with large objects, moving is much more efficient than copying, as it avoids copying data and simply transfers the pointer.

Forgetting to Account for Iterators in Algorithms

When using algorithms like `std::sort`, `std::reverse`, or `std::for_each`, you often pass iterators to represent the range of data you want to work with. A common mistake is to forget to pass the correct iterators or to incorrectly modify the container while iterating, which can lead to undefined behavior.

- **Mistake:** Modifying a container while using iterators on it, which invalidates the iterators.
- **Fix:** Ensure that the container is not modified while it is being iterated over.

```
std::vector<int> vec = {1, 2, 3, 4};  
std::for_each(vec.begin(), vec.end(), [](int& x) { x *= 2; });
```

In this example, iterating and modifying the vector is safe. However, if you add or remove elements from `vec` during iteration, iterators might become invalid, leading to undefined behavior.

8.2.3 Mistakes with Memory Management and Efficiency

Memory management is another area where mistakes can degrade performance and cause resource leaks.

Memory Reallocation with `std::vector`

As previously mentioned, `std::vector` grows dynamically, and reallocating its memory can become a performance bottleneck. Reallocations occur when the vector's capacity is exceeded, which can lead to copying all of its elements into a new memory block. However, you can avoid these unnecessary reallocations by **reserving capacity** ahead of time when you know the size of the container.

- **Mistake:** Failing to call `reserve` when you know the number of elements that will be inserted.
- **Fix:** Use `reserve` to pre-allocate memory in `std::vector`.

```
std::vector<int> vec;  
vec.reserve(1000); // Preallocate memory for 1000 elements
```

This avoids the cost of reallocation when inserting elements into the vector.

Not Using `std::move` with Large Objects

When working with large objects or containers, it is essential to move objects rather than copy them. Using `std::move` with large objects or containers avoids the cost of copying data and can significantly improve performance.

- **Mistake:** Copying large objects when moving would be more efficient.
- **Fix:** Use `std::move` when you no longer need the original object.

```
std::vector<int> source = {1, 2, 3, 4};  
std::vector<int> destination = std::move(source); // Move ownership  
↪ instead of copying
```

This avoids the expensive operation of copying the entire vector and simply transfers ownership of the resources to the new vector.

8.2.4 Mistakes with Iterators

Iterators are powerful but can also introduce bugs if not used correctly.

Using Invalidated Iterators

Many STL algorithms require passing iterators to ranges of data. Modifying the container while iterating through it (e.g., adding/removing elements) invalidates iterators, leading to undefined behavior. This is a common mistake when using iterators.

- **Mistake:** Modifying a container while iterating through it.
- **Fix:** Avoid modifying the container while iterating over it or use different iteration methods when modifications are necessary.

```
std::vector<int> vec = {1, 2, 3, 4};
for (auto it = vec.begin(); it != vec.end(); ++it) {
    if (*it == 2) {
        vec.erase(it); // Incorrect: erasing while iterating invalidates
        ↪ the iterator
    }
}
```

To fix this, you can increment the iterator manually after the erase operation or use algorithms like `std::remove_if` to handle the removal safely.

```
vec.erase(std::remove(vec.begin(), vec.end(), 2), vec.end()); // Correct
↪ way
```

Conclusion

Avoiding common mistakes with STL containers and algorithms is essential for writing efficient, bug-free code. By understanding container types, the proper use of algorithms, iterator safety, and memory management, developers can avoid inefficiencies and errors.

The key takeaways for avoiding mistakes are:

- Choose the right container based on your needs (e.g., `std::vector` vs. `std::list`).
- Leverage STL algorithms instead of writing custom loops or logic.
- Use `std::move` for large objects to avoid unnecessary copying.
- Be cautious with iterators and ensure containers are not modified during iteration.
- Use `reserve` for `std::vector` to optimize memory usage and prevent unnecessary reallocations.

By following these best practices, developers can fully harness the power of STL, avoiding common pitfalls and maximizing performance and maintainability.

Chapter 9

Resources and Documentation

9.1 Links to Advanced Tutorials and Official STL Documentation

In this section, we will provide a comprehensive collection of resources, including advanced tutorials and official documentation, that will help you deepen your understanding of STL (Standard Template Library) in C++23. Whether you are a beginner or an experienced developer, these resources will help you explore more complex topics, stay up to date with the latest changes, and gain insight into best practices for using STL effectively.

9.1.1 Official Documentation

The **official C++ Standard Library documentation** is the most authoritative source for all things related to STL. It provides comprehensive details about all containers, algorithms, iterators, utilities, and more. For C++23, it is essential to access the most recent version of the documentation to ensure you are working with the newest features and enhancements.

cppreference.com

Cppreference is one of the most commonly used resources for accessing up-to-date, community-maintained documentation of the C++ Standard Library. It provides clear, concise, and exhaustive details on the various aspects of STL, including code examples and detailed descriptions of all functions, classes, and templates.

- **Link:** cppreference.com
- Highlights:
 - Comprehensive reference for STL containers (like `std::vector`, `std::map`, `std::unordered_map`).
 - Algorithms, iterators, and utilities.
 - C++23-specific changes and new features.
 - Examples of how to use templates, algorithms, and iterators.

ISO C++ Standard

The official ISO C++ standards document is the most authoritative resource for the C++ language, including STL. While it's often not as user-friendly as other sources, it is the final word on all specifications of the C++ language, and you can rely on it for detailed, precise information.

- **Link:** [ISO C++ Official Site](http://www.iso.org/iso/committee/iso/sc22/15418.html)
- Highlights:
 - Access to the full ISO C++ standard.
 - In-depth specifications for the language and library, including STL.
 - Changes in C++23 and related standard updates.

While ISO documentation is authoritative, it can sometimes be challenging to navigate, so for practical use cases, **cppreference.com** may be more accessible.

9.1.2 Advanced Tutorials

The following tutorials offer in-depth learning materials, examples, and best practices for using STL effectively in advanced programming scenarios. These resources will help you master new C++23 features and leverage STL to its fullest potential.

C++ Programming at GeeksforGeeks

GeeksforGeeks is a well-known platform for learning all levels of programming concepts, from beginner to advanced. It provides several detailed tutorials on C++ STL, covering both fundamental topics as well as complex concepts like iterators, custom containers, and algorithm optimizations.

- **Link:** [GeeksforGeeks C++ STL Tutorials](#)
- Highlights:
 - In-depth articles on containers like `std::vector`, `std::map`, `std::set`, etc.
 - Tutorials on more advanced STL topics such as **custom allocators**, **memory management**, and **algorithm optimizations**.
 - C++ STL challenges and quizzes to test your knowledge.
 - Detailed explanations and use cases for each STL container, iterator, and algorithm.

C++ STL Tutorials by Learncpp.com

Learncpp.com offers a series of well-structured tutorials that cover the fundamentals and more advanced features of C++ STL. Their content is perfect for developers transitioning from basic to advanced STL usage.

- **Link:** [Learncpp C++ STL](#)
- **Highlights:**
 - An easy-to-follow guide on STL containers and algorithms.
 - Advanced tutorials that cover **complex iterators**, **ranges**, and **new C++23 features**.
 - Deep dives into **performance optimization** and **efficiency** when using STL.
 - Insightful tips on choosing the right container for different use cases and ensuring correctness with algorithms.

C++ Standard Library by Nicolai M. Josuttis

Nicolai M. Josuttis' book, "The C++ Standard Library: A Tutorial and Reference," is considered a classic resource for C++ developers. It covers the STL in great detail, explaining the various components, their complexities, and providing advanced examples.

- **Link:** [The C++ Standard Library by Nicolai M. Josuttis \(Book\)](#)
- **Highlights:**
 - Advanced explanations of how STL works internally, including **allocators**, **iterators**, and **container design**.
 - Detailed discussion of **complexity analysis** and the **trade-offs** involved in using different STL components.
 - Practical tips on using STL in real-world applications.

Although this is a paid resource, it is one of the most trusted books for mastering the C++ Standard Library.

C++ Weekly YouTube Channel

For developers who prefer video tutorials, **C++ Weekly** by Jason Turner is a fantastic YouTube channel that focuses on best practices, advanced C++ techniques, and STL. It is an excellent resource for developers who want to stay current with C++23 and STL updates.

- **Link:** [C++ Weekly on YouTube](#)
- Highlights:
 - Short, digestible video lessons that cover both beginner and advanced topics.
 - Examples on efficient STL usage, new C++23 features, and optimization techniques.
 - Video explanations of complex concepts like **C++ template metaprogramming**, **ranges**, and **parallel algorithms** in STL.

This channel is a great way to keep your knowledge up-to-date with practical, real-world examples.

C++ STL Book by David R. Musser and Atul Saini

Another classic text, "**C++ STL: The Standard Template Library**" by David R. Musser and Atul Saini, provides detailed discussions and examples of STL features, from basic containers to more complex algorithms.

- **Link:** [C++ STL by Musser and Saini \(Book\)](#)
- Highlights:
 - Thorough coverage of STL containers and algorithms with a focus on understanding the **design principles** behind them.
 - Detailed sections on **algorithm optimization** and **complexity analysis**.
 - Provides a foundation for learning how to design efficient, maintainable C++ programs using STL.

9.1.3 Additional Resources

Apart from tutorials, there are other valuable resources to enhance your C++ STL knowledge:

Stack Overflow

When you're encountering specific problems or bugs related to STL, **Stack Overflow** is a go-to resource. Many C++ developers discuss and resolve issues related to STL containers, algorithms, and their performance characteristics.

- **Link:** [Stack Overflow - C++ STL](#)
- Highlights:
 - Large community of developers providing answers to a wide variety of STL-related questions.
 - Solutions to edge cases and performance pitfalls that may not be covered in traditional tutorials.

GitHub Repositories

Several GitHub repositories contain open-source projects and implementations of STL-related projects. Examining these can provide deeper insights into how STL is used in complex, real-world software.

- **Link:** [GitHub - C++ STL Projects](#)
- Highlights:
 - Explore codebases that use STL in production systems.
 - View examples of STL container optimizations and advanced usages.

Conclusion

By utilizing the resources outlined above, you can deepen your understanding of STL and C++23. Whether through official documentation, advanced tutorials, or community-driven platforms, these resources will help you become proficient in using STL for performance, readability, and maintainability in your applications.

Key resources to remember:

- **cppreference.com** for official documentation and quick reference.
- **GeeksforGeeks** and **Learncpp.com** for detailed tutorials.
- **YouTube channels** like C++ Weekly for practical video lessons.
- **Books** by Nicolai Josuttis and David Musser for deep dives into STL concepts.
- **Stack Overflow** and **GitHub** for community solutions and open-source examples.

Together, these resources will enable you to master STL in C++23 and apply it to real-world software development challenges.

9.2 Books and Training Courses for Further Learning

In this section, we will focus on **books** and **training courses** that provide deep insights into C++ and the Standard Template Library (STL), including specific resources that cover the **C++23** updates. Whether you prefer a traditional book format or prefer interactive online learning, these resources will help you solidify your knowledge, broaden your understanding, and learn advanced techniques to optimize your C++ programming using STL.

9.2.1 Books for Learning STL and C++

Books are an excellent resource for diving deep into the specifics of STL and its applications in C++. They offer structured, well-documented content with examples, explanations, and exercises. Below are some of the most influential and authoritative books on STL and C++ that will provide valuable learning paths for both beginners and advanced developers.

”The C++ Programming Language” by Bjarne Stroustrup

Bjarne Stroustrup, the creator of C++, provides the definitive guide to C++ programming in his book, **”The C++ Programming Language”**. This book is comprehensive and covers the full breadth of C++ features, including STL. Although it does not focus exclusively on STL, it is an indispensable resource for any serious C++ programmer.

- **Link:** [The C++ Programming Language - 4th Edition](#)
- **Highlights:**
 - In-depth coverage of C++ fundamentals and advanced features.
 - Detailed explanations of STL containers, iterators, and algorithms.
 - Thorough examples illustrating real-world applications of STL.
 - C++11, C++14, C++17, and C++20 features covered, with considerations for C++23.
 - Suitable for both novice and experienced C++ developers.

This book is a great reference, especially for understanding the relationship between STL and other advanced C++ concepts.

”Effective STL” by Scott Meyers

Scott Meyers is one of the most respected C++ experts, and his book **”Effective STL”** is a must-read for developers who want to become proficient in using STL. This book provides

detailed guidelines for efficient STL usage, tips for avoiding common pitfalls, and techniques to maximize the performance of STL containers and algorithms.

- **Link:** [Effective STL - Scott Meyers](#)

- Highlights:

- Focuses on best practices and tips for using STL effectively.
- Offers advice on selecting the right container for your use case.
- Discusses the performance implications of various STL operations.
- Covers C++98 through C++03 STL, but many of the principles are relevant to C++11, C++17, and C++23 as well.
- Aimed at intermediate to advanced C++ developers.

This book will give you a solid foundation for making informed, performance-conscious decisions when using STL in your projects.

”The C++ Standard Library: A Tutorial and Reference” by Nicolai M. Josuttis

This classic book by Nicolai Josuttis is one of the most comprehensive references for learning the **Standard Library** and **STL**. The book is an excellent resource for understanding how to use STL components in a practical and efficient way.

- **Link:** [The C++ Standard Library by Nicolai M. Josuttis](#)

- Highlights:

- Extensive coverage of STL containers, iterators, algorithms, and other components.
- Thorough exploration of **container types**, **memory management**, and **performance considerations**.

- Provides a detailed explanation of **newer features** of STL that were introduced in C++11 and C++14 (still relevant to C++23).
- Real-world examples and insights into **algorithmic complexity** and **use cases** for each STL component.
- Aimed at intermediate to advanced programmers.

This book is a fantastic resource for mastering the STL and understanding its design philosophy and internal workings.

”C++ Standard Library: A Tutorial and Reference” (C++17 Edition) by Nicolai Josuttis

While the previous edition focused on older standards, the **C++17 edition** of this book is an updated version that includes C++14 and C++17 features. It can still be very relevant for C++23 developers, as many of the foundational concepts haven't changed.

- **Link:** [C++ Standard Library \(C++17 Edition\)](#)

- Highlights

:

- Focus on **C++14** and **C++17** updates, with applications still relevant to **C++23**.
- Detailed examples and hands-on exercises, ideal for developers who prefer practical learning.
- Covers **ranges**, **memory management**, **custom allocators**, and **container optimizations**.

This book offers a deep dive into the STL for developers familiar with the basics of C++ and C++11 and is an excellent follow-up to the C++03 edition.

”C++ Template Metaprogramming” by David Abrahams

Though not exclusively about STL, **C++ Template Metaprogramming** is an essential book for those who want to harness advanced template techniques. It explains how to work with templates in a more advanced way, which is valuable when working with STL, as many STL components rely heavily on template-based programming.

- **Link:** [C++ Template Metaprogramming](#)
- Highlights:
 - Teaches advanced techniques related to **template programming**.
 - Provides a foundation for understanding **template-based STL** components.
 - Aimed at advanced C++ developers, especially those interested in **generic programming** and **compile-time computations**.

This book is highly recommended for developers who want to go beyond basic STL usage and start creating powerful, reusable, and efficient template-based code.

9.2.2 Training Courses for C++ and STL

In addition to books, online courses are an excellent way to learn C++ and STL in an interactive format. These platforms offer video lectures, assignments, quizzes, and projects that can enhance your learning experience. Here are some top-rated courses focused on C++ and STL.

”C++ Programming for Beginners” - Udemy

This course on Udemy provides an excellent foundation in C++ programming, including a focus on STL. It starts with the basics of C++ and gradually covers more advanced topics, including STL containers, algorithms, and iterators.

- **Link:** C++ Programming for Beginners
- Highlights:

- Focus on C++ syntax, structures, and memory management.
- In-depth introduction to **STL containers** and **algorithm libraries**.
- Practical examples, quizzes, and hands-on exercises.
- Ideal for beginners and intermediate learners.

This course provides a gentle introduction to C++ with enough coverage of STL to get you started using it in your own projects.

”Mastering C++” - Pluralsight

Pluralsight offers a comprehensive course on C++ programming that includes both theory and practice. The ”**Mastering C++**” course dives deep into the latest versions of C++ (including C++17 and C++23) and covers STL extensively.

- **Link:** Mastering C++ - Pluralsight
- Highlights:
 - Detailed coverage of **STL containers**, **algorithms**, and **iterators**.
 - Performance optimizations for STL-based code.
 - Tutorials on working with **ranges** and **concepts** in C++23.
 - Great for intermediate to advanced learners.

This course is ideal if you are already familiar with basic C++ but want to level up and explore STL in more depth.

”C++ STL” - Coursera (University of California, Santa Cruz)

This online course, offered by Coursera, dives specifically into STL and covers a wide range of STL components and advanced techniques. You will learn about **containers**, **algorithms**, and **iterators**, with a focus on practical applications and performance.

- **Link:** [C++ STL on Coursera](#)
- Highlights:
 - Focus on understanding **containers**, **iterators**, and **algorithmic operations**.
 - Real-world examples of STL usage in software development.
 - Ideal for both beginners and intermediate learners.

This course offers an excellent learning path for mastering STL, with hands-on coding challenges that reinforce the concepts learned.

”Advanced C++” - LinkedIn Learning

LinkedIn Learning offers a range of courses that go beyond the basics of C++ programming. Their ”Advanced C++” course focuses on **advanced STL features**, including **ranges**, **concepts**, and **algorithm optimization** in C++23.

- **Link:** [Advanced C++ - LinkedIn Learning](#)
- Highlights:
 - Focus on advanced STL topics, including **ranges** and **algorithm efficiency**.
 - Learn best practices for writing efficient and scalable code.
 - Great for experienced C++ developers who want to explore more advanced topics in STL.

Conclusion

Books and training courses are invaluable resources for mastering STL and advancing your C++ skills. The books listed above offer in-depth theory, examples, and best practices, while online courses provide interactive learning experiences with hands-on practice.

Whether you choose to learn through a structured book, a course, or a combination of both, these resources will help you leverage STL's full potential in your C++23 projects.

- **Books:** Essential for deep-dive exploration, covering both basic and advanced STL topics.
- **Courses:** Interactive and practical, suitable for learning at your own pace.

Together, these resources will set you on the path to mastering STL in C++23 and optimizing your C++ programming skills for real-world applications.

Appedices

Appendix A: Key C++23 STL Features at a Glance

Here's a quick reference to the most important updates in C++23 that impact the **Standard Template Library (STL)**:

1. New Templates and Concepts

- **`std::expected<T, E>`**: A template designed to simplify error handling in C++ by providing a way to represent and handle errors in a more expressive manner.
- **`std::ranges`**: This feature expands the capabilities of STL algorithms by adding range-based algorithms, making it easier to work with ranges (sequences of elements) without relying on raw iterators.
- **`std::span`**: A lightweight, non-owning view of a sequence of objects, offering better safety and flexibility when dealing with arrays and contiguous data.
- **`std::sem` and `std::latch`**: These synchronization primitives simplify concurrent programming by providing easy-to-use tools for signaling and thread synchronization.

2. Modified Containers

- **`std::vector`**: Now has improvements to performance, with better memory management and optimized algorithms.
- **`std::list`**: Introduced enhancements to support more efficient algorithms for list manipulation.

3. Enhanced Algorithms

- **`std::ranges` algorithm library**: The new algorithm suite supports **views** and **ranges** and provides more flexibility and expressive power.
- **`std::find`** and other algorithms now support **ranges**, allowing for cleaner and more declarative code.

4. Backward Compatibility and Performance

- C++23 continues to maintain strong backward compatibility while providing performance enhancements such as better **algorithmic complexity**, memory management improvements, and support for **`constexpr`** in more STL components.

Appendix B: STL Code Examples

This appendix includes a set of practical code examples that demonstrate the key concepts and features covered in the book. These examples will help reinforce your understanding of STL's functionality and provide a reference for how to apply STL components in real-world code.

Example: `std::expected` for Error Handling

```
#include <iostream>
#include <expected>

std::expected<int, std::string> divide(int numerator, int denominator) {
    if (denominator == 0) {
        return std::unexpected("Division by zero error.");
    }
    return numerator / denominator;
}

int main() {
    auto result = divide(10, 0);

    if (result) {
        std::cout << "Result: " << *result << std::endl;
    } else {
        std::cout << "Error: " << result.error() << std::endl;
    }
    return 0;
}
```

- **Explanation:** The `std::expected` type is used here to either return a valid result or an error string. This approach simplifies error handling and makes your code more readable and maintainable.

Example: Using `std::span`

```
#include <iostream>
#include <span>
```

```

void print_span(std::span<int> s) {
    for (int elem : s) {
        std::cout << elem << " ";
    }
    std::cout << std::endl;
}

int main() {
    int arr[] = {1, 2, 3, 4, 5};
    std::span<int> sp(arr, 5); // Span points to array 'arr'

    print_span(sp);
    return 0;
}

```

- **Explanation:** `std::span` is a lightweight wrapper for a contiguous array, providing a safer way to handle arrays without managing their memory. It does not take ownership of the data and allows for flexible slicing and passing of arrays to functions.

Example: Using `std::ranges`

```

#include <iostream>
#include <ranges>
#include <vector>

int main() {
    std::vector<int> nums = {1, 2, 3, 4, 5};

    auto result = nums | std::ranges::transform([](int n) { return n * 2;
    ↪ }) | std::ranges::to<std::vector>();
}

```

```
for (auto num : result) {  
    std::cout << num << " ";  
}  
std::cout << std::endl;  
  
return 0;  
}
```

- **Explanation:** This example demonstrates using the `std::ranges` library to transform a collection of integers. `std::ranges` provides more declarative syntax for manipulating data ranges, enabling cleaner, easier-to-read code.

Appendix C: Reference Material and Official Documentation

Here are several official sources and references to help you explore the C++23 Standard Library and STL components further:

1. ISO/IEC Standard C++23

- Official specifications for C++23 can be found on the [ISO C++ website](#). The standard document provides precise descriptions of all new features and updates in the language.

2. C++ Reference

- [cppreference.com](#) is a comprehensive and up-to-date resource for all aspects of C++ programming, including detailed documentation for STL components. It covers every function, template, class, and algorithm available in the standard library.

3. C++ Standard Library Documentation ([cppreference.com](#))

- STL Reference provides specific details about containers, algorithms, iterators, and ranges within the STL, along with code examples.

4. C++ Core Guidelines

- The C++ Core Guidelines is an essential document, providing best practices for writing modern, efficient, and safe C++ code. It covers a wide range of C++ features, including STL usage.

5. C++23 Features

- To keep track of the latest developments in C++23, refer to various blog posts and release notes from C++ standards meetings and conferences like CppCon and others:
 - C++23 New Features

Appendix D: Glossary of STL Terms

This glossary provides a quick reference to key STL terminology, helping you to understand the foundational concepts and jargon used throughout the book.

- **Container:** A data structure that stores a collection of elements, such as `std::vector`, `std::list`, or `std::map`.
- **Iterator:** An object that allows you to access the elements of a container. Iterators abstract away the container's internal implementation and provide a uniform interface to access elements.
- **Algorithm:** A function that performs an operation on a range of elements. Examples include `std::sort`, `std::find`, and `std::accumulate`.

- **Range:** A new abstraction introduced in C++20 (and extended in C++23) that represents a sequence of elements. Ranges are more flexible than iterators and allow algorithms to be written in a more declarative style.
- **Concept:** A constraint placed on template parameters that defines a set of requirements. Concepts simplify template programming by making templates easier to understand and ensuring that types used in templates meet specific criteria.

Appendix E: Further Learning Resources

In addition to the books, online courses, and documentation referenced throughout this book, the following additional learning materials are useful for C++23 STL enthusiasts:

1. YouTube Channels

- [TheCherno](#): A popular YouTube channel that provides tutorials and explanations for C++ development, including STL usage.
- [CppCon](#): The official YouTube channel for the C++ conference, with in-depth talks on modern C++ features, including STL and C++23 updates.

2. Forums and Communities

- **Stack Overflow:** The largest online community for C++ programmers where you can ask questions and share knowledge related to STL and C++23.
- **Reddit C++:** The C++ subreddit is a great place to find discussions and links to useful resources on STL and new C++ features.
- **C++ Discord Servers:** Many online C++ communities on Discord offer real-time help, tutorials, and discussions on the latest in C++ programming.

3. C++ Blogs

- [C++ Weekly](#): Offers tutorials and discussions on modern C++ practices, STL, and C++23.
- [Bjarne Stroustrup's Blog](#): A blog by the creator of C++ that provides insights into the evolution of the language, including STL.