

From C++ to Rust



RUST





"You fight with the compiler, so you don't fight with runtime bugs."



A little bit of history ...





What is Rust

Rust is an open-source general programming language that focuses on performance and safety (memory safety / type safety). It is primarily used for building command line tools, web applications, server apps or to be used in embedded systems.

Resource:

- Linux & Mac/OSX: run curl --proto '=https' --tlsv1.3 https://sh.rustup.rs -sSf | sh
- GitHub repo: https://github.com/rust-lang/rust
- Windows install link: https://www.rust-lang.org/tools/install
- Documentation: <u>https://doc.rust-lang.org/book/</u>
- Quick install: <u>https://rustup.rs/</u>
- Official site: <u>https://www.rust-lang.org/</u>

- 2006 \rightarrow started as a project develop in Mozilla by Graydon Hoare
- 2010 \rightarrow officially announced as a project
- 2015 \rightarrow Rust 1.0 (first stable released announce)
- 2021 → Rust Foundation is formed, and the project is no longer maintained solely by Mozilla. Companies that are part of Rust Foundations are: AWS, Google, Huawei, Microsoft and Mozilla
- 2022 → Linus Torvalds announce that Rust is probably going to be used in Linux Kernel in the near future

 2020 → Amazon announced its implication in using Rust as a language for various project (AWS FireCracker being one of them)



 2020/Sep → While not confirmed by Apple, there are roomers that Apple is also using Rust internally



• 2021 → Google joins Rust Foundation with the director of

Engineering for the Android Platform – Lars Bergstrom



Rust nation UK (2024)

https://www.youtube.com/watch?v=6mZRWFQRvmw&t=27012s



I'm super excited to have the opportunity to join the Rust Foundation board to both continue to suppor Rust and help to grow its usage at Google. I am currently the Director of Engineering for Android Platform Programming languages, where I work with teams supporting C++, Java, Kotlin, and Rust development. Like many other software projects, improving memory safety in our most performancesensitive code is a critical need on Android to both keep our users safe and reduce the number of

• 2022/Sep \rightarrow Rust for Linux Kernel is announced to be released in Linux kernel 6.1



• 2022/Sep \rightarrow Azure announce its support for Rust programming



programming languages commonly used for native applications that require high performance.

• Close after that event, Microsoft started to change some of its internal code to Rust.



memory-safe code is already reaching developers.

 And after Microsoft Build Conference from 2023, Microsoft announces its first kernel components written in Rust as part of their ecosystem.



THENEWSTACK

PODCASTS EBOOKS EVENTS NEWSLETTER CONTRIBUTE

ARCHITECTURE ENGINEERING OPERATIONS

 In May.2024, Microsoft donates 1M USD to Rust Foundation to confirm company interest in this language.

RUST / SOFTWARE DEVELOPMENT

Microsoft's \$1M Vote of Confidence in Rust's Future

Microsoft has made an unrestricted \$1 million donation to the Rust Foundation, demonstrating its commitment to the Rust programming language and its ecosystem.

May 7th, 2024 9:30am by Darryl K. Taft



• Additionally, NSA has issued a document that suggest using memory safety languages (such as Rust)





 Finally, it is worth mention that Discord uses Rust on several backend projects that require memory safety and increase performance: <u>https://discord.com/blog/search?query=rust</u>



WHY DISCORD IS SWITCHING FROM GO TO RUST

This service was a great candidate to port to **Rust** since it was small and self-contained, but we also hoped that **Rust** would fix these latency spikes....

(areuc) shives....



USING RUST TO SCALE ELIXIR FOR 11 MILLION CONCURRENT USERS

But Discord has been using **Rust** to make things go fast, and we posed a question: "Could we use **Rust** to go faster?". **Rust** is not a functional language, and will happily let you mutate data structures....



HOW DISCORD STORES TRILLIONS OF MESSAGES

Rust! We'd used it for a few projects previously, and it lived up to the hype for us. It gave us fast C/C++ speeds without having to sacrifice safety....

Other memorable notions:

- 2022 → CloudFlare announced Pingore (their proxy that connects Cloudflare to Internet – written in Rust): https://blog.cloudflare.com/how-we-built-pingora-the-proxy-that-connectscloudflare-to-the-internet/
- 2022 → Facebook announced their support for Rust for server side components: https://www.zdnet.com/article/the-rust-programming-language-just-got-a-big-boost-from-meta/
- 2022 → Google announced that they started to use Rust for Android to mitigate risks: https://www.zdnet.com/article/google-after-using-rust-we-slashed-android-memory-safety-vulnerabilities/
- 2023 → Github switch to a new search engine (BlackBird) written completely in Rust: <u>https://www.zdnet.com/article/github-builds-a-search-engine-for-code-from-scratch-in-rust/</u>
- 2023 → Meta announces Buck2 (a build system written in Rust): <u>https://engineering.fb.com/2023/10/23/developer-tools/5-things-you-didnt-know-about-buck2/</u>

Or other a little bit different:

Rust is rolling off the Volvo assembly line



embedded automotive 🔵 rust





astronaut lives. This article examines how NASA and SpaceX implement Rust, with



Rust IDEs

- Visual Studio Code: <u>https://marketplace.visualstudio.com/items?itemName=rust-lang.rust-analyzer</u>
- IntelliJ (RostOver):

https://blog.jetbrains.com/rust/2023/09/13/introducing-rustrover-a-standalone-rust-ide-by-jetbrains/

• Eclipse:

https://www.eclipse.org/downloads/packages/release/2019-09/r/eclipse-ide-rust-developersincludes-incubating-components

- Online IDE (Rust playground): https://play.rust-lang.org/
- Other online compilers: https://www.tutorialspoint.com/compile_rust_online.php https://replit.com/languages/rust https://www.onlinegdb.com/online_rust_compiler https://rust.godbolt.org/



Rust Characteristics

- Strong-typed & statically typed language
- LLVM backend (native compiler) gcc backend also a possibility in the future
- Ownership and lifetimes for variables
- Memory safety (allocation / access)
- No garbage collector
- Zero cost abstraction
- Move semantics
- Traits (for polymorphism)
- Package manager and build mechanisms



First Rust program

First RUST Program



• C-like syntax

First RUST Program



- C-like syntax
- However, there are some differences:
 - A function in C is defined by writing the return type first, while in Rust a function is defined using a special keyword fn
 - "printf" is a function in C/C++, while "print!" is a macro in Rust

First RUST Program



- C-like syntax
- However, there are some differences:
 - A function in C is defined by writing the return type first, while in Rust a function is defined using a special keyword fn
 - "printf" is a function in C/C++, while "print!" is a macro in Rust
- To specify the return value of a function, use the following syntax:
 "-> <type>"

<u>1. Using rustc (rust compiler) command line:</u>

- Make sure that rust is installed
- Create a file in a folder named "first.rs" and insert into it the "hello world example (the one with a main function)
- Run the following command from command line: rustc first.rs
- An executable file (e.g. first.exe if you run this command in Windows) should appear in the first.rs file
- Run the executable file created on the precedent step (e.g. run first.exe if you are on Windows)

2. Using cargo (rust package manager) from command line:

- Make sure that rust is installed
- Run the following command from command line: cargo new first
- You should see a new folder (named first) that was created in the current folder with the following structure:

١	Current folder	
\first	A folder that contains your first project	
\first\.git	A hidden folder with a git integration data	
\first\.gitignore	Ignore rules for git repo	
\first\Cargo.toml	Configuration file for first project (INI like format)	
\src	A folder with all rust sources	
\src\main.rs	The main file of the rust project	

2. Using cargo (rust package manager) from command line:

- Make sure that rust is installed
- Run the following command from command line: cargo new first
- You should see a new folder (named first) that was created in the current folder with the following structure:

١	[package]
\first	<pre>name = "first" we made = "0.1.0"</pre>
\first\.git	edition = "2022"
\first\.gitignore	
\first\Cargo.toml	# See more keys and their definitions at
\src	[dependencies]
\src\main.rs	

2. Using cargo (rust package manager) from command line:

- Make sure that rust is installed
- Run the following command from command line: cargo new first
- You should see a new folder (named first) that was created in the current folder.
- Modify the "\src\main.rs" to contain the hello world example
- In folder "\first" execute the following command: cargo run

```
\first>cargo run
   Compiling first v0.1.0 (E:\Lucru\Rust\temp2\first)
    Finished dev [unoptimized + debuginfo] target(s) in 0.81s
    Running `target\debug\first.exe`
Hello, world!
```

3. Try it using rust playground:

- Open a browser and go to https://play.rust-lang.org/
- Write the hello world code
- Hit the Run button from the top-left side of the web-page

S Rust Playground × +	× –	o x
\leftarrow \rightarrow C $ ightarrow$ play.rust-lang.org	፼ ☆ □	G :
RUN ▶ ··· DEBUG ∨ STABLE ∨ ··· SHARE TOOLS ∨	CONFIG	~ ⑦
<pre>1 fn main() { 2 print!("Hello world"); 3 }</pre>		
Execution		Close
Standard Error		
Compiling playground v0.0.1 (/playground) Finished dev [unoptimized + debuginfo] target(s) in 0. Running `target/debug/playground`	58s	
Standard Output		
Hello world		•
4		•

4. Use different features from specialized IDEs:

- Use features such as create new project (IntelliJ) or various command/prompts from Visual Studio Code
- In the backend the cargo utility is usually used





Rust vs other languages



Language	Learning Curve	Key Challenges	Ideal For
Python	Easy	Minimal; focuses on readability and simplicity	Beginners, data science, web development
JavaScript	Moderate	Asynchronous programming, dynamic typing	Web development, full-stack applications
Go	Moderate	Understanding goroutines and channels for concurrency	Backend services, cloud applications
Java	Moderate	Verbose syntax, object-oriented concepts	Enterprise applications, Android development
C++	Difficult	Manual memory management, complex syntax	Systems programming, game development
Rust	Very Difficult	Ownership model, lifetimes, borrow checker	Systems programming, performance-critical apps



Learning Difficulty

Language	Difficulty
HTML / CSS	10 / 100
Python	20 / 100
Ruby	30 / 100
JavaScript	40 / 100
C#	45 / 100
Switft	50 / 100
PHP	55 / 100
Go	60 / 100
Kotlin	65 / 100
Java	70 / 100
C++	75 / 100
Rust	80 / 100

https://www.crossover.com/resources/12-programming-languages-ranked-by-difficulty-chart



UN-INITIALIZED VARIABLES: In rust all variables, data members from structure, enum variants, etc MUST be initialized (if they are not the code will not compile)



MISSING RETURN TYPE: In rust all functions / methods must return a value if it is being specified in the method / function definition. Failing to do this results in a compiler error.

Rust	<i>C/C++</i>
<pre>fn foo(x: i32, y: i32) {</pre>	<pre>int foo(int x, int y) {</pre>
}	<pre>} int main() {</pre>
fn main() { let a: $i32 - foo(10, 20)$:	<pre>int a = foo(10, 20); }</pre>
<pre>}</pre>	
Won't compile → "foo" must return a i32	Compiles OK ! (release mode)

IMPLICIT CASTS: In rust you CAN NOT convert a number of a different type to another type **IMPLICITELLY**. You need to explicately specify the conversion.


EXHAUSTIVENESS CHECKING: In Rust, in a *match* (*switch in C++*) *block* all possible values must be addressed. Failing to do so results in a compiler error.





Rust vs C++ (Variables & Functions)

Feature	C++	Rust
Variable implicit mutable state	Mutable	Immutable
Const type	Means a constant numerical value or a value that can not be changed	Means a constant numerical.
Overflow operations	Allowed, not managed	Managed on DEBUG
Uninitialized variables	Yes (possible)	Not Possible
Missing return type for functions	Yes (possible)	Not Possible
Basic types with clear size	Partial (e.g. int, char might have different representations)	Yes (i8, i16, i32, u32, u64,)
Larger types (128 bytes)	Νο	Yes (i128 and u128)
Implicit casts between basic types	Yes	No (not allowed)
String support	Multiple (Ascii, WTF-16/32)	UTF-8 (guaranteed)
Exhaustiveness Checking	Not enforced	Enforced at compile time

OWNERSHIP: A memory zone / a resource in Rust has ONLY ONE owner (or more simply put, you can not have two variables that share the same memory).



MOVE SEMANTICS: Most of the types in Rust use move-semantics (an object is moved (ownership is transferred) and its lifetime ends (the source object is destroyed).



BORROWING LOGIC: An object can borrow either multiple immutable references or **ONE** and only **ONE** mutable reference (every other combination is not allowed).

```
Rust
                                                   Rust
fn main()
                                                   fn main()
     let mut s = String::from("ABC");
                                                        let s = String::from("ABC");
     let r1 = \&s;
                                                        let r1 = \&s;
     let r2 = &mut s;
                                                        let r^2 = \&s;
     println!("{r1}");
                                                        println!("{r1}");
              Won't compile \rightarrow we have both a mutable
                                                                              Compiles OK !
              and immutable references at the same time
                                                                       (multiple immutable references)
```

LIFETIME LOGIC: References have lifetime associated with lifetime rules and relations that Rust uses the validate if a reference exists (is being used outside its lifetime / scope)

```
C/C++
Rust
fn foo<'a>(a: &'a str,
                                                      const char* foo(const char* a, const char* b)
              b: &'a str) -> &'a str {
                                                           return a;
     а
fn main() {
                                                     int main() {
     let result;
                                                           const char* result;
          let s1 = String::from("h");
let s2 = String::from("w");
                                                                std::string s1 = "h";
std::string s2 = "w";
                                                                result = foo(s1.c_str(), s2.c_str());
          result = foo(\&s1, \&s2);
     println!("{}", result);
                                                           std::cout << result;</pre>
                                                                                        Compiles OK !
              Won't compile \rightarrow because result lives more
                                                                         (even if result points to an invalid memory address)
                   than s1 and s2 (longer lifetime)
```



Rust vs C++ (Memory Overview)

Feature	C++	Rust
Manual Memory	new/delete or smart pointers (unique_ptr, etc.)	Ownership and borrowing system (no manual free)
Garbage Collector	No (manual delete object)	No (Automatic clean up)
Undefined Behavior	Yes (possible)	Not Possible
Use after free	Yes (possible)	Not Possible
NULL pointer associated risks	Yes (possible) (via NULL or nullptr)	Not Possible (there is no null pointer in Rust)
Dangling pointers	Yes (possible)	Not Possible
Lifetime support	No	Yes (implicit for each reference)
Thread safety	Manual	Enforced at compile time
Exception support	Yes	Νο
Error handling	Not enforced	Enforced at compile time via Result <t,r> or Option<t></t></t,r>

ALGEBRAIC DATA TYPE: An *enum* in Rust can have different data types for each of its fields. It can also be a regular enum (just like the ones from C/C++)

Rust	C/C++
<pre>enum Color { Red, Green, Blue, } enum ComplexColor { Red, Green, Blue, RGB(u32), ARGB(u8,u8,u8), Named(String) }</pre>	<pre>enum Color { Red, Green, Blue, }</pre>

SELF-CONSUME: In Rust you can create a method in a struct that can consume the instance (e.g. make sure that after the execution of that method the instance is destroyed).





Rust vs C++ (Classes)

Feature	C++	Rust
Encapsulation	private, protected, public modifiers	via pub and modules
Extending existing class	No	Yes
Inheritance	Yes	Νο
Method overloading	Yes (possible)	Not Possible
Virtual methods / Polymorphism	Yes (possible)	Yes (highly efficient)
Interfaces	Partial (through abstract classes)	Yes (with traits)
Constructors	Yes	No
Copy/Move constructors/assign.	Manual or Implcit (unsafe)	Yes (safe and fast)
Destructors	Yes	Yes (with traits – via drop)
Friend functions	Yes	Νο
Operator overloading	Yes	Yes
Self consume capacity	No	Yes

TUPLES: Rust supports tuples (just like Python) that can be used to return any multiple data from any function. The same can be achived in C++ using std::pair (but it is not language specific)

Rust
<pre>fn foo(x: i32, y: i32) -> (i32, i32) { (x + y, x * y) }</pre>
<pre>fn main() { let (s,p) = foo(10, 20); }</pre>

C/C++
<pre>#include <iostream> #include <utility></utility></iostream></pre>
<pre>std::pair<int, int=""> foo(int x, int y) { return {x + y, x * y}; }</int,></pre>
<pre>int main() { auto [s, p] = foo(10, 20); }</pre>

LAMBDA FUNCTIONS: In Rust you can create a method in a struct that can consume the instance (e.g. make sure that after the execution of that method the instance is destroyed).

Rust	C/C++
<pre>fn main() { let mut x = 1; let print_x = { println!("x={}",x); <u>x+=</u>1; }; println!("x from main = {}",x); print_x(); }</pre>	<pre>int main() { int x = 1; auto print_x = [&x]() { std::cout << "x=" << x</pre>
Won't compile → because is is a mutable reference in a lambda (can not be used as a immutable reference in main)	<pre>Compiles OK ! (even if x is shared)</pre>

ITERATOR CHAIN: Rust allows the result of one iterator to be the input of another one. Using this techniques sequences of data can be easily converted / process.

Rust	<i>C/C++</i>
<pre>fn main() { let a = vec![1,2,3,4,5,6,7,8,9]; let s:i32 = a.iter() .skip(3) .take(4) .inspect(x println!{"{:?}",*x}) .sum(); println!("sum is {}",s); }</pre>	<pre>int main() { std::vector<int> a = {1,2,3,4,5,6,7,8,9}; auto begin = std::next(a.begin(), 3); auto end = std::next(begin, 4); int sum = 0; for (auto it = begin; it != end; ++it) { std::cout << *it << "\n"; sum += *it; } std::cout << "sum is " << sum << "\n";</int></pre>



Rust vs C++ (Data structures)

Feature	C++	Rust
Tuples	Yes (via std::pair)	Implicit (as part of the language)
Type alias	Yes (typedef & using)	Yes
New type idiom	Partial (through wrapper classes)	Yes
Tree structures	Yes	Yes
Graph (maybe bi-oriented), Double linked list	Yes	Complicated (due to ownership requirements)
Vectors, Maps, Sets	Yes (from std)	Yes (implicit - part of the Rust core)
Iterator chaining	Partial (but supported)	Yes
Templates/Generics	Yes (With a limited constrains)	Yes (with a complex constrain)
Smart pointers	Not implicit (you need to use std)	Implicit (as part of the language)
Lambda	Yes (but unsafe)	Yes (safe)

Compiles but prints 8

(8 = 1 + 2*3 + 1 = 1+6+1)

MACROS: In Rust , macros are evaluated over the AST (Abstract Syntax Tree) and as such avoid several pitfalls that come with a simple subsitusion. Parameters must be explained (e.g. expression, const, etc)

```
Rust

macro_rules! multiply {
    ($p1: expr , $p2: expr) => {
        $p1 * $p2
      };
}
fn main() {
    let z = multiply!(1 + 2, 3 + 1);
    println!("{z}");
}

Kuther

C/C++

#include <iostream>

#define MULTIPLY(x,y) x*y

void main() {
    let z = MULTIPLY(1+2,3+1);
    printf("%d,%d,%d",x,y,z);
}
```

Compiles OK (and prints 12)

RECURSIVE MACROS: Rust macros support recursion. There are some limitations on how many such calls the compiler will do, but they can be overwritten by the use of "256"] (with different values).

Rust macro rules! count { **Compiles OK** () => {0usize}; (\$first:tt) => {1usize}; (\$first:tt,\$(\$tail:tt),*)=> { 1usize + count!(\$(\$tail),*) fn main() { let x = count!(1,2,3,4,5);println!("{x}");

PROCEDURAL MACROS: In Rust , you can write a form of compiler

extension that creates Rust code that will further be compiled and executed



Rust (resulted code)

```
#[derive(Copy, Clone, Debug)]
pub struct Test_16Bits { value: u16 }
impl Test 16Bits {
    pub const V1: Test 16Bits = Test 16Bits { value: 0x1u16 };
    pub const V2: Test_16Bits = Test_16Bits { value: 0x2u16 };
    pub const V3: Test 16Bits = Test 16Bits { value: 0x4u16 };
    pub const V4: Test 16Bits = Test 16Bits { value: 0x8000u16 };
    pub const None: Test 16Bits = Test 16Bits { value: 0 };
    pub fn contains(&self, obj: Test_16Bits) -> bool {...}
    pub fn contains one(&self, obj: Test 16Bits) -> bool {...}
    pub fn is empty(&self) -> bool {...}
    pub fn clear(&mut self) {...}
    pub fn remove(&mut self, obj: Test_16Bits) {...}
    pub fn set(&mut self, obj: Test 16Bits) {...}
    pub fn get value(&self) -> u16 {...}
impl std::ops::BitOr for Test_16Bits {...}
impl std::ops::BitOrAssign for Test_16Bits {...}
impl std::ops::BitAnd for Test_16Bits {...}
impl std::ops::BitAndAssign for Test 16Bits {...}
impl std::cmp::PartialEq for Test_16Bits {...}
impl std::default::Default for Test 16Bits {...}
impl std::fmt::Display for Test 16Bits {...}
```



Rust vs C++ (Macros)

Feature	C++	Rust
Туре	Preprocessor (text substitution)	Hygienic (work at AST level)
Compile-time execution	Νο	Yes (expend into valid Rust code)
Hygiene (name safety)	No	Yes (prevents name collisions)
Debug support	Νο	Yes (compiler aware)
Error messages	Poor — errors reported after expansion	Excellent — errors often point to macro definition
Parameter parsing	Very weak — no parsing, just token substitution	Strong pattern matching and repetition with \$() syntax
Use cases	Constants, small functions, platform-specific code	DSLs, declarative code, boilerplate reduction, metaprogramming
Procedural macros (compiler extensions)	Νο	Yes (powerful support)
Recursive calls	Νο	Yes (configurable recursive depth)

DATA RACES: Rust can statically (at compile time) analyze cases where a reference is used in multiple threads (that could lead to data races) and not allow them.

Rust	C/C++
use std::rc::Rc; use std::thread;	<pre>#include <thread> int main() { int data = 42:</thread></pre>
<pre>fn main() { let data = Rc::new(42);</pre>	<pre>std::thread handle([&data]() { std::cout << "Value = "</pre>
<pre>let h = thread::spawn(move { println!("Value = {}", data); });</pre>	<pre> data << std::endl; }); </pre>
<pre>h.join().unwrap(); }</pre>	<pre>handle.join(); }</pre>
Won't compile \rightarrow because data can not be moved from the main thread to another thread)	Compiles and Run ! (but allows access to a stack variable from a different thread)

Rust Channels

A channel is MPSC (Multiple Producers Single Consumer) mechanism that can be used to send data from one thread to another.



Rust Channels

Using channels (example)

Rust (main)

```
use std::sync::mpsc;
struct Message { id: u32, text: String }
impl Message { fn new(id: u32, text: &str) -> Self {...} }
fn main() {
     let mut v = Vec::new();
           let (producer, receiver) = mpsc::channel::<Message>();
v.push(std::thread::spawn(move || loop {...}));
           for i in 0..5 {
                 let th_producer = producer.clone();
                 v.push(std::thread::spawn(move |[´{...}));
     for th in v { th.join().unwrap(); }
println!("All threads finished");
```

Output (possible)

Received : Hello from 0 Received : Hello from 2 Received : Hello from 2 Received : Hello from 3 Received : Hello from 3 Received : Hello from 4 Received : Hello from 4 Received : Hello from 4 Received : Hello from 1 Terminate the receiver thread All threads finished



Rust vs C++ (Synchronization)

Feature	C++	Rust
Data races	Possible and undefined behavior	Not possible (safe)
Mutex & Locks	Yes	Yes
Atomic types	Yes	Yes
Conditional variables	Yes	Yes
Channels	Νο	Yes (part of the standard library)
Async/await syntax	No (maybe in C++26)	Yes
Async runtimes	Νο	Yes (tokio, async-std, etc)
Green threads	Νο	Yes - via async runtimes (non- blocking tasks, cooperative scheduling)



Conditional Compilation

Let's see a simple example (with code and cargo.toml).



There are two options to run this program:

1. Run "cargo run --features METHOD_A"

```
C++ (equivalent code)
#ifdef METHOD A
    void foo() {
        printf("Method A");
#endif
#ifdef METHOD B
    void foo() {
        printf("Method B");
#endif
void main() {
    foo();
```

Rust unit tests

Rust allows to write unit tests within your own codebase that validate that your code runs as expected

Rust (main.rs)

```
pub fn add(x: u8, y: u8) -> Option<u8> {
    let result = (x as u32) + (y as u32);
    if result > 255 {
        return None;
        }
        return Some(result as u8);
    }
    #[test]
    fn check_add() {
        assert_eq!(add(100,155),Some(255));
        assert_ne!(add(0,0),Some(1));
        }
    #[test]
    fn check_overflow() {
        assert_eq!(add(200,200),None);
        assert_ne!(add(100,100),None);
        }
    }
```



Rust Documentation (Function/Module)

To write a documentation, use <mark>///</mark> characters (on multiple lines) to explain (in Markdown format what that function / module is doing).

Rust

```
Divides `x` to `y`. If `y` is 0 than it returns None,
    otherwise it returns Some(x/y)
                                                  ▶ Run | Debug
    # Example
                                                  fn main() {
                                                      let x = div()
                                                              expected 2 arguments, found 0 rust-analyzer(E0107)
    if let Some(result) = div(5/2) {
        println!("Result is {result}");
                                                              rust tester
      else {
                                                              fn div(x: i32, y: i32) -> Option<i32>
        println!("Division by 0");
                                                              Divides x to y. If y is 0 than it returns None, otherwise it returns Some(x/y)
fn div(x: i32, y: i32) -> Option<i32> {
                                                              Example
    if y != 0 {
                                                              if let Some(result) = div(5/2) {
         Some(x / y)
                                                                  println!("Result is {result}");
    } else {
                                                              } else {
         None
                                                                  println!("Division by 0");
```



Rust vs C++ (CI/CD)

Feature	C++	Rust
Build System	Fragmented (Make, CMake, Ninja, etc.)	Unified: [cargo build]
Cross Platform Builds	Νο	Yes
Crate Publishing	Νο	Yes (crates.io)
Version pinning	No enforcements	Strictly versioned via cargo file
Doc generation	Not implicit (3 rd party tools)	Yes, auto generated (rustdoc)
Testing documentation	Νο	Yes, automatic via cargo test
Hosted documentation	No (manually hosted / updated)	Yes, on docs.rs
Unit testing	Not implicit (only via 3 rd party tools)	Yes, implicit, via #[test]
Code coverage	Not implicit (only via 3 rd party tools)	Yes, implicit, via cargo tarpaulin
Benchmarking	Not implicit (only via 3 rd party tools)	Built-in but nightly #[bench]
Linter support / Static analysis	3 rd party tools (clang-tidy, cppcheck)	Yes, implicit, via cargo clippy
Code formatting	clang-format (needs config per project)	Yes, implicit, via cargo fmt



Optimizations



Rust has several optimizations for mutable references since at one moment of time there could be only one mutable reference.

Let's consider the following C/C++ code and its Rust equivalent:





Rust has several optimizations for mutable references since at one moment of time there could Why do we have this difference? ference.

Let's consider the following C/C++ code and its Rust equivalent:





Rust has several optimizations for mutable references since at one moment of time there could Why do we have this difference? ference.

Let's consider the following C/C++ code and its Rust equivalent:

<i>C/C++</i>	Rust
<pre>void foo(const unsigned int * *output += *input; *output += *input; }</pre>	However, <u>there is no guarantee</u> that the output pointer can't be access from a different thread. As such, the compiler has to write the new value to output pointer so that if another thread is trying to read it, it will read a correct value. This also means that it has to perform a similar write for the second operation !
<pre>mov eax, dword ptr [rsi] add eax, dword ptr [rdi] mov dword ptr [rsi], eax add eax, dword ptr [rdi] mov dword ptr [rsi], eax</pre>	<pre>mov eax, dword ptr [rdi] add eax, eax add dword ptr [rsi], eax</pre>



Rust has several optimizations for mutable references since at one moment of time there could be only one mutable reference.

C/C
 On the other hand, Rust knows that since output is a mutable reference, there is
 only one such reference and no other thread can access it. Because of this, it
 does not have to write the value after the first call to += operator. Furthermore,
 since there can not be mutable reference towards the input variable (as there is already an immutable one), it can reuse its value for the second operation.



mov	<pre>eax, dword ptr [rsi]</pre>
add	eax, dword ptr [rdi]
mov	dword ptr [rsi], eax
add	eax, dword ptr [rdi]
mov	dword ptr [rsi], eax

IOV	eax, dword ptr [rdi]
dd	eax, eax
dd	dword ptr [rsi], eax



Map comparation between C++ and Rust

Let's compare how various types of maps work on Rust and C++.

For this we will use:

- std::map (C++)
- std::unordered_map (C++)
- HashMap (Rust)
- BTreeMap (Rust)

The same algorithm will be written in both Rust and C++ and tested in Debug and Release mode. We will use GetTickCount API to measure time. Each variation of the build will be executed for 10 times and the average will be compute.

So ... lets see the testing algorithm:

```
C++
Rust
extern "system" { fn GetTickCount64() -> u64; }
fn get time() -> u64 { unsafe { GetTickCount64() } }
use std::collections::{BTreeMap, HashMap};
#[derive(Debug, Copy, Clone)]
struct Test { v1: u64, v2: f32, v3: bool }
                                                              };
fn main() {
    let mut m: HashMap<u32, Test> = HashMap::new();
    let start = get time();
    for i in 0..1 000 000 {
        let t = Test { v1: i as u64, v2: 1.5,v3: i%2==0};
       m.insert(i, t);
    let end = get_time();
    println!("{}", end - start);
```

```
#include <Windows.h>
#include <map>
#include <unordered map>
struct Test {
   unsigned long long v1;
   float v2;
   bool v3;
void main() {
    std::unordered map<unsigned int, Test> m;
    auto start = GetTickCount64();
   for (auto i = 0; i < 1000000; i++) {
        m[i] = Test{ (unsigned long long)i,
                      1.5, i % 2 == 0 };
    auto end = GetTickCount64();
   printf("%d", (int)(end - start));
```

So ... lets see the testing algorithm:

<pre>extern "system" { fn GetTickCount64() -> u64; } fn get_time() -> u64 { unsafe { GetTickCount64() } } use { We will run the same algorithm using: We will run the same algorithm using: We will run the same algorithm using:</pre> #include <windows.h> #include <map> time() We will run the same algorithm using:</map></windows.h>	Rust	C++
We will run the same algorithm using:	<pre>extern "system" { fn GetTickCount64() -> u64; } fn get_time() -> u64 { unsafe { GetTickCount64() } }</pre>	<pre>#include <windows.h> #include <map> #include <unordered map=""></unordered></map></windows.h></pre>
#[der • HashMap struc • BTreeMap b • std::map	 We will run the same algorithm using: #[der • HashMap struc • BTreeMap 	<pre>struc U We will run the same algorithm using: f • std::unordered_map b • std::map</pre>
<pre>fn main() { let mut m: HashMap<u32, test=""> = HashMap::new(); let start = get_time(); for i in 01_000_000 { let t = Test { v1: i as u64, v2: 1.5,v3: i%2==0}; m.insert(i, t); } let end = get_time(); println!("{}", end - start); }; </u32,></pre>	<pre>fn main() { let mut m: HashMap<u32, test=""> = HashMap::new(); let start = get_time(); for i in 01_000_000 { let t = Test { v1: i as u64, v2: 1.5,v3: i%2==0}; m.insert(i, t); } let end = get_time(); println!("{}", end - start); }</u32,></pre>	<pre>}; void main() { std::unordered_map<unsigned int,="" test=""> m; auto start = GetTickCount64(); for (auto i = 0; i < 1000000; i++) { m[i] = Test{ (unsigned long long)i,</unsigned></pre>

So ... lets see the testing algorithm:

	#1	#2	#3	#4	#5	#6	#7	#8	#9	#10	Average
C++ (Debug) std:unordered_map	938	1266	1390	1328	1250	1297	1250	1344	1485	1437	1298
C++ (Release) std:unordered_map	297	234	235	281	266	266	234	235	234	234	251
C++ (Debug) std::map	1312	1765	1953	1875	1875	1859	1813	1812	1797	1828	1788
C++ (Release) std::map	156	141	172	157	141	171	172	156	172	156	159
Rust (Debug) HashMap	1141	1297	1265	1250	1281	1312	1359	1297	1343	1297	1284
Rust (Release) HashMap	78	78	63	78	78	94	93	94	94	94	84
Rust (Debug) BTreeMap	2703	3156	3078	2906	2765	2875	2937	2844	2860	2781	2890
Rust (Release) BTreeMap	93	93	109	110	125	110	125	141	125	125	115
C++ vs Rust (on maps)

The general conclusion after these tests is:

- Rust is slower the C++ when it comes to debug mode (due to many checks)
- In terms of Release mode, Rust is faster (however, it should be noted that we are not comparing the same algorithms and as such these tests might NOT be correct). However, since we've compared the standard algorithms from each (Rust and C++) libraries, the results are however relevant.
- The tests were performed on Windows 11 (using Microsoft compiler). To produce accurate results, other C++ compilers (such as clang and gcc) should be tested as well.



Vector comparation between C++ and Rust

Let's compare how efficient vector push method is for both C++ and Rust.

Rust	C++
extern "system" {	<pre>#include <windows.h></windows.h></pre>
<pre>fn GetTickCount64 () -> u64;</pre>	<pre>#include <vector></vector></pre>
}	<pre>struct Test {</pre>
fn get_time () -> u64 {	int v1;
<pre>unsafe { GetTickCount64() }</pre>	float v2;
}	char32_t v3;
<pre>#[derive(Debug,Copy,Clone)]</pre>	uint8_t v4[256];
<pre>struct Test { v1: i32, v2: f32, v3: char, v4: [u8;256] }</pre>	};
<pre>fn main() {</pre>	<pre>void main() {</pre>
let mut <u>v</u> : Vec <test> = Vec::new();</test>	<pre>std::vector<test> v;</test></pre>
<pre>let t = Test{v1:5,v2:1.3,v3:'A',v4:[48u8;256]};</pre>	Test t;
<pre>let start = get_time();</pre>	<pre>auto start = GetTickCount64();</pre>
for i in 010_000_000 {	for (auto i = 0; i < 10000000; i++) {
<u>v</u> . <u>push</u> (t);	v.push_back(t);
}	}
<pre>let end = get_time();</pre>	<pre>auto end = GetTickCount64();</pre>
<pre>println!("{}",end-start);</pre>	<pre>printf("%d", (int)(end - start));</pre>



Both codes were teste in the same environment, for 10 times and the average was recorded. All tests were run on x64 architecture (Debug and Release). Times are measures in milliseconds.

Keep in mind that **GetTickCount** function has an error margin of 16ms.

	#1	#2	#3	#4	#5	#6	#7	#8	#9	#10	Average
C++ (Debug)	2562	2562	2640	2578	2578	2515	2562	2578	2562	2547	2568
Rust (Debug)	1922	1844	1860	1843	1812	1813	1797	1781	1828	1813	1831
C++ (Release)	1828	1781	1797	1781	1781	1781	1797	1797	1821	1797	1796
Rust (Release)	1750	1750	1688	1719	1687	1703	1719	1687	1703	1687	1709



As a general conclusion, when it comes to vectors (and copying object not moving them), Rust is faster than C/C++ (in both debug and release modes).

We should point out that the build that was tested for C++ was compiled with Microsoft compiler (cl.exe) and it does not reflect results for gcc or clang (that might optimize the C++ code in a different way).

However, the question still remains on what's different in Rust vs C++ in terms of how vector works ?

So ... the difference lies in how growth algorithm works for those two cases (Rust and C++).



So ... lets see the behavior if we reserve the memory from the start.

```
C++
Rust
                                                               #include <Windows.h>
extern "system" {
                                                               #include <vector>
         fn GetTickCount64 () -> u64;
                                                               struct Test {
                                                                  int v1;
fn get time () -> u64 {
                                                                  float v2;
         unsafe { GetTickCount64() }
                                                                   char32 t v3;
                                                                  uint8 t v4[256];
#[derive(Debug,Copy,Clone)]
                                                               };
struct Test { v1: i32, v2: f32, v3: char, v4: [u8;256] }
                                                               void main() {
fn main() {
                                                                   std::vector<Test> v;
   let mut v: Vec<Test> = Vec::with_capacity(10_000_000);
                                                                  Test t:
   let t = Test{v1:5,v2:1.3,v3: 'A',v4:|48u8;256|};
                                                                  v.reserve(10000000);
   let start = get time();
                                                                   auto start = GetTickCount64();
    for i in 0..10 000 000 {
                                                                   for (auto i = 0; i < 10000000; i++) {
       v.push(t);
                                                                       v.push back(t);
    let end = get_time();
                                                                   auto end = GetTickCount64();
    println!("{}",end-start);
                                                                   printf("%d", (int)(end - start));
```



Test were performed in a similar manner like the previous ones (Debug and Release, 10 iterations and we compute the average).

	#1	#2	#3	#4	#5	#6	#7	#8	#9	#10	Average
C++ (Debug)	782	766	781	766	781	766	797	797	828	781	784
Rust (Debug)	984	1094	1063	1031	875	1032	1016	860	906	859	972
C++ (Release)	547	532	531	515	500	515	516	500	531	531	521
Rust (Release)	532	500	531	516	562	531	547	547	547	515	532

Keep in mind that there is an error margin of 16 ms for GetTickCount API. This means that the difference between C++ and Rust is insignificant (we can consider both at the same level).



Other tests (Mandelbrot)

The Mandelbrot test is a benchmarking exercise that measures the performance of a programming language or compiler by computing and rendering the Mandelbrot fractal.



What It Measures:

- Raw CPU performance
- Loop optimization
- Floating-point performance
- Parallelism or threading efficiency
- Compiler code generation quality

Language	Time	Peak Memory	Version
<u>Rust</u>	247ms	4.9MB	rustc 1.89.0-nightly
<u>Rust</u>	291ms	4.8MB	rustc 1.87.0
<u>C</u>	332ms	6.0MB	zigcc 0.14.1
<u>C-Sharp</u>	332ms	37.3MB	dotnet 9.0.300
<u>C</u>	452ms	6.5MB	clang 14.0.0-1ubuntu1.1
<u>C</u>	543ms	6.6MB	gcc 15.1.0
<u>Nim</u>	578ms	4.5MB	nim 2.2.4
<u>Java</u>	1162ms	55.7MB	openjdk 23
<u>Java</u>	1167ms	54.6MB	openjdk 21
Java	1194ms	108.9MB	graal/jvm 17.0.8
Go	3230ms	7.7MB	go 1.24.3

https://programming-language-benchmarks.vercel.app/problem/mandelbrot



Other tests (Mandelbrot)

The Mandelbrot test is a benchmarking exercise that measures the performance of a programming language or compiler by computing and rendering the Mandelbrot fractal.



What It Measures:

- Raw CPU performance
- Loop optimization
- Floating-point performance
- Parallelism or threading efficiency
- Compiler code generation quality

#	Language	Time	Memory	#	Language	Time	Memory
		(sec)	(MB)			(sec)	(MB)
1	Rust	0.95	35,598	14	Dart	4.29	45,175
2	Chapel	1.18	42,156	15	Haskell	6.64	51,057
3	Julia	1.54	357,138	16	F#	7.17	49,979
4	C gcc	1.64	35,582	17	Swift	7.27	49,312
5	C++ g++	2.36	38,281	18	OCaml	7.60	64,643
6	Intel Fortran	2.72	85,975	19	Erlang	53.86	98,140
7	Go	3.77	37,970	20	РНР	68.29	53,531
8	Free Pascal	3.91	35,529	21	Ruby	143.13	118,436
9	Java	3.96	58,348	22	Lua	159.01	652,796
10	Ada 2012	4.01	41,099	23	Python 3	182.94	62,173
11	C#	4.02	40,743	24	РНР	258.19	16,437
12	Node.js	4.05	144,757	25	Smalltalk	>5 min	175,542
13	Lisp	4.20	60,654	26	Perl	>8 min	114,569

https://benchmarksgame-team.pages.debian.net/benchmarksgame/performance/mandelbrot.html

(code that have possible hand-written vector instructions or "unsafe" was removed)

