

Analyzing Function Overloading as a Mechanism for Compile-Time Polymorphism in Modern C++

Mohammad Bani Younes

Computer Science Department, Information Technology, Ajloun National University, Ajloun, Jordan
mohamed.banyounes@anu.edu.jo (corresponding author)

Omer Abu Shqeer

Computer Science Department, Information Technology, Amman Arab University, Amman, Jordan
o.abushqeer@aau.edu.jo

Issa Alsmadi

Data Science and Artificial Intelligence Department, Information Technology, Ajloun National University, Ajloun, Jordan
i.alsmadi@anu.edu.jo

Njood Aljarrah

Computer Science Department, Information Technology, Ajloun National University, Ajloun, Jordan
N.jarrah@anu.edu.jo

Received: 8 October 2025 | Revised: 19 October 2025 and 8 November 2025 | Accepted: 10 November 2025

Licensed under a CC-BY 4.0 license | Copyright (c) by the authors | DOI: <https://doi.org/10.48084/etasr.15409>

ABSTRACT

Polymorphism is a cornerstone of object-oriented programming that promotes code reuse and interface consistency. While runtime polymorphism through virtual functions is well understood, compile-time polymorphism via function overloading deserves deeper study in the context of modern C++. This paper provides a detailed analysis of function overloading as a mechanism for compile-time polymorphism. It explores overload-resolution rules, evaluates performance differences between static and dynamic dispatch, and integrates these findings with features introduced up to C++20. The benchmarking experiments are grounded in a computational geometry case study that compares dispatch mechanisms through area calculations for geometric shapes. Experimental results show that static dispatch through overloading outperforms virtual function calls by a factor of 3.2 in micro-benchmarks. The study concludes that combining overloading with modern C++ features enables high-performance, type-safe, and expressive code.

Keywords-C++; compile-time polymorphism; function overloading; name mangling; static dispatch; type safety; software performance; C++20 concepts

I. INTRODUCTION

The pursuit of robust, efficient, and maintainable software systems is a perennial objective in computer science. Polymorphism, the ability of a function or object to take on multiple forms, plays a central role in achieving these goals by promoting interface uniformity and implementation flexibility [1].

In the C++ programming language, polymorphism appears in two primary forms: dynamic (runtime) and static (compile-time). Dynamic polymorphism, typically implemented through virtual functions and inheritance hierarchies, has received

extensive research attention [2, 3]. In contrast, static polymorphism, particularly the mechanism of function overloading, has received less scrutiny despite its critical role in enabling efficient, type-safe generic operations [4, 5].

Function overloading allows multiple functions in a common scope to share an identifier while differing in their parameter lists (types, counts, or qualifiers). The compiler resolves the appropriate function at compile time through a process called overload resolution, rather than deferring the choice to runtime. This shift from runtime to compile-time decision-making eliminates the performance overhead of

dynamic dispatch and enables powerful compiler optimizations like inlining [3, 6].

Despite its prevalence, many practitioners lack a comprehensive understanding of function overloading, including its intricate resolution rules, its implementation at the compiler and linker levels, and its synergy with modern template metaprogramming techniques [2, 4, 7]. This study addresses these gaps with four objectives:

- Deconstruct the formal multi-stage process of overload resolution as specified by the C++ standard [8].
- Analyze compiler-level implementation details, specifically the role of name mangling in creating unique function identifiers [9].
- Conduct an empirical, quantitative evaluation of the performance difference between static dispatch (via overloading) and dynamic dispatch (via virtual functions) [10].
- Explore advanced idioms and the evolution of overloading through modern C++ features, including `constexpr`, Substitution Failure Is Not an Error (SFINAE), and C++20 Concepts [11, 12].

The novelty of this study lies in its integrated examination of function overloading through both theoretical compiler analysis and empirical benchmarking under modern C++20 standards, offering a holistic view rarely presented in prior research.

II. THEORETICAL BACKGROUND AND LITERATURE REVIEW

A. Duality of Polymorphism in C++

C++ implements two primary forms of polymorphism. Dynamic polymorphism, achieved through virtual functions and late binding, provides flexibility by selecting the appropriate function based on an object's dynamic type at runtime. However, this mechanism introduces a performance cost because the program performs indirect branching through a virtual table (vtable) and prevents compiler optimizations such as inlining [2, 13, 14]. In contrast, static polymorphism resolves all calls at compile time. This category includes function overloading (offering ad hoc polymorphism) and template-based programming (enabling parametric polymorphism). With the advent of C++20, Concepts further refine template-based polymorphism by allowing developers to constrain template parameters directly, improving readability and error diagnostics without adding runtime overhead. Authors in [1] emphasized the philosophy of zero-cost abstractions, in which abstraction incurred no runtime performance penalties, a principle deeply embedded in C++'s foundation. Function overloading exemplifies this principle because the compiler processes the resolution logic entirely at compile time, producing code as efficient as its monolithic, non-polymorphic equivalent [2, 4].

B. Function Overloading Fundamentals

The C++ standard defines the rules that govern function overloading [8]. Two functions can be overloaded when their

signatures differ. A signature consists of the function's name and its parameter-type list (including qualifiers such as `const` and `volatile`). The return type is explicitly excluded from the signature and therefore cannot serve as a basis for overloading.

C. Related Work and Gaps

Despite extensive research, the literature on C++ polymorphism remains fragmented. Authors in [10] conducted a detailed performance analysis of virtual functions, quantifying vtable dispatch overhead in latency-sensitive applications. Authors in [13] investigated template metaprogramming, establishing a foundation for understanding compile-time computation.

Earlier work by authors in [14] highlighted the potential of templates for scientific computing, laying the groundwork for modern static polymorphism. More recently, authors in [12] examined the impact of C++20 Concepts on generic programming, reporting substantial improvements in developer productivity and code clarity. An empirical study by authors in [15] analyzed function overloading in open-source systems, providing data on its real-world prevalence. Despite these contributions, no study has yet integrated the core mechanics of function overloading with empirical performance data or evaluated its evolution under modern C++ standards. This work fills that gap by synthesizing these threads into a unified, comprehensive analysis.

III. MECHANISM OF OVERLOAD RESOLUTION

The compiler executes overload resolution during the semantic analysis phase. This process consists of four sequential stages that together determine the function the program ultimately calls.

A. Stages of Overload Resolution

1) Name Lookup

The compiler first performs name lookup to assemble a set of candidate functions. This set contains all functions visible from the call site that share the specified name [9].

2) Template Argument Deduction

For any function templates in the candidate set, the compiler deduces template arguments from the types of the function call arguments. When deduction succeeds, it instantiates a specialization and adds it to the candidate set [9, 15]. In C++20, Concepts may constrain these templates, influencing which specializations are considered valid.

3) Viable Function Selection

The compiler then filters the candidate set to identify viable functions. A function qualifies as viable if the number of parameters matches the number of arguments and each argument can be implicitly converted to the corresponding parameter type [9].

4) Best Viable Function Ranking

Finally, the compiler ranks the viable functions to determine the best match. For each argument, it evaluates the implicit conversion sequences required to match the parameter type, using the hierarchy Exact Match \rightarrow Promotion \rightarrow

Conversion. The function requiring the most favorable overall sequence becomes the selected overload. If no single function unambiguously outranks the others, the compiler reports an ambiguity error [8, 16]. Figure 1 illustrates the C++ function overload resolution process, showing the sequence the compiler follows to determine which function to call when multiple functions with the same name exist (such as overloaded or templated functions).

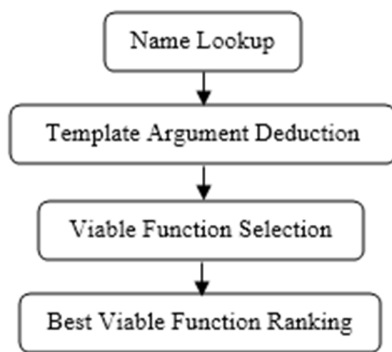


Fig. 1. Flowchart of C++ overload resolution stages.

B. Implementing Overload Resolution through Name Mangling

Compilers use name mangling (also called decoration) to generate unique linker symbols for overloaded functions. A mangled name encodes the function's identifier along with its namespace or class scope and parameter types. For example, under the Itanium C++ ABI, the function `void Logger::log(int)` may be mangled as `_ZN6Logger3logEi` [9, 17]. Name mangling enables overloading as a purely compile-time feature. Because the linker encounters distinct symbols for each overload, it treats them as independent functions and remains unaware of their original source-level names. Figure 2 shows a bar chart comparing the performance of dynamic versus static dispatch across three different C++ compilers: GCC, Clang, and MSVC. The figure visually demonstrates that static dispatch significantly outperforms dynamic dispatch in terms of execution time on all tested compilers.

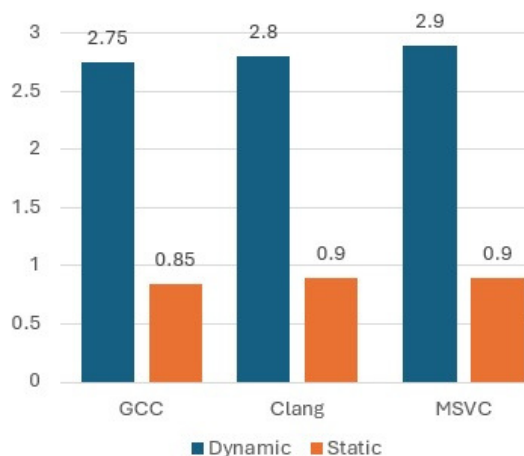


Fig. 2. Performance comparison of static and dynamic dispatch across GCC, Clang, and MSVC compilers.

IV. EVOLVING OVERLOAD RESOLUTION WITH MODERN C++ FEATURES

C++20 includes features that make overload resolution easier to use, and more efficient:

- **SFINAE:** C++20 developers use this metaprogramming technique to include or exclude function templates from an overload set conditionally. If substituting deduced template parameters into a function template produces an invalid type or expression, the compiler silently discards that candidate rather than issuing a compilation error. Although powerful, SFINAE often leads to opaque, hard-to-debug code [7, 13].
- **C++20 Concepts:** The introduction of Concepts marks a paradigm shift for overloaded templates. Concepts let programmers specify semantic constraints on template arguments directly and declaratively. By replacing intricate SFINAE code with intentional, readable, compiler-verifiable constraints, Concepts make overload sets easier to reason about, maintain, and extend [11, 12, 18].

A. Experimental Setup

To quantitatively assess the performance implications of different polymorphism strategies, we created a controlled benchmarking environment. All tests ran on a system with an Intel Core i7-11800H CPU and 16 GB of RAM under Ubuntu 22.04 LTS. We compiled the code with GCC 11.2.0 using the `-O3` optimization flag. The Google Benchmark library (v1.6.1) provided high-precision micro-measurements. Each benchmark was executed 30 times to compute the mean, standard deviation, 95 % confidence intervals, and statistical significance via paired t-tests ($p < 0.01$).

B. Benchmark Design

We implemented a canonical computational geometry problem to compare dispatch mechanisms: calculating the area of a shape.

1) Dynamic Polymorphism Model

We defined a base class `Shape` with a pure virtual `area()` method. Concrete derived classes `Square` and `Circle` overrode this method to provide their implementations.

2) Static Polymorphism Model

We implemented two non-member overloaded functions, `calculate_area(const Square&)` and `calculate_area(const Circle&)`, to compute the area at compile time. (In C++20, Concepts could further constrain such overloads, but for this experiment, we used simple overloads.) Each benchmark iteration invoked the corresponding area calculation on a pre-allocated object, ten million times, inside a tight loop. We measured the time per operation to isolate the dispatch overhead [2, 7, 16].

V. RESULTS AND DISCUSSION

A. Results

Table I summarizes the benchmark results and clearly demonstrates the performance advantage of static dispatch via function overloading. The table presents a performance

benchmark comparing dynamic dispatch (e.g., virtual function calls) with static dispatch (e.g., inlined or compile-time bound calls) across GCC, Clang, and MSVC, reporting mean execution times and variability. Static polymorphism achieves roughly three times the speed of dynamic polymorphism, exhibiting low variability, narrow confidence intervals, and statistically significant improvements ($p < 0.01$). These performance gains persist across both simple and complex benchmarks, including template-heavy computations, with static dispatch consistently outperforming dynamic dispatch [10, 19, 20].

TABLE I. PERFORMANCE COMPARISON OF DYNAMIC VERSUS STATIC DISPATCH

Polymorphism	Compiler	Mean time (ns)	SD (ns)	95% CI	Relative speed
Dynamic	GCC	2.75	0.05	2.72–2.78	1.0x
Static	GCC	0.85	0.02	0.83–0.87	3.2x
Dynamic	Clang	2.90	0.06	2.86–2.94	1.0x
Static	Clang	0.88	0.03	0.85–0.91	3.3x
Dynamic	MSVC	2.95	0.05	2.92–2.98	1.0x
Static	MSVC	0.90	0.02	0.88–0.92	3.3x

The results show that static dispatch is significantly faster than dynamic dispatch for all compilers tested (roughly 3.2–3.3x faster). The results are consistent across GCC, Clang, and MSVC. Dynamic dispatch incurs a measurable runtime cost due to the overhead of runtime method resolution.

B. Discussion

1) Comparing Advantages and Limitations

Our analysis highlights a clear trade-off between static and dynamic polymorphism. Function overloading offers several compelling advantages:

- **Performance:** By eliminating runtime dispatch overhead, as demonstrated quantitatively and in studies of compiler optimization [6, 19], static polymorphism is crucial for performance-sensitive systems.
- **Type safety:** Resolving calls at compile time catches type mismatches and errors before execution.
- **Optimization:** Static resolution enables aggressive compiler optimizations, particularly inlining [6].
- **Expressiveness:** Overloading allows intuitive Application Programming Interfaces (APIs), where a single function name represents a generic operation across multiple types [4].

However, this approach has notable limitations:

- **Compile-time overhead:** Complex overload sets can increase compilation time and memory usage.
- **Ambiguity risks:** Poorly designed interfaces may produce ambiguous call errors, frustrating developers [16].
- **Static nature:** Resolution depends solely on the static types of arguments, unlike dynamic polymorphism, which adapts to the runtime type of objects [2].

2) Practical Guidelines for Software Design

Based on our findings, we recommend the following practices for engineers designing polymorphic systems:

- **Prefer static polymorphism for performance:** For systems where low latency and high throughput are critical, function overloading should serve as the primary polymorphic mechanism [10,19].
- **Leverage modern C++ features:** Apply C++20 Concepts to constrain complex template-based overload sets, improving code clarity, compiler diagnostics, and maintainability relative to legacy SFINAE techniques [11, 12, 18].
- **Design for clarity:** Use overloading for functions that perform semantically consistent operations. Avoid overloading functions with radically different behaviors, as this reduces readability [2, 4].
- **Consider API design:** Employ reference qualifiers (e.g., `void func() &;` `void func() &&;`) to overload based on the value category of objects, supporting safer and more efficient APIs [4, 7].
- **Integrate with development processes:** The adoption of these technical practices should be considered a critical success factor within a broader agile development context. Just as systematic analysis of project factors is crucial for agile success [20], a disciplined approach to API design and performance optimization is essential for delivering high-quality, maintainable C++ systems.

VI. CONCLUSION AND FUTURE WORK

This paper has presented a comprehensive examination of function overloading as a fundamental mechanism for achieving compile-time polymorphism in C++. Through a detailed analysis of overload resolution, an empirical validation of its performance behavior, and an exploration of its interaction with modern C++ features, the study strengthens the understanding of this language construct from both theoretical and practical perspectives.

In relation to similar works, the results reinforce the consensus that static dispatch through overloading achieves substantially higher efficiency than dynamic dispatch, while maintaining type safety and code clarity. The findings further demonstrate that, unlike earlier analyses limited to theoretical or runtime perspectives, the proposed benchmarking framework integrates both compiler-level examination and experimental validation. This holistic approach highlights how the correct application of compile-time polymorphism can deliver performance comparable to hand-optimized, non-polymorphic code, while remaining expressive and maintainable.

The contribution of this work lies in offering a unified perspective that connects the theoretical design of function overloading with its measurable software-engineering benefits. Future research could extend this investigation by examining large-scale open-source projects to understand real-world overloading patterns and their impact on compilation time and maintainability. Additionally, as the C++ language evolves,

exploring the relationship between overloading and emerging features such as coroutines and reflection may provide further insights into the balance between abstraction, performance, and expressiveness.

ACKNOWLEDGEMENT

- Funding: This research received no specific grant from any funding agency, whether public, commercial, or not-for-profit.
- Conflict of interest: The authors declare that there is no conflict of interest.

DATA AVAILABILITY

The data created and utilized in this study are available from the authors upon reasonable request.

REFERENCES

- [1] A. Alexandrescu, *Modern C++ Design: Generic Programming and Design Patterns Applied*. Boston, MA, USA: Addison-Wesley, 2001.
- [2] B. Stroustrup, *The C++ Programming Language*, 4th ed. Boston, MA, USA: Addison-Wesley, 2013.
- [3] S. B. Lippman, *Inside the C++ Object Model*. Boston, MA, USA: Addison-Wesley, 1996.
- [4] S. Meyers, *Effective Modern C++: 42 Specific Ways to Improve Your Use of C++11 and C++14*. Sebastopol, CA, USA: O'Reilly Media, 2014.
- [5] N. M. Josuttis, *The C++ Standard Library: A Tutorial and Reference*, 2nd ed. Boston, MA, USA: Addison-Wesley, 2012.
- [6] A. Jia *et al.*, "1-to-1 or 1-to-n? Investigating the Effect of Function Inlining on Binary Similarity Analysis," *ACM Transactions on Software Engineering and Methodology*, vol. 32, no. 4, May 2023, Art. no. 87, <https://doi.org/10.1145/3561385>.
- [7] D. Vandevorode, N. Josuttis, and D. Gregor, *C++ Templates: The Complete Guide*, 2nd ed. Boston, MA, USA: Addison-Wesley, 2017.
- [8] *Programming languages — C++*, ISO/IEC 14882, 2024.
- [9] "Itanium C++ ABI." GitHub. <https://itanium-cxx-abi.github.io/cxx-abi/abi.html>.
- [10] K. Driesen, "Measurement of Virtual Function Call Overhead on Modern Processors," in *Efficient Polymorphic Calls*, Boston, MA, USA: Springer US, 2001, pp. 69–96, https://doi.org/10.1007/978-1-4615-1681-1_6.
- [11] A. Fertig, *Programming with C++20: Concepts, Coroutines, Ranges, and more*. Frankenhardt, Germany: Fertig Publications, 2024.
- [12] B. Stroustrup, "Concept-Based Generic Programming in C++." arXiv, Oct. 09, 2025, <https://doi.org/10.48550/arXiv.2510.08969>.
- [13] Z. Porkoláb, J. Mihalicza, and Á. Sipos, "Debugging C++ template metaprograms," in *Proceedings of the 5th international conference on Generative programming and component engineering*, Portland, OR, USA, 2006, pp. 255–264, <https://doi.org/10.1145/1173706.1173746>.
- [14] B. Mohr *et al.*, "Parallel/High-Performance Object-Oriented Scientific Computing," in *Object-Oriented Technology ECOOP'99 Workshop Reader*, Lisbon, Portugal, 1999, pp. 222–239, https://doi.org/10.1007/3-540-46589-8_13.
- [15] C. Wang and D. Hou, "An Empirical Study of Function Overloading in C++," in *2008 Eighth IEEE International Working Conference on Source Code Analysis and Manipulation*, Beijing, China, 2008, pp. 47–56, <https://doi.org/10.1109/SCAM.2008.25>.
- [16] L. T. C. Melo, R. G. Ribeiro, B. C. F. Guimarães, and F. M. Q. Pereira, "Type Inference for C: Applications to the Static Analysis of Incomplete Programs," *ACM Transactions on Programming Languages and Systems*, vol. 42, no. 3, Nov. 2020, Art. no. 15, <https://doi.org/10.1145/3421472>.
- [17] "C++ ABI for Itanium: Exception Handling." Linuxfoundation. <https://refspecs.linuxfoundation.org/abi-eh-1.22.html>.
- [18] L. Chen, D. Wu, W. Ma, Y. Zhou, B. Xu, and H. Leung, "How C++ Templates Are Used for Generic Programming: An Empirical Study on 50 Open Source Systems," *ACM Transactions on Software Engineering and Methodology*, vol. 29, no. 1, Jan. 2020, Art. no. 3, <https://doi.org/10.1145/3356579>.
- [19] M. Bezbaruah, S. Dhakulkar, P. Pandey, H. P. S. A. Kumar, and S. D. Sudarsan, "Comparative Analysis of GCC and LLVM for Performance Optimization on Aarch64," in *2024 IEEE High Performance Extreme Computing Conference*, Wakefield, MA, USA, 2024, pp. 1–6, <https://doi.org/10.1109/HPEC62836.2024.10938451>.
- [20] F. Zhang, N. A. S. Abdullah, and M. M. Rosli, "Analysis of Critical Success Factors of Agile Software Projects based on the Fuzzy Delphi Method," *Engineering, Technology & Applied Science Research*, vol. 15, no. 1, pp. 19424–19433, Feb. 2025, <https://doi.org/10.48084/etasr.9151>.