

A Comparative Analysis of C++, Java and Python: Strengths and Limitations and Future Trends of Language Execution

Contents

Introduction.....	2
Performance.....	3
Memory Management.....	4
Language Features.....	7
Future Trends.....	8
Conclusion.....	9
References.....	11

Introduction

Programming languages play an important role in software development, offering a plethora of features, paradigms, and trade-offs. Among the multitude of programming languages that have emerged in recent decades that have shaped the software landscape, C++, Java, and Python stand out as three distinct choices, each catering to different requirements and needs.

Python is a high-level object-oriented language and the most popular programming language in today's world. Comparatively, C++ is a low-level and object-oriented programming language and is usually used to create most operating systems. Java is another object-oriented programming language that holds a lot of influence on C++. However, to enable simplicity, it excludes features such as pointers and operator overloading (Gherardi, L., Brugali, D., & Comotti, D. (2012)).

In this essay, I will dive into the intricacies of the language of C++ and compare its merits and drawbacks to both Java and Python. The exploration will include the analysis and comparison of each language's performance, memory management, and language features. Additionally, I will embark on discussing what the future of these languages may look like. This involves considering the potential adoption of each other's features among languages, suggesting a future where the distinct boundaries between them could become less defined. I will aim to decipher which of these three languages stands out with the most robust and compelling features. By examining the strengths and limitations of these languages, insights can be gained into when each language excels and where it might have constraints.

Performance

Performance is a critical need in programming languages, particularly in scenarios where execution speed is crucial to the code. Examples are game development, high-frequency trading applications as well as system-level programming. The term 'performance' could include many aspects such as execution speed,

code size, data size, and memory footprint at run-time. Many programmers prioritise the speed at which their programs execute (W. T. L. P. Lavrijsen and A. Dutta. (2016)).

C++ has multiple high-performing advantages. Its low-level features, such as direct memory access and pointer arithmetic, allow developers to optimise code for efficiency. It has been designed with a strong focus on efficiency from its inception. The principle of "zero overhead" is a key guiding principle for C++ (Goldthwaite, L. (2006)). This principle states that any language feature not used in a program should not impose additional costs in terms of code size, memory usage, or runtime performance. This principle underscores that extraneous features must not lead to unnecessary performance drawbacks or increased memory usage.

Java, while not as performant as C++ in certain scenarios due to the Java Virtual Machine (JVM) overhead, has made strides in performance optimization (Martin, R. C. (1997)). Just-In-Time (JIT) compilation and advanced runtime optimisations bridge the gap between Java and C++ in many use cases. Java could achieve better performance than C++ by utilising dynamic compilation, which allows the compiler to access runtime information to optimise the code during execution. This would give the language an advantage by providing the compiler with data about the program's behaviour as it runs. This information enables more advanced optimisations that are not available to traditional C++ compilers (Reinholtz, K. (2000)).

Python in comparison to these two languages performs the worst due to its interpreted nature. Compared to C++ and Java, two compiler languages, Python code is executed line by line by the Python interpreter, which adds overhead. A study from the University of Toronto in 2022 revealed that Python (CPython) incurs substantial overhead, performing 8.01x and 29.50x slower on average than C++ counterparts. Moreover, CPython struggles to effectively utilise multiple cores, primarily due to its Global Interpreter

Lock (GIL), which hampers parallelism and scalability in multi-core systems (Lion, D., Chiu, A., Stumm, M., & Yuan, D. (2022)).

Memory Management

Part of a good language performance is its memory system which involves allocating, deallocating, and managing memory (Zehra, F., Javed, M., Khan, D., & Pasha, M. (2020)). Additional code is needed to manage changing memory requirements as the amount of memory required for running an application cannot be estimated.

C++ provides an efficient approach to dynamically allocate memory. This language uses new and deleted operators or smart pointers to handle memory allocation and deallocation (Zehra, F., Javed, M., Khan, D., & Pasha, M. (2020)). As C++ offers manual memory management developers gain detailed control over what memory they can allocate and deallocate. This implies that once a memory is allocated explicitly using a new keyword, it cannot be reused until and unless it is deallocated with a deleted keyword. This approach allows for efficient memory utilisation. There are many advantages to this type of memory management such as allowing for optimisation of memory usage and performance. Additionally, this time of management would be well-suited for applications where predictable memory behaviour is crucial to its performance such as game engines and embedded systems. Nevertheless, this style of memory management is inclined to have memory-related bugs such as memory leaks and invalid memory accesses. This consequently occurs because memory that is allocated may never be deallocated even after all pointers to it have been destroyed (Vassev, E., & Paquet, J. (2006)). Additionally, developers would be required to have a deep understanding of memory management in order for errors such as these could be avoided. Overall, this style of management may be more time-consuming and error-prone compared to automatic memory management.

In contrast to C++, Java does not support the delete operator and instead employs a technique known as garbage collection to manage memory (Martin, R. C. (1997)). This is done by a process within the Java Virtual Machine (JVM). This form of memory management frees sections of memory automatically once they are no longer being used and an object with no reference variables pointing to it is categorised as "garbage" (Vassev, E., & Paquet, J. (2006)). Subsequently, the Java runtime system takes the responsibility of eliminating this "garbage" at a later point in time, thereby reclaiming the memory it occupied and returning it to the memory heap. This technique is efficient as it will simplify the development of certain types of applications and software developers are spared the concern of meticulously cleaning up unwanted memory. The key advantage of garbage collection is that it automates memory management, reducing the likelihood of memory-related bugs unlike C++. However, overhead may still occur due to the periodic garbage collection process which can result in unpredictable pauses during garbage collection cycles. Overall, this technique may not be as efficient in situations where precise control over memory allocation is needed.

Python uses automatic memory management similarly to Java but with different mechanisms. Python utilises reference counting and a cyclic garbage collector (Zehra, F., Javed, M., Khan, D., & Pasha, M. (2020)). Objects are created and stored in a private heap managed by the Python memory manager. The memory manager handles the allocation and deallocation of memory through a garbage collector. Objects that are no longer reachable are automatically marked for deallocation and released by the garbage collector, unlike manual memory management in C++. Reference counting tracks the number of references to an object, and when the count drops to zero, the object is immediately deallocated (Zehra, F., Javed, M., Khan, D., & Pasha, M. (2020)). The cyclic garbage collector identifies and cleans up circular references that reference counting alone cannot handle. Additionally, objects are tracked across generations, and once allocation minus deallocation hits the threshold, an object may be deleted or moved to the next generation if it still has references. This technique has many advantages including securing a balance between ease of use and memory management and reducing the risk of memory leaks by

deallocating objects when they are no longer referenced. Nonetheless, in contrast to C++, Python streamlines memory management, albeit at the cost of potential overhead linked with reference counting and garbage collection. It is pertinent for developers to remain conscious of circular references that elude the grasp of solitary reference counting. Consequently, while Python's memory management framework epitomises simplicity, it might not match the efficiency of manual memory management concerning memory consumption and execution speed.

Language Features

The design of language features influences how developers write code and solve problems. Each language's set of features caters to different programming paradigms and styles. Since 1990, C++ has seen multiple large revisions of the language, adding object-oriented language features such as classes and inheritance as well as many other extensions such as template types, operator overloading, an exception mechanism, and others (Prechelt, L. (2003)). Most of this language's strengths and limitations are identical to those of C, except that it is much larger and more powerful which could be an additional con as it may require exceptional developers to work with it. Well-written C++ programs can be very efficient and have good access to low-level features of the machine. C++ encompasses a range of programming paradigms, including procedural, object-oriented, and generic programming. Its comprehensive standard library and template metaprogramming offer robust abstractions to address diverse tasks.

Java is also primarily object-oriented and enforces strong encapsulation through its class-based design. In comparison to C++, it is more compact, deliberately excluding certain complex constructs. The language is bundled with an extensive standardised library that encourages sound design practices (Prechelt, L. (2003)). Unlike being directly compiled into machine code, Java is transformed into machine-independent bytecode, which is further compiled during load time through a "just-in-time compiler" (JIT). These traits

of Java facilitate the creation of well-structured programs, yet they pose challenges to achieving the same level of efficiency as C++, both for programmers and compilers.

Python emphasises simplicity and readability, offering a concise and expressive syntax (Lo, C. A., Lin, Y. T., & Wu, C. C. (2015, April)). It supports multiple programming paradigms, including procedural, object-oriented, and functional approaches. The combination of dynamic typing and interpreted nature makes Python an excellent choice for rapid prototyping and scripting tasks. Python's syntax is simpler, and it offers high-level data structures that enable the creation of concise programs. The language's diverse paradigms also provide learners with the opportunity to explore various programming language features, making Python an increasingly popular choice as a first language for learning. Similarly to Java, Python's well-organised constructs allow seamless combination without requiring special governing rules. This design parallels Java, where distinct modules house built-in features for explicit and individual utilisation, aiding incremental learning. In essence, Python's emphasis on simplicity and learnability contrasts C++, which caters to experienced developers with its advanced features and control.

Future Trends

In addition to the well-established languages of C++, Java, and Python, the future of language capabilities are being shaped by innovative tools such as PyPy and Cling (Lavrijsen, W. T., & Dutta, A. (2016, November)).

As mentioned previously, Python's performance is slower than competitor languages such as C++ and Java. When combined with high-performance libraries, Python's slower speed isn't a big disadvantage as most of the time-consuming tasks are handled by these libraries and not the Python interpreter itself. Incorporating Python as a high-level productivity language alongside high-performance C++ libraries necessitates the development of effective, functional, and user-friendly cross-language connections. As a

result, the overall impact of Python on the final performance is minimal. The existence of well-made connections between Python and these libraries has caused a rise in the number of scientists who choose to use Python exclusively for their programming (Lavrijsen, W. T., & Dutta, A. (2016, November)). As more and more code is written in Python, how fast Python performs becomes more significant.

The PyPy project has become a pivotal factor in this problem. PyPy offers an alternative Python interpreter that works well and is compatible with existing Python code (Kundu, B., Vassilev, V., & Lavrijsen, W. (2023)). It improves runtime performance through just-in-time (JIT) compilation. By dynamically compiling Python code, PyPy aims to accelerate execution speeds beyond what traditional interpreters offer. This approach challenges the conventional trade-off between Python's ease of use and execution efficiency, making it an intriguing option for performance-critical applications.

Cling, on the other hand, introduces a novel way to interact with C++ code by providing a C++ interpreter. Traditionally, C++ has been known for its compilation process, which can be time-consuming (Kundu, B., Vassilev, V., & Lavrijsen, W. (2023)). Cling allows developers to experiment with C++ snippets interactively, enabling rapid prototyping and exploration without the overhead of compilation. This tool is particularly useful for testing ideas and algorithms flexibly and interactively. While Cling itself focuses on C++, it can be used to create a bridge between Python and C++ by providing a runtime environment where both languages can communicate.

Finally, one specific tool that facilitates the seamless integration of C++ code with Python is 'cppyy', enabling Python programs to use C++ classes, functions, and libraries as if they were native Python objects. It is executed by automatically generating Python bindings for C++ code, allowing Python developers to interact with C++ functionalities directly from their Python scripts (Kundu, B., Vassilev, V., & Lavrijsen, W. (2023)). In essence, Cppyy serves as a bridge that bridges the gap between the two

languages, enabling interoperability and leveraging the strengths of both. These tools exemplify the broader trend of merging languages to create more versatile and efficient programming environments.

Conclusion

To conclude, the comparison and analysis of C++, Java, and Python have revealed distinct strengths, limitations, and potential futures for each language. Each of these languages addresses unique needs within the software development landscape. C++ is the most notable of the three languages for its high performance, enabled by low-level features like direct memory access and pointer arithmetic. Its "zero overhead" principle reinforces efficient code generation, making it a robust choice for applications where execution speed is crucial. While Java does not hold a similar performance to C++ in certain scenarios, it still bridges the gap through JIT compilation and runtime optimizations. This approach, combined with its extensive standardised library, facilitates well-structured programs and allows for the creation of versatile applications. Contrastly, Python's interpreted nature introduces performance challenges, making it less suitable for performance-critical applications. However, its simplicity, readability, and dynamic typing offer unparalleled ease of use and rapid prototyping capabilities, making it an ideal choice for scripting and exploration as well as being a good language for beginners.

Memory management strategies also highlight the language's strengths. C++ provides manual memory management, offering detailed and flexible control over memory allocation and deallocation. While efficient, it requires a deep and experienced understanding of memory management and is likely to have memory-related bugs. Java's use of garbage collections and automating memory management significantly reduces the likelihood of memory-related bugs. Java and Python hold the most similarity in this area as Python similarly uses reference counting and a cyclic garbage collector to achieve a balance between ease of use and memory management.

Looking to the future, innovative tools like PyPy and Cling are reshaping language execution. PyPy's JIT compilation demonstrates an increase in performance for Python, mitigating its historical performance challenges. Cling introduces interactive C++ experimentation, potentially creating a bridge between Python and C++. Moreover, the 'cppy' tool showcases the potential of seamless integration between Python and C++, enabling interoperability and leveraging the strengths of both languages.

In a rapidly evolving tech landscape, language convergence is vital. As challenges grow in complexity, integrating languages gains significance, leveraging strengths and innovating norms. The analysis of C++, Java, and Python calls attention to the importance of selecting languages for project alignment. The future offers blended approaches, surmounting limitations for heightened software development efficiency and versatility.

Word Count: 2561

References

Gherardi, L., Brugali, D. and Comotti, D., 2012. A java vs. c++ performance evaluation: a 3d modelling benchmark. In Simulation, Modeling, and Programming for Autonomous Robots: Third International Conference, SIMPAR 2012, Tsukuba, Japan, November 5-8, 2012. Proceedings 3 (pp. 161-172). Springer Berlin Heidelberg.

Goldthwaite, L., 2006. Technical report on C++ performance. ISO/IEC PDTR, 18015.

Kundu, B., Vassilev, V. and Lavrijsen, W., 2023. Efficient and Accurate Automatic Python Bindings with cppy & Cling. arXiv preprint arXiv:2304.02712.

Lavrijsen, W.T. and Dutta, A., 2016, November. High-performance Python-C++ bindings with PyPy and Cling. In 2016 6th Workshop on Python for High-Performance and Scientific Computing (PyHPC) (pp. 27-35). IEEE.

Lion, D., Chiu, A., Stumm, M. and Yuan, D., 2022. Investigating Managed Language Runtime Performance: Why {JavaScript} and Python are 8x and 29x slower than C++, yet Java and Go can be Faster?. In 2022 USENIX Annual Technical Conference (USENIX ATC 22) (pp. 835-852).

Lo, C.A., Lin, Y.T. and Wu, C.C., 2015, April. Which programming language should students learn first? A comparison of Java and Python. In 2015 International Conference on Learning and Teaching in Computing and Engineering (pp. 225-226). IEEE.

Martin, R.C., 1997. Java and C++ A critical comparison. Technical Note, Object Mentor.

Prechelt, L., 2003. Are scripting languages any good? A validation of Perl, Python, Rexx, and Tcl against C, C++, and Java. Adv. Comput. 57, pp.205-270.

Reinholtz, K., 2000. Java will be faster than C++. ACM Sigplan Notices, 35(2), pp.25-28.

Zehra, F., Javed, M., Khan, D. and Pasha, M., 2020. Comparative analysis of c++ and python in terms of memory and time.