

Stable Inline
Assembly

Florob

What?

Why?

Assembly
Primer

Prior Art

gcc-style
DSL-style

History

Stabilized
Syntax

Questions

Stable Inline Assembly

Florian “Florob” Zeitz

2022-06-20

1 What?

2 Why?

3 Assembly Primer

4 Prior Art

5 History

6 Stabilized Syntax

1 What?

2 Why?

3 Assembly Primer

4 Prior Art

5 History

6 Stabilized Syntax

What is Inline Assembly?

Stable Inline
Assembly

Florob

What?

Why?

Assembly
Primer

Prior Art

gcc-style
DSL-style

History

Stabilized
Syntax

Questions

- Handwritten assembly to compile into the final binary
- Written *inline* in function bodies
- Integrates with surrounding code
- Available in some systems programming languages

1 What?

2 Why?

3 Assembly Primer

4 Prior Art

5 History

6 Stabilized Syntax

Why do we need Inline Assembly?

Stable Inline
Assembly

Florob

What?

Why?

Assembly
Primer

Prior Art

gcc-style
DSL-style

History

Stabilized
Syntax

Questions

- low-level control
 - configuration registers
 - hardware not otherwise accessible
- performance
- predictable timing
 - cryptography (constant time)
 - embedded
- convenience
 - no separate assembly file
 - no need for a `build.rs`
- Rust philosophy: safe with **unsafe** `{}` escape hatches where needed

Why do we need Inline Assembly?

Stable Inline
Assembly

Florob

What?

Why?

Assembly
Primer

Prior Art

gcc-style
DSL-style

History

Stabilized
Syntax

Questions

- low-level control
 - configuration registers
 - hardware not otherwise accessible
- performance
- predictable timing
 - cryptography (constant time)
 - embedded
- convenience
 - no separate assembly file
 - no need for a `build.rs`
- Rust philosophy: safe with **unsafe** `{}` escape hatches where needed

Why do we need Inline Assembly?

Stable Inline
Assembly

Florob

What?

Why?

Assembly
Primer

Prior Art

gcc-style
DSL-style

History

Stabilized
Syntax

Questions

- low-level control
 - configuration registers
 - hardware not otherwise accessible
- performance
- predictable timing
 - cryptography (constant time)
 - embedded
- convenience
 - no separate assembly file
 - no need for a `build.rs`
- Rust philosophy: safe with **unsafe** `{}` escape hatches where needed

Why do we need Inline Assembly?

Stable Inline
Assembly

Florob

What?

Why?

Assembly
Primer

Prior Art

gcc-style
DSL-style

History

Stabilized
Syntax

Questions

- low-level control
 - configuration registers
 - hardware not otherwise accessible
- performance
- predictable timing
 - cryptography (constant time)
 - embedded
- convenience
 - no separate assembly file
 - no need for a `build.rs`
- Rust philosophy: safe with `unsafe {}` escape hatches where needed

Why do we need Inline Assembly?

Stable Inline
Assembly

Florob

What?

Why?

Assembly
Primer

Prior Art

gcc-style
DSL-style

History

Stabilized
Syntax

Questions

- low-level control
 - configuration registers
 - hardware not otherwise accessible
- performance
- predictable timing
 - cryptography (constant time)
 - embedded
- convenience
 - no separate assembly file
 - no need for a `build.rs`
- Rust philosophy: safe with **unsafe** `{}` escape hatches where needed

Why not use intrinsics?

Stable Inline
Assembly

Florob

What?

Why?

Assembly
Primer

Prior Art

gcc-style
DSL-style

History

Stabilized
Syntax

Questions

- “Just” using an intrinsic per instruction has often been suggested
- Can be re-ordered
- Instructions can be inserted in between
- Hard to scale
- Unreasonable for rare use-cases

However: Definitely use them when applicable

1 What?

2 Why?

3 Assembly Primer

4 Prior Art

5 History

6 Stabilized Syntax

Assembly Primer

Stable Inline
Assembly

Florob

What?

Why?

Assembly
Primer

Prior Art

gcc-style

DSL-style

History

Stabilized
Syntax

Questions

- Textual representation of CPU instructions
- Usually a mnemonic followed by operands
- Modifies registers, memory, flags, ...
- Labels to jump to
- Immediate values

Examples

Stable Inline
Assembly

Florob

What?

Why?

Assembly
Primer

Prior Art
gcc-style
DSL-style

History

Stabilized
Syntax

Questions

x86 (AT&T)

```
1 jmp foo
2 mov $5, %eax
3 foo:
4 mov 5(%eax, %ebx, 2), %edx
```

ARM

```
1 ldmiaeq sp!, {r4-r7, pc}
2 ldrb r0, [sp, #32]
```

x86 (Intel)

```
1 jmp foo
2 mov eax, 5
3 foo:
4 mov edx, [eax + ebx*2 + 5]
```

PowerPC

```
1 lwz r0, 8(r1)
2 rlwimi r11, r1, 0, 30, 28
```

1 What?

2 Why?

3 Assembly Primer

4 Prior Art

5 History

6 Stabilized Syntax

Prior Art

Stable Inline
Assembly

Florob

What?

Why?

Assembly
Primer

Prior Art

gcc-style
DSL-style

History

Stabilized
Syntax

Questions

- Available in some systems programming languages
- Some use a template style (e. g. GNU-C, LLVM-IR, ldc)
- Some use an embedded DSL (e. g. D, MSVC)

1 What?

2 Why?

3 Assembly Primer

4 Prior Art

- gcc-style
- DSL-style

5 History

6 Stabilized Syntax

Example: gcc

Stable Inline
Assembly

Florob

What?

Why?

Assembly
Primer

Prior Art

gcc-style

DSL-style

History

Stabilized
Syntax

Questions

- 1 `asm [volatile] ("template" : outs : ins : clobber)`
 - Uses a template strings
 - Compiler fills template based on constraints
 - Template result interpreted by the assembler (compiler doesn't have to know instructions)

gcc Template

Stable Inline
Assembly

Florob

What?

Why?

Assembly
Primer

Prior Art

gcc-style

DSL-style

History

Stabilized
Syntax

Questions

- `%0`, or `%[name]` as placeholders
- Escapes: `%%`, `%{`, `%}`, `%|`

gcc Arguments

Stable Inline
Assembly

Florob

What?

Why?

Assembly
Primer

Prior Art

gcc-style
DSL-style

History

Stabilized
Syntax

Questions

- Specify outputs, inputs, and clobbers
- Constraints determine how to fill the template (register, address, immediate)
- All inputs have to be read before any output is written
- Clobbers specify touched state, that can not be inferred from inputs and outputs
- Compiler moves arguments into registers, saves clobbered registers, knows to reload changed memory

gcc Constraints

Stable Inline
Assembly

Florob

What?

Why?

Assembly
Primer

Prior Art

gcc-style

DSL-style

History

Stabilized
Syntax

Questions

- Constraint codes are usually a single letter
- Output
 - start with = or + (input and output)
 - can be modified with & to allow write before all inputs are read (early-clobber)
- Input
 - can be tied to outputs by specifying the operand number as constrained
- Clobbers
 - Used registers
 - Modified flags
 - If memory was modified

gcc Examples: Add 5 to variable

Stable Inline
Assembly

Florob

What?

Why?

Assembly
Primer

Prior Art

gcc-style
DSL-style

History

Stabilized
Syntax

Questions

```
1 int var = 0;  
2 asm ("add $5, %0" : "+r"(var));
```

gcc Examples: Get L1 cache size

Stable Inline
Assembly

Florob

What?

Why?

Assembly
Primer

Prior Art

gcc-style
DSL-style

History

Stabilized
Syntax

Questions

```
1  int ebx, ecx;
2  asm (
3      "mov $4, %%eax;"
4      "xor %%ecx, %%ecx;"
5      "cpuid;"
6      "mov %%ebx, %0;"
7      : "=r"(ebx), "=c"(ecx)
8      :
9      : "eax", "ebx", "edx"
10 );
11 printf("L1 Cache: %i\n", ((ebx >> 22) + 1)
12      * (((ebx >> 12) & 0x3ff) + 1)
13      * ((ebx & 0xfff) + 1)
14      * (ecx + 1));
```

gcc Examples: Increment

Stable Inline
Assembly

Florob

What?

Why?

Assembly
Primer

Prior Art

gcc-style

DSL-style

History

Stabilized
Syntax

Questions

```
1  int in = 3, out;  
2  asm (  
3      "mov $1, %0\r\n"  
4      "add %1, %0\r\n"  
5      : "=r"(out)  
6      : "r"(in)  
7  );
```

Is this correct?

gcc Examples: Increment

Stable Inline
Assembly

Florob

What?

Why?

Assembly
Primer

Prior Art

gcc-style

DSL-style

History

Stabilized
Syntax

Questions

```
1  int in = 3, out;  
2  asm (  
3      "mov $1, %0\r\n"  
4      "add %1, %0\r\n"  
5      : "=r"(out)  
6      : "r"(in)  
7  );
```

Is this correct? No, in and out could be in the same register.

gcc-style Pros/Cons

Stable Inline
Assembly

Florob

What?

Why?

Assembly
Primer

Prior Art

gcc-style
DSL-style

History

Stabilized
Syntax

Questions

Pros

- Existing implementations support all architectures
- Arbitrary assembly supported
- Can be optimized

Cons

- Hard to write
- Easy to get constraints wrong
- Missing constraints lead to undefined behavior

1 What?

2 Why?

3 Assembly Primer

4 Prior Art

- gcc-style
- DSL-style

5 History

6 Stabilized Syntax

Example: D

Stable Inline
Assembly

Florob

What?

Why?

Assembly
Primer

Prior Art

gcc-style
DSL-style

History

Stabilized
Syntax

Questions

```
1 int var = 0;  
2 asm {  
3     mov RAX, RDX;  
4     add var, 3 + 2;  
5     add var, RAX;  
6 }
```

- Implemented as DSL inside `asm {}`
- Standardized per CPU family
- Many D expressions allowed

DSL-style Pros/Cons

Stable Inline
Assembly

Florob

What?

Why?

Assembly
Primer

Prior Art

gcc-style
DSL-style

History

Stabilized
Syntax

Questions

Pros

- Easy to write
- Tight integration with surrounding code
- Can be optimized
- Impossible to forget constraints

Cons

- Much work to support a new architecture
- Existing implementations only support x86(_64)
- Existing implementations don't support some instructions and registers
- Neither assembly nor language syntax

Inline Assembly 2017

Stable Inline Assembly

Florob

What?

Why?

Assembly Primer

Prior Art

gcc-style

DSL-style

History

Stabilized Syntax

Questions

Inline Assembly

Florob



Why?

- low-level control
 - config registers
 - hardware not otherwise accessible
- performance
 - predictable timing
 - embedded
 - cryptography (side-channel)
- convenience
 - no separate assembly file
 - no need for a `build.rs`
- Rust philosophy: safe with `unsafe {}` escape hatches where needed

Playlists: rustcgb's videos starting here / audio / related events

90 min

2017-06-07

2017-06-13

312

Fahrplan

Rust's support for inline assembly is currently unstable. The talk will give an overview of the current unstable implementation in Rust, as well as inline assembly support available in other programming languages. This will lead into a discussion about a sensible design for this feature and a way to stabilize this eventually.

<https://media.ccc.de/v/rustcgn.2017.06.inline-assembly>

1 What?

2 Why?

3 Assembly Primer

4 Prior Art

5 History

6 Stabilized Syntax

Global ASM

Stable Inline
Assembly

Florob

What?

Why?

Assembly
Primer

Prior Art

gcc-style
DSL-style

History

Stabilized
Syntax

Questions

- Stable since Rust 1.59, but specified since 2016
- Include arbitrary assembly in the `.text` section
- Only in “global scope” (outside functions)
- Replaces linking to externally compiled assembly
- Can use `#[cfg()]` directives

Global ASM

Stable Inline
Assembly

Florob

What?

Why?

Assembly
Primer

Prior Art

gcc-style
DSL-style

History

Stabilized
Syntax

Questions

```
1  std::arch::global_asm!(r#"
2      .globl add_5
3      add_5:
4          mov rax, rdi
5          add rax, 5
6          ret
7  "#);
8
9  extern {
10     fn add_5(n: u64) -> u64;
11 }
12
13 fn main() {
14     println!("{}", unsafe { add_5(37) });
15 }
```

llvm_asm!()

Stable Inline
Assembly

Florob

What?

Why?

Assembly
Primer

Prior Art

gcc-style
DSL-style

History

Stabilized
Syntax

Questions

```
1  unsafe {  
2      llvm_asm!("template"  
3          : outs  
4          : ins  
5          : clobbers  
6          : options  
7      );  
8  }
```

- Used to be `asm!()`
- Straight binding to LLVM IR (gcc-style)
- Ever so slightly different from both LLVM IR and GNU-C
- Renamed, deprecated, now removed

llvm_asm!()

Stable Inline
Assembly

Florob

What?

Why?

Assembly
Primer

Prior Art

gcc-style
DSL-style

History

Stabilized
Syntax

Questions

```
1 let mut var = 0;  
2 unsafe {  
3     llvm_asm!("add $$5, $0" : "+r"(var));  
4 }
```

- No good route to stabilization
- Had various known bugs
- Hard to support on non-LLVM backends

(Pre-)RFCs

Stable Inline
Assembly

Florob

2014 “RFC 129: refine the asm! extension” by pczarn

What?

2018 “[Pre-RFC]: Inline assembly” by Florob

Why?

2019 “[Pre-RFC #2]: Inline assembly” by Amanieu

Assembly
Primer

2019 “RFC 2836: Introduce the ASM project group” by Amanieu

Prior Art

gcc-style
DSL-style

2020 “RFC 2850: Inline assembly” by Amanieu

History

2020 “RFC 2843: Add llvm_asm! and deprecate asm!” by Amanieu
(merged)

Stabilized
Syntax

Questions

2020 “RFC 2873 Inline assembly” by Amanieu (merged)

1 What?

2 Why?

3 Assembly Primer

4 Prior Art

5 History

6 Stabilized Syntax

Design Goals (of the initial Pre-RFC)

Stable Inline
Assembly

Florob

What?

Why?

Assembly
Primer

Prior Art

gcc-style
DSL-style

History

Stabilized
Syntax

Questions

- Easy to use
- Portable
 - Between platforms (x86, AMD64, ARM, AVR, PowerPC, RISC-V, ...)
 - Between backends (LLVM, Cretone, gcc, ...)
- Support all instructions (with all operand types)

asm!()

Stable Inline
Assembly

Florob

What?

Why?

Assembly
Primer

Prior Art

gcc-style
DSL-style

History

Stabilized
Syntax

Questions

- 1 `asm!("template", operands, options);`
 - Stabilized in Rust 1.59
 - Fundamentally gcc-style
 - Easier syntax than gcc
 - Safer defaults than gcc

Template

Stable Inline
Assembly

Florob

What?

Why?

Assembly
Primer

Prior Art

gcc-style
DSL-style

History

Stabilized
Syntax

Questions

- Designed after Rust format strings
- Multiple template arguments allowed, each representing one line
- `{}` or `{name}` as placeholder
- Modifiers via `{:h}`
- Escapes: `{{, }}`

Arguments

Stable Inline
Assembly

Florob

What?

Why?

Assembly
Primer

Prior Art

gcc-style
DSL-style

History

Stabilized
Syntax

Questions

- Outputs and inputs, in any order
- Clobbers are just ignored outputs
- Inputs can be read before outputs

Constraints

Stable Inline
Assembly

Florob

What?

Why?

Assembly
Primer

Prior Art

gcc-style
DSL-style

History

Stabilized
Syntax

Questions

- Register classes (e. g. `reg`, `xmm_reg`)
- Specific registers as string literals (e. g. `"eax"`)
- No memory, supported by using an address in a register

Example: Outputs

Stable Inline
Assembly

Florob

What?

Why?

Assembly
Primer

Prior Art

gcc-style
DSL-style

History

Stabilized
Syntax

Questions

```
1 use std::arch::asm;  
2  
3 let x: u64;  
4 unsafe {  
5     asm!("mov {}, 42", out(reg) x);  
6 }
```

Example: Inputs

Stable Inline
Assembly

Florob

What?

Why?

Assembly
Primer

Prior Art

gcc-style
DSL-style

History

Stabilized
Syntax

Questions

```
1  use std::arch::asm;
2
3  let input: u64 = 3;
4  let output: u64;
5  unsafe {
6      asm!(
7          "mov {0}, 1",
8          "add {0}, {1}",
9          out(reg) output,
10         in(reg) input,
11         );
12 }
```

Is this correct?

Example: Inputs

Stable Inline
Assembly

Florob

What?

Why?

Assembly
Primer

Prior Art

gcc-style
DSL-style

History

Stabilized
Syntax

Questions

```
1 use std::arch::asm;
2
3 let input: u64 = 3;
4 let output: u64;
5 unsafe {
6     asm!(
7         "mov {0}, 1",
8         "add {0}, {1}",
9         out(reg) output,
10        in(reg) input,
11        );
12 }
```

Is this correct? Yes, outputs are early-clobber by default

Example: Late Outputs

Stable Inline
Assembly

Florob

What?

Why?

Assembly
Primer

Prior Art

gcc-style
DSL-style

History

Stabilized
Syntax

Questions

```
1 use std::arch::asm;
2
3 let input: u64 = 3;
4 let output: u64;
5 unsafe {
6     asm!(
7         "mov {0}, {1}",
8         "add {0}, 1",
9         lateout(reg) output,
10        in(reg) input,
11    );
12 }
```

- Outputs that may only be written after all inputs are read
- Allows for more optimized code

Example: Combined Input and Output

Stable Inline
Assembly

Florob

What?

Why?

Assembly
Primer

Prior Art

gcc-style
DSL-style

History

Stabilized
Syntax

Questions

```
1 use std::arch::asm;
2
3 let mut val: u64 = 6;
4 let val_plus5: u64;
5 unsafe {
6     asm!("add {}, 1", inout(reg) val);
7
8     asm!(
9         "add {}, {}",
10        inout(reg) 5u64 => val_plus5,
11        in(reg) val,
12    );
13 }
```

Example: Get L1 Cache Size

Stable Inline
Assembly

Florob

What?

Why?

Assembly
Primer

Prior Art

gcc-style
DSL-style

History

Stabilized
Syntax

Questions

```
1  let rbx: u64;
2  let rcx: u64;
3  unsafe { asm!(
4      "push rbx",
5      "cpuid",
6      "mov {scratch}, rbx",
7      "pop rbx",
8      scratch = lateout(reg) rbx, // rbx is reserved by LLVM
9      inlateout("rax") 4u64 => _,
10     inlateout("rcx") 0u64 => rcx,
11     out("rdx") _
12 ) };
13 println!("L1 Cache: {}", ((rbx >> 22) + 1)
14     * (((rbx >> 12) & 0x3ff) + 1)
15     * ((rbx & 0xfff) + 1) * (rcx + 1));
```


Safer defaults

Stable Inline
Assembly

Florob

What?

Why?

Assembly
Primer

Prior Art

gcc-style
DSL-style

History

Stabilized
Syntax

Questions

```
1 let mut a: u64 = 4;  
2 let b: u64 = 4;  
3 unsafe { asm!(  
4     "add {0}, {1}",  
5     inlateout(reg) a, in(reg) b,  
6     options(pure, nomem, nostack),  
7 ) };  
8 assert_eq!(a, 8);
```

- Assumes volatile by default, `option(pure)` to override
- Assumes memory access by default, `option(nomem)` to override
- Assumes stack access by default, `option(nostack)` to override

Goals?

Stable Inline
Assembly

Florob

What?

Why?

Assembly
Primer

Prior Art

gcc-style
DSL-style

History

Stabilized
Syntax

Questions

- Currently supported architectures: x86, AMD64, ARM, AArch64, RISC-V
- Full support only in LLVM backend
- Can fall back to treating inline assembly like a function call

Thank you for your attention.
Any questions?



<https://babelmonkeys.de/~florob/talks/RC-2022-06-20-rust-inline-asm.pdf>