

Automatic verification of low-level code: C, assembly and binary

| | | |
|--------------------------|---------------------|------------|
| Frédéric Recoules | <i>CEA, List</i> | |
| Marie-Laure Potet | <i>Grenoble INP</i> | Director |
| Richard Bonichon | <i>Nomadic Labs</i> | Supervisor |
| Sébastien Bardin | <i>CEA, List</i> | Supervisor |

Today's challenge :
mixed C & **inline assembly** code


Inline assembly example (bits/strings.h@glibc_2.19)

```
1563 # ifdef __PIC__
1564 __STRING_INLINE size_t
1565 __strcspn_g (const char *__s, const char *__reject)
1566 {
1567     register unsigned long int __d0, __d1, __d2;
1568     register const char *__res;
1569     __asm__ __volatile__
1570     ("pushl    %%ebx\n\t"
1571      "movl     %4,%%edi\n\t"
1572      "cld\n\t"
1573      "repne; scasb\n\t"
1574      "notl     %%ecx\n\t"
1575      "leal    -1(%%ecx),%%ebx\n\t"
1576      "i:\n\t"
1577      "lodsrb\n\t"
1578      "testb   %%al,%%al\n\t"
1579      "je      2f\n\t"
1580      "movl    %4,%%edi\n\t"
1581      "movl    %%ebx,%%ecx\n\t"
1582      "repne; scasb\n\t"
1583      "jne     1b\n\t"
1584      "2:\n\t"
1585      "popl    %%ebx"
1586      : "=S" (__res), "=&a" (__d0), "=&c" (__d1), "=&D" (__d2)
1587      : "x" (__reject), "0" (__s), "1" (0), "2" (0xffffffff)
1588      : "memory", "cc");
1589     return (__res - 1) - __s;
1590 }
1618 # endif
```

Inline assembly is well spread



7k packages



Found **3107** x86 chunks
in 202 packages



1264
projets

355
28%¹



ALSA

GMP

libyuv

¹according to Rigger et al.

Adapting formal methods to
common software is **challenging**

Inline assembly makes C analyzers ineffective



```
WARNING: function "main" has inline asm
ERROR: inline assembly is unsupported
NOTE: ignoring this error at this location
```

```
done: total instructions = 161
done: completed paths = 1
done: generated tests = 1
```



```
done for function main
===== VALUES COMPUTED =====
Values at end of function mid_pred:
  i ∈ [--..--]    i ∈ [-5..5] expected
Values at end of function main:
  a ∈ {0; 1; 2; 3; 4; 5}
  b ∈ [-5..10]
  c ∈ [-10..0]
  i ∈ [--..--]    i ∈ [-5..5] expected
```

Incomplete

Imprecise

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

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

A few known inline assembly bugs 🚫

- `strcspn`
`glibc` – Mars 1998 .. January 1999
- `compare_double_and_swap_double`
`libatomic_ops` – February 2008 .. Mars 2012
- `compare_double_and_swap_double`
`libatomic_ops` – Mars 2012 .. September 2012
- `bswap`
`libtomcrypt` – April 2005 .. November 2012

GNU-style interface is **really** error-prone

Goals & challenges

Interface compliance

must ensure that no bug lies in the interface

Enable formal verification

must allow to perform verification of mixed C & inline assembly code

Widely applicable

must be as much architecture, compiler and analysis agnostic



arm



etc.

Prior work on inline assembly

| | <u>Manual</u> | <u>Goanna¹</u> | <u>Vx86²</u> | <u>Inception³</u> | <u>Goal</u> |
|-----------------------------------|---------------|---------------------------|-------------------------|------------------------------|-------------|
| Interface compliance | ✓ | ✓ | N/A | ✗ | ✓ |
| Enable formal verification | ✓ | ✗ | ✓ | ✓ | ✓ |
| Widely applicable | ✗ | ✗ | ✗ | ✓ | ✓ |

¹Fehnker et al. Some Assembly Required - Program Analysis of Embedded System Code

²Schulte et al. Vx86: x86 Assembler Simulated in C Powered by Automated Theorem Proving

³Corteggiani et al. Inception: System-Wide Security Testing of Real-World Embedded Systems Software

Contributions

A **novel** operational semantics for inline assembly

- an operational semantics between C & binary
- a method to automatically extract inline assembly semantics ([TInA-core](#))

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

- comprehensive formalization of **interface compliance** (**Framing** conditions & **Unicity** condition)
- thorough experiments with [RUSTInA](#) over **2.6k⁺** real-world chunks (**986** severe issues found, **803** patches, **7** package patch accepted)
- a study of current bad coding practices (**6** recurrent patterns yield **90%** of issues, including **5** **fragile** patterns)

[ICSE 2021]



A **trustworthy**, **verification-oriented** lifting method

- first **verification friendly** lifting
- tailored post-lifting **validation pass**
- experiments with [TInA](#) over KLEE and Frama-C

[ASE 2019]

The **interface compliance** challenge

Inline assembly example (atomic_ops/sysdeps/gcc/x86.h)

```
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 */
                          : "=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

"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";

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
- code-interface mismatches can lead to bugs

Goals & challenges

Define interface compliance

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 as much compiler agnostic



C compiler

Contributions (1/2)

A **formalization** of interface of compliance

- 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

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** on the **compiler choices***

Interface compliance properties

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

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

Frame-read. $\text{exec}(S_1, C^l \langle T \rangle) \stackrel{\diamond}{\cong}_{B^0, F}^{T} \text{exec}(S_2, C^l \langle T \rangle)$

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

Unicity. $\text{exec}(S_1, C^l \langle T_1 \rangle) \stackrel{\diamond}{\cong}_{B^0, F}^{T_1, T_2} \text{exec}(S_2, C^l \langle T_2 \rangle)$

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

Contributions (2/2)

Thorough experiments of our prototype

- 2.6k⁺ 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)


Checking and patching statistics

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

Over 2656 chunks



Over 202 packages



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

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 |
| | | | | 792 80% |

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

Real-life impact of RUSTiNA

Submitted patches

- 114 faulty chunks in **8 packages** (7 applied)
- **538** severe issues

libtomcrypt

xfstt

haproxy

UDPCast

 **FFMPEG**

x264

ALSA

libatomic_ops

Verification-oriented lifting

Inline assembly makes C analyzers ineffective



```
WARNING: function "main" has inline asm
ERROR: inline assembly is unsupported
NOTE: ignoring this error at this location
```

```
done: total instructions = 161
done: completed paths = 1
done: generated tests = 1
```



```
done for function main
===== VALUES COMPUTED =====
Values at end of function mid_pred:
  i ∈ [--..--]    i ∈ [-5..5] expected
Values at end of function main:
  a ∈ {0; 1; 2; 3; 4; 5}
  b ∈ [-5..10]
  c ∈ [-10..0]
  i ∈ [--..--]    i ∈ [-5..5] expected
```

Incomplete

Imprecise

Common workarounds

```
int mid_pred (int a, int b, int c) {
    int i = b;
    #ifndef DISABLE_ASM
    __asm__
        ("cmp    %2, %1 \n\t"
         "cmovg  %1, %0 \n\t"
         "cmovg  %2, %1 \n\t"
         "cmp    %3, %1 \n\t"
         "cmovl  %3, %1 \n\t"
         "cmp    %1, %0 \n\t"
         "cmovg  %1, %0 \n\t"
         : "+&r" (i), "+&r" (a)
         : "r" (b), "r" (c));
    #else
    i = max(a, b);
    a = min(a, b);
    a = max(a, c);
    i = min(i, a);
    #endif
    return i;
}
```

Manual handling

manpower intensive

error prone

Dedicated analyzer

substantial engineering effort

Our proposition

Automatically **lift** ASM to **equivalent C**

```
int mid_pred (int a, int b, int c)
{
  int i = b;
  __asm__ ("cmp  X2, X1 \n\t"
          "cmovg X1, X0 \n\t"
          "cmovg X2, X1 \n\t"
          "cmp  X3, X1 \n\t"
          "cmovl X3, X1 \n\t"
          "cmp  X1, X0 \n\t"
          "cmovg X1, X0 \n\t"
          : "+&r" (i), "+&r" (a)
          : "r" (b), "r" (c));
  return i;
}
```

C + ASM

Lift

```
int mid_pred (int a, int b, int c)
{
  int i = b;
  {
    int __tina_tmp3, __tina_tmp2;
    int __tina_tmp1, __tina_tmp4;
    __TINA_BEGIN_1__ : ;
    if (a > b) __tina_tmp3 = a;
    else __tina_tmp3 = i;
    if (a > b) __tina_tmp2 = b;
    else __tina_tmp2 = a;
    if (__tina_tmp2 < c) __tina_tmp1 = c;
    else __tina_tmp1 = __tina_tmp2;
    if (__tina_tmp3 > __tina_tmp1)
      __tina_tmp4 = __tina_tmp1;
    else __tina_tmp4 = __tina_tmp3;
    i = __tina_tmp4;
    __TINA_END_1__ : ;
  }
  return i;
}
```

C only

Analyze



Reuse C tools

Goals & challenges

Verification friendly

decent enough analysis outputs for verification process

Trustable

usable in sound formal method context

Widely applicable

must be generic and verification technique agnostic



frama[©]

EVA

frama[©]

WP etc.

Contributions

Dedicated high-level structure recovery mechanism

- identify 3 main threats to verifiability
- dedicated rewriting steps

Tailored validation pass

- preserve control flow graph isomorphism
- SMT based basic block equivalence checking

Thorough experiments of our prototype

- **100%** validation of lifted chunks
- positive impact of **TINA** for 3 standard verification tools (KLEE, Frama-C EVA, Frama-C WP)

Verification-oriented lifting

original

```
__asm__  
(  
    "cmp    %0, %1 \n\t"  
    "cmovg %1, %0 \n\t"  
    /* [ ... ] */  
    : "+&r" (i), "+&r" (a)  
    : /* [ ... ] */  
    : /* no clobbers */  
);
```

- T1. low-level data & computation
- T2. low-level packing & representation
- T3. unusual & unstructured control flow

basic lifting

```
__eax__ = (unsigned int)i;  
__ebx__ = (unsigned int)a;  
__res32__ = __ebx__ - __eax__;  
__zf__ = __res32__ == 0u;  
__sf__ = (int)__res32__ < 0;  
__of__ = ((__ebx__ >> 31)  
         != (__eax__ >> 31)  
         & ((__ebx__ >> 31)  
         != (__res32__ >> 31));  
if (!__zf__ & __sf__ == __of__)  
    goto l1;  
else goto l2;  
l1: __tmp__ = __ebx__; goto l3;  
l2: __tmp__ = __eax__; goto l3;  
l3: __eax__ = __tmp__;  
i = (int)__eax__;
```

TINA lifting

```
int __tmp__;  
if (a > i)  
    __tmp__ = a;  
else  
    __tmp__ = i;  
i = __tmp__;
```

- high-level predicate
- unpacking
- expression propagation
- loop normalization

Verifiability of lifted code

| Analysis | KLEE symbolic execution | Frama-C EVA abstract interpretation | Frama-C WP deductive verification |
|-----------|---|--|---|
| Criterion | Number of explored paths in 10m timeout | Number of functions without alarms | Number of fully discharged proofs |
| Lifting | NONE | 0 / 58 | 0 / 12 |
| | BASIC | 12 / 58 | 1 / 12 |
| | TInA | 19 / 58 | 12 / 12 |

Summary

A **novel** operational semantics for inline assembly

- an operational semantics between C & binary
- a method to automatically extract inline assembly semantics (**TInA-core**)

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

- comprehensive formalization of **interface compliance**
(**Framing** conditions & **Unicity** condition)
- thorough experiments with **RUSTInA** over **2.6k⁺** real-world chunks
(**986** severe issues found, **803** patches, **7** package patch accepted)
- a study of current bad coding practices
(**6** recurrent patterns yield **90%** of issues, including **5 fragile** patterns)

[ICSE 2021]



A **trustworthy**, **verification-oriented** lifting method

- first **verification friendly** lifting
- tailored post-lifting **validation pass**
- experiments with **TInA** over KLEE and Frama-C

[ASE 2019]

**Thank you
for your attention**

Summary

A **novel** operational semantics for inline assembly

- an operational semantics between C & binary
- a method to automatically extract inline assembly semantics ([TInA-core](#))

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

- comprehensive formalization of **interface compliance** (**Framing** conditions & **Unicity** condition)
- thorough experiments with [RUSTInA](#) over **2.6k⁺** real-world chunks (**986** severe issues found, **803** patches, **7** package patch accepted)
- a study of current bad coding practices (**6** recurrent patterns yield **90%** of issues, including **5** **fragile** patterns)

[ICSE 2021]



A **trustworthy**, **verification-oriented** lifting method

- first **verification friendly** lifting
- tailored post-lifting **validation pass**
- experiments with [TInA](#) over KLEE and Frama-C

[ASE 2019]