


# Interface Compliance of Inline Assembly:

Automatically Check, Patch and Refine



---

**Frédéric Recoules**

Univ. Paris-Saclay, CEA, List

Sébastien Bardin

Univ. Paris-Saclay, CEA, List

Richard Bonichon

Tweag I/O

Matthieu Lemerre

Univ. Paris-Saclay, CEA, List

Laurent Mounier

Univ. Grenoble Alpes, VERIMAG

Marie-Laure Potet



Univ. Grenoble Alpes, VERIMAG

International Conference on Software Engineering, 2021

# Inline assembly example in C code

```
AO_INLINE int
AO_compare_double_and_swap_double_full(volatile AO_double_t *addr,
                                       AO_t old_val1, AO_t old_val2,
                                       AO_t new_val1, AO_t new_val2)
{
    char result;
    [...]
    __asm__ __volatile__(
        "xchg %%ebx,%6;" /* swap GOT ptr and new_val1 */
        "lock; cmpxchg8b %0; setz %1;"
        "xchg %%ebx,%6;" /* restore ebx and edi */
        : "=m"(*addr), "a"(result)
        : "m"(*addr), "d" (old_val2), "a" (old_val1),
          "c" (new_val2), "D" (new_val1) : "memory");
    [...]
    return (int) result;
}
```

# Inline assembly is well spread



- full access to hardware
- hand-crafted optimization
- security / obfuscation

<sup>a</sup>according to Rigger et al., 2018

**“GCC-style inline assembly is notoriously  
hard to write correctly”**

**Oliver Stannard,  
ARM Senior Software Engineer on llvm threads, 2018**

# Inline assembly example in C code

```
AO_INLINE int
AO_compare_double_and_swap_double_full(volatile AO_double_t *addr,
                                       AO_t old_val1, AO_t old_val2,
                                       AO_t new_val1, AO_t new_val2)
{
    char result;
    [...]
    __asm__ __volatile__(
        "xchg %%ebx,%6;" /* swap GOT ptr and new_val1 */
        "lock; cmpxchg8b %0; setz %1;"
        "xchg %%ebx,%6;" /* restore ebx and edi */
        : "=m"(*addr), "a"(result)
        : "m"(*addr), "d" (old_val2), "a" (old_val1),
          "c" (new_val2), "D" (new_val1) : "memory");
    [...]
    return (int) result;
}
```

# Inline assembly example in C code

```
AO_INLINE int
AO_compare_double_and_swap_double_full(volatile AO_double_t *addr,
                                       AO_t old_val1, AO_t old_val2,
                                       AO_t new_val1, AO_t new_val2)
{
    char result;
    [...]
    __asm__ __volatile__("xchg %%ebx,%6;" /* swap GOT ptr and new_val1 */
                        "lock; cmpxchg8b %0; setz %1;"
                        "xchg %%ebx,%6;" /* restore ebx and edi */
                        : "=m"(*addr), "=a"(result)
                        : "m"(*addr), "d" (old_val2), "a" (old_val1),
                          "c" (new_val2), "D" (new_val1) : "memory");
    [...]
    return (int) result;
}
```

Assembly template

# Inline assembly example in C code

```
AO_INLINE int
AO_compare_double_and_swap_double_full(volatile AO_double_t *addr,
                                       AO_t old_val1, AO_t old_val2,
                                       AO_t new_val1, AO_t new_val2)
{
    char result;
    [...]
    __asm__ __volatile__(
        "xchg %%ebx,%6" /* swap GOT ptr and new_val1 */
        "lock; cmpxchg8b %0 setz %1"
        "xchg %%ebx,%6" /* restore ebx and edi */
        : "=m"(*addr), "=a"(result)
        : "m"(*addr), "d" (old_val2), "a" (old_val1),
          "c" (new_val2), "D" (new_val1) : "memory");
    [...]
    return (int) result;
}
```

Assembly template

# Inline assembly example in C code

```
AO_INLINE int
AO_compare_double_and_swap_double_full(volatile AO_double_t *addr,
                                       AO_t old_val1, AO_t old_val2,
                                       AO_t new_val1, AO_t new_val2)
{
    char result;
    [...]
    __asm__ __volatile__(
        "xchg %%ebx,%6" /* swap GOT ptr and new_val1 */
        "lock; cmpxchg8b %0 setz %1"
        "xchg %%ebx,%6" /* restore ebx and edi */
        : "=m"(*addr), "=a"(result)
        : "m"(*addr), "d" (old_val2), "a" (old_val1),
          "c" (new_val2), "D" (new_val1) : "memory");
    [...]
    return (int) result;
}
```

Assembly template

Output list

Input list

Clobber list



# Inline assembly example in C code

```
AO_INLINE int
AO_compare_double_and_swap_double_full(volatile AO_double_t *addr,
                                       AO_t old_val1, AO_t old_val2,
                                       AO_t new_val1, AO_t new_val2)
{
    char result;
    [...]
    __asm__ __volatile__(
        "xchg %%ebx,%6;" /* swap GOT ptr and new_val1 */
        "lock; cmpxchg8b %0; setz %1;" %eax
        "xchg %%ebx,%6;" /* restore ebx and edi */
        : "=m"(*addr), "=a"(result)
        : "m"(*addr), "d"(old_val2), "a"(old_val1),
          "c"(new_val2), "D"(new_val1) : "memory");
    [...]
    return (int) result;
}
```

Assembly template

Output list

Input list

Clobber list

%ecx

%edi

%edx

This code works fine prior to GCC 5.0,  
then suddenly crashes with a Segmentation fault

- compiler knowledge is limited to the interface
- register allocation and optimizations rely on it
- mismatches code-interface can lead to bugs

# A few known inline assembly bugs 🦋

- `strcspn`  
[glibc](#) – January 1999, commit 7c97add
- `compare_double_and_swap_double`  
[libatomic\\_ops](#) – Mars 2012, commit 30cea1b
- `compare_double_and_swap_double`  
[libatomic\\_ops](#) – September 2012, commit 64d81cd
- `bswap`  
[libtomcrypt](#) – November 2012, commit cefff85

Interface compliance does matter

Today's challenge :  
**Interface Compliance**

**Define**

–

**Check**

–

**Patch**

# Goal & challenges

## Define

must be built on a currently missing proper formalization  
*indeed there is not even a complete documentation...*

## Check, Patch & Refine

must be able to check whether an assembly chunk is compliant  
*ideally, should suggest a patch for the non compliant ones*

## Widely applicable

must be compiler & architecture agnostic



# Our contributions (1/2)

## A **novel semantics** and comprehensive **formalization**

- support GCC, Clang and mostly icc
- **Framing** condition & **Unicity** condition

## A method to **check**, **patch** and **refine** the interface

- dataflow analysis + dedicated optimizations
- infer an over-approximation of the ideal interface

# Our contributions (2/2)

## Thorough experiments of our prototype

- **2.6k<sup>+</sup>** real-world assembly chunks (**Debian**)
- **2183** issues, including **986 severe** issues
- **2000** patches, including **803 severe** fixes
- **7** packages have already accepted the fixes

<https://github.com/binsec/icse2021-artifact992>

DOI [10.5281/zenodo.4601172](https://doi.org/10.5281/zenodo.4601172)



## A study of current inline assembly bad coding practices

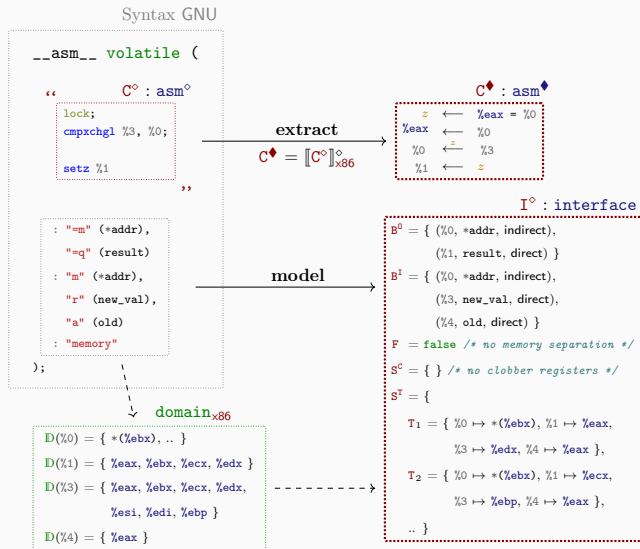
- 6 recurrent patterns yield **90%** of issues
- 5 patterns rely on **fragile** assumptions  
(**80%** of severe issues)

# GNU documentation is informal & incomplete

- no standard, only based on GCC implementation
- non documented behaviors may change at any time
- Clang and icc follow “what they understood”



# Looking for the missing formalism



# Interface compliance properties

## Frame-write :

*“Only **clobber** registers and **output** location are allowed to be **modified** by the assembly template”*

## Frame-read :

*“All **read** values must be **initialized** – only **input** dependent values are allowed in output productions, memory addressing and branching condition”*

## Unicity :

*“The instruction behavior **must not depend** of the **compiler choices**”*

# Interface compliance properties

**Frame-write** :  $\forall \mathbf{l} \notin \mathbf{B}^0 \cup \mathbf{S}^c; \mathbf{S}(\mathbf{l}) = \text{exec}(\mathbf{S}, \mathbf{C}^l \langle \mathbf{T} \rangle)(\mathbf{l})$

*“Only **clobber** registers and **output** location are allowed to be **modified** by the assembly template”*

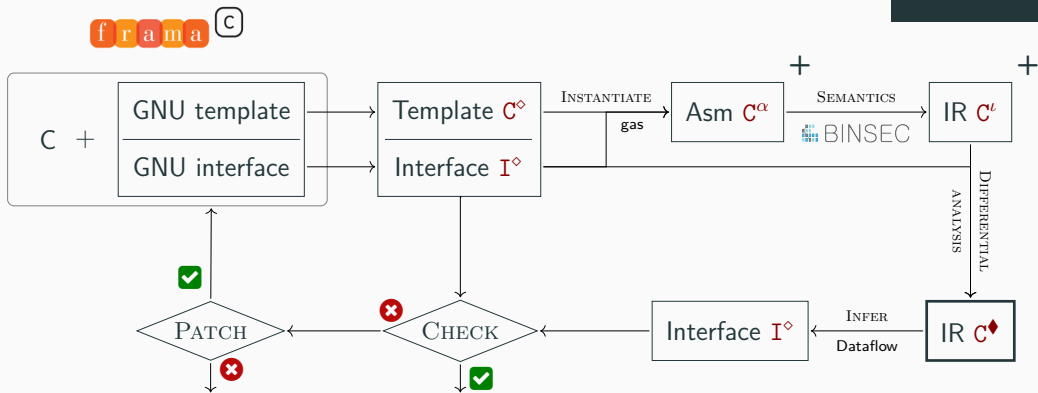
**Frame-read** :  $\text{exec}(\mathbf{S}_1, \mathbf{C}^l \langle \mathbf{T} \rangle) \overset{\diamond}{\cong}_{\mathbf{B}^0, \mathbf{F}}^{\mathbf{T}} \text{exec}(\mathbf{S}_2, \mathbf{C}^l \langle \mathbf{T} \rangle)$

*“All **read** values must be **initialized** – only **input** dependent values are allowed in output productions, memory addressing and branching condition”*

**Unicity** :  $\text{exec}(\mathbf{S}_1, \mathbf{C}^l \langle \mathbf{T}_1 \rangle) \overset{\diamond}{\cong}_{\mathbf{B}^0, \mathbf{F}}^{\mathbf{T}_1, \mathbf{T}_2} \text{exec}(\mathbf{S}_2, \mathbf{C}^l \langle \mathbf{T}_2 \rangle)$

*“The instruction behavior **must not depend** of the **compiler choices**”*  
(Unicity implies Frame-read)

# Our prototype RUSTInA





# Experimental evaluation

- How does perform RUSTINA at checking and patching?
- Why so many issues do not turn more often into bugs?
- What is the real impact of the reported issues?

(more research questions addressed in the paper)

# Checking and patching statistics

	Initial code	Patched code
<b>Found issues</b>	<b>2183</b>	183
significant issues	986	183
<b>frame-write</b>	<b>1718</b>	0
🚩 – flag register clobbered	1197	0
❌ – read-only input clobbered	17	0
❌ – unbound register clobbered	436	0
❌ – unbound memory access	68	0
<b>frame-read</b>	<b>379</b>	183
❌ – non written write-only output	19	0
❌ – unbound register read	183	183
❌ – unbound memory access	177	0
<b>unicity</b>	<b>86</b>	0



Total time: 2min – Average time per chunk: 40ms

**Common** issues (90%)  
do not break very often

Are they somehow under  
“implicit protections”?



What if we **stress out** the compilation  
process? (“copy-paste”, -O3, -lto, etc.)

# Common bad coding practices

6 recurrent patterns yield **90%** of issues

5 of them can lead to **bugs**

Pattern	Omitted clobber	Implicit protection	Robust?	# issues
P1 –	"cc"	compiler choice	✓	1197
P2 –	%ebx register	compiler choice	✗ (GCC ≥ 5) + 🐛	30
P3 –	%esp register	compiler choice	✗ (GCC ≥ 4.6) + 🐛	5
P4 –	"memory"	function embedding	✗ (inlining, cloning) + 🐛	285
P5 –	MMX register	ABI	✗ (inlining, cloning)	363
P6 –	XMM register	compiler option	✗ (cloning)	109
				<b>792</b> 80%

✓ : does not break – ✗ : has been broken – 🐛 : known bug



# Real-life impact of RUSTInA

## Submitted patches (applied or in review)

- 114 faulty chunks in **8 packages**
- **538** severe issues (55%)

libtomcrypt

xfstt

haproxy

UDPCast




**FFMPEG**

x264

ALSA

libatomic\_ops

- Have a look @ the paper
- Have a look @ the artifact
- Have a look @  BINSEC

Interface compliance is **hard**,  
it **matters** but it is **no longer** a problem  
thanks to RUSTINA

If you have any question,  
do not hesitate!