



The **rev.ng**

binary analysis framework

This presentation is available at:

rev.ng/presentation

All the demos are available at:

github.com/revng/demos

rev.ng Labs



1. 10 people
2. Partly in Milan, Italy, partly in rest of Europe
3. Compiler engineers/security researchers

Outline

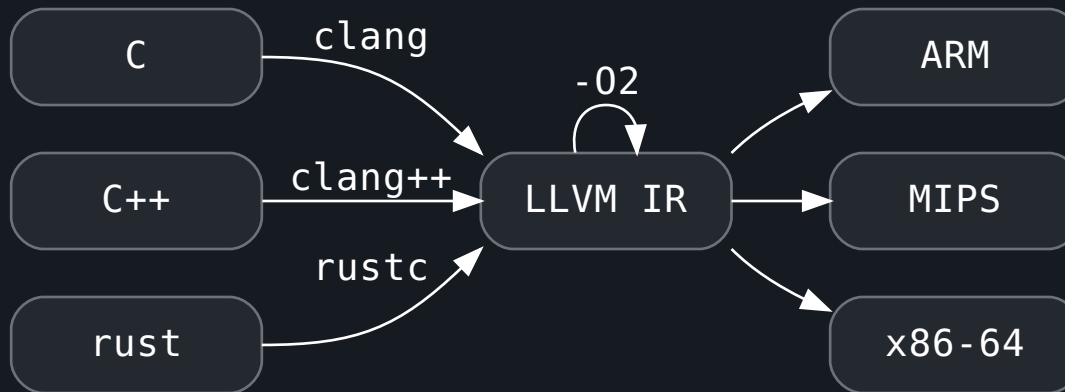
1. rev.ng **design** overview
2. **Demo**: **how to interact** with rev.ng
3. Let's put our hands into **LLVM IR**
4. Automatic **bug finding**
5. The **state of rev.ng**

rev.ng design overview

rev.ng is an **open source**
binary analysis framework
and **interactive** decompiler
for native code

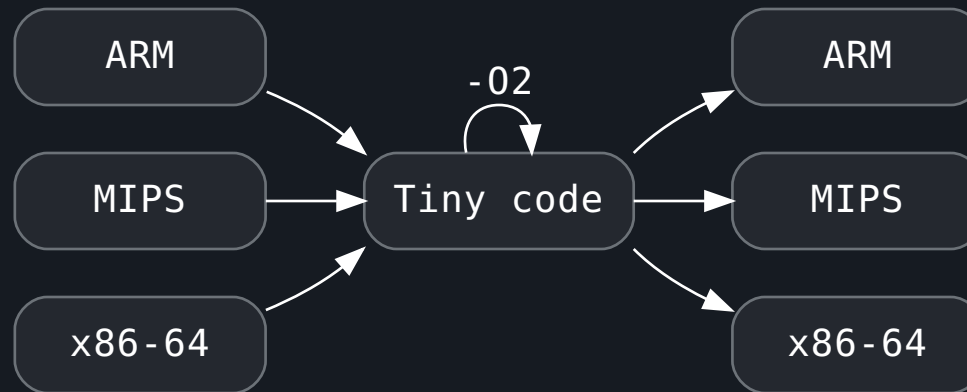
This is a compiler

LLVM



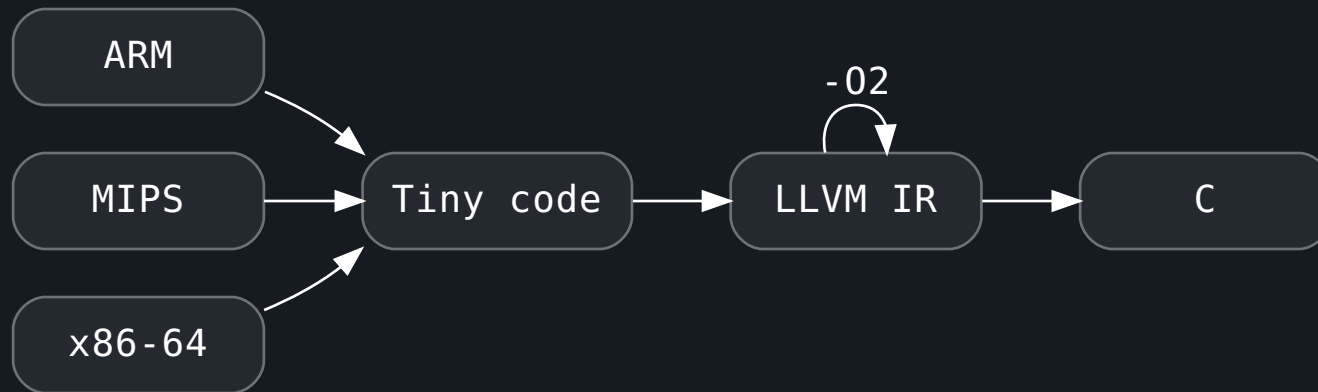
This is a **dynamic binary translator**

QEMU



This is a decompiler

rev.ng





Using QEMU as a lifter

We use QEMU but
we **do not** run any code

We just use it as a lifter

Writing an **accurate lifter** is hard



An amazing QEMU disk image every day!
Brightening your days in the winter holiday season.
This advent calendar is brought to you by the QEMU community.



Day 1 - TinyCore Linux

TinyCore linux

Size of download is 22M bytes.

[Download](#)

Day 2 - Bootable PDF holiday card

Bootable PDF holiday card

Size of download is 8M bytes.

[Download](#)

Day 3 - Bootable Assembly word game: FLORDLE

Bootable Assembly word game: FLORDLE

Size of download is 2.5k bytes.

[Download](#)

QEMU Advent Calendar

Supporting **many architecture** is hard

QEMU supports:

Alpha, **ARM/AArch64**, CRIS, HPPA, **i386/x86-64**,
Hexagon, LatticeMico32, 68K, MicroBlaze, **MIPS**,
Moxie, Nios2, OpenRISC, PowerPC, RISC-V, SH4,
Sparc, **s390x**, TileGX, TriCore, Unicore32, Xtensa

QEMU is **good** for:

- accuracy
- supporting many architectures
- lifting
- dynamic binary translation

QEMU is **bad** for:

- anything else

Tiny code



LLVM IR

Here be dragons



LLVM

- Enable us to focus on building a **decompiler**, not a compiler framework
- Well known and big community
- Well defined semantics
- Many tools build on top of it (e.g., **KLEE**)
- High performance (C++)

A note on **symbolic execution**

We do not use symbolic execution in the pipeline.
However, we deem it appropriate for bug hunting!

A note on fuzzing

We've done it.

We're no longer focused on it.

There are a lot of effective alternative approaches.

How do I **interact**
with rev.ng?

The model

- basically rev.ng's project file
- a YAML document
- contains everything the user can customize


Example

```
Architecture: x86_64
DefaultABI: SystemV_x86_64
Segments:
  - StartOffset: 0
    FileSize: 7
    StartAddress: "0x400000:Generic64"
Functions:
  - Entry: "0x400000:Code_x86_64"
TypeDefinitions:
  - Kind: StructType
    ID: 1
    Size: 8
    Fields:
      - Offset: 4
        Type: ...
```


Interaction option #1: the CLI

```
revng \  
  artifact \  
  --analyze \  
  decompile-to-single-file \  
  /bin/df
```

Interaction option #2: the daemon



tl;dr: **users stay out-of-process**

They can:

- use any language having YAML + HTTP/system
- use any version of the language they want
- use a different version of LLVM
- be on a different machine
- crash independently from rev.ng

No more



Products >

Solutions

Partners

Shop

Support >

Company >



IDAPython and Python 3

As of now (IDA version 7.3), IDA ships with an IDAPython plugin, that is compiled against, and compatible with Python 2.7.

The problem: Python 2.x end-of-life

The Python authors [have decided](#) that Python 3 has been available for long enough, to drop support for Python 2.x.

That effectively means that since Python 2.x will be unmaintained, it will gradually disappear from the landscape.

Moving IDAPython to Python 3

Work has begun (in fact, work is even finished) here at Hex-Rays to make IDAPython compilable, and compatible with Python 3.

users stay out-of-process

developers stay in-process

They have to buy into our dev stack (docs)

Demo time!

Try it!

tl;dr there are only two types of actions:

1. Request an **artifact** (LLVM IR, valid C, ...)
2. Run an **analysis**/make changes to the model

It's **IR** time!

Example program

```
long myfunction(long value) {  
    long result = value;  
    result = result * 2;  
    return result;  
}
```

Disassembly

```
1  myfunction:  
2  push    rbp  
3  mov     rbp, rsp  
4  mov     QWORD PTR [rbp-0x8], rdi  
5  mov     rax, QWORD PTR [rbp-0x8]  
6  mov     QWORD PTR [rbp-0x10], rax  
7  mov     rax, QWORD PTR [rbp-0x10]  
8  shl     rax, 0x1  
9  mov     QWORD PTR [rbp-0x10], rax  
10 mov     rax, QWORD PTR [rbp-0x10]  
11 pop     rbp  
12 ret
```



```

***
define void @local_uyfunction() {
call @llvm.lifetime.start@plt(%@"Code_x86_64", @.lba.1, @.lba.1, @.lba.0,
ptr null, ptr null)
%0 = load i64, ptr @rbp
%1 = load i64, ptr @rbp
%2 = add i64 %1, -8
%3 = inttoptr @.lba.2 to ptr
store i64 %0, ptr %3
store i64 %2, ptr @rbp
call @llvm.lifetime.start@plt(%@"Code_x86_64", @.lba.1, @.lba.0, @.lba.0,
ptr null, ptr null)
%4 = load i64, ptr @rbp
store i64 %4, ptr @rbp
call @llvm.lifetime.start@plt(%@"Code_x86_64", @.lba.4, @.lba.0, @.lba.0,
ptr null, ptr null)
%5 = load i64, ptr @rbp
%6 = add i64 %5, -8
%7 = load i64, ptr @rdi
%8 = inttoptr @.lba.5 to ptr
store i64 %7, ptr %8
call @llvm.lifetime.start@plt(%@"Code_x86_64", @.lba.4, @.lba.0, @.lba.0,
ptr null, ptr null)
%9 = load i64, ptr @rbp
%10 = add i64 %9, -8
%11 = inttoptr @.lba.6 to ptr
%12 = load i64, ptr %11
store i64 %12, ptr @rax
call @llvm.lifetime.start@plt(%@"Code_x86_64", @.lba.4, @.lba.0, @.lba.0,
ptr null, ptr null)
%13 = load i64, ptr @rbp
%14 = add i64 %13, -16
%15 = load i64, ptr @rax
%16 = inttoptr @.lba.7 to ptr
store i64 %15, ptr %16
call @llvm.lifetime.start@plt(%@"Code_x86_64", @.lba.4, @.lba.0, @.lba.0,
ptr null, ptr null)
%17 = load i64, ptr @rbp
%18 = add i64 %17, -16
%19 = inttoptr @.lba.8 to ptr
%20 = load i64, ptr %19
store i64 %20, ptr @rax
call @llvm.lifetime.start@plt(%@"Code_x86_64", @.lba.4, @.lba.0, @.lba.0,
ptr null, ptr null)
%21 = load i64, ptr @rax
%22 = shl i64 %21, 1
store i64 %22, ptr @rax
store i64 %22, ptr @cc_dst
call @llvm.lifetime.start@plt(%@"Code_x86_64", @.lba.4, @.lba.0, @.lba.0,
ptr null, ptr null)
%23 = load i64, ptr @rbp
%24 = add i64 %23, -16
%25 = load i64, ptr @rax
%26 = inttoptr @.lba.9 to ptr
store i64 %25, ptr %26
call @llvm.lifetime.start@plt(%@"Code_x86_64", @.lba.4, @.lba.0, @.lba.0,
ptr null, ptr null)
%27 = load i64, ptr @rbp
%28 = add i64 %27, -16
%29 = inttoptr @.lba.10 to ptr
%30 = load i64, ptr %29
store i64 %30, ptr @rax
call @llvm.lifetime.start@plt(%@"Code_x86_64", @.lba.1, @.lba.0, @.lba.0,
ptr null, ptr null)
%31 = load i64, ptr @rbp
%32 = inttoptr @.lba.11 to ptr
%33 = load i64, ptr %32
%34 = add i64 %33, 8
store i64 %34, ptr @rbp
store i64 %33, ptr @rbp
call @llvm.lifetime.start@plt(%@"Code_x86_64", @.lba.1, @.lba.0, @.lba.0,
ptr null, ptr null)
%35 = load i64, ptr @rbp
%36 = inttoptr @.lba.12 to ptr
%37 = load i64, ptr %36
%38 = add i64 %37, 8
store i64 %38, ptr @rbp
store i64 %37, ptr @cc
ret void
}

```

```

define i64 @local_myfunction(i64 %rdi_x86_64) {
  %0 = call i64 @__init_rbp()
  call @newpc(ptr nonnull @"revng.const.0x401130:Code_x86_64", i64 1, i32 1, i32 0,
ptr null, ptr null)
  %1 = load i64, ptr @rsp
  %2 = add i64 %1, -8
  %3 = inttoptr i64 %2 to ptr
  store i64 %0, ptr %3
  store i64 %2, ptr @rsp
  call @newpc(ptr nonnull @"revng.const.0x401131:Code_x86_64", i64 3, i32 0, i32 0,
ptr null, ptr null)
  %4 = load i64, ptr @rsp
  call @newpc(ptr nonnull @"revng.const.0x401134:Code_x86_64", i64 4, i32 0, i32 0,
ptr null, ptr null)
  %5 = add i64 %4, -8
  %6 = inttoptr i64 %5 to ptr
  store i64 %rdi_x86_64, ptr %6
  call @newpc(ptr nonnull @"revng.const.0x401138:Code_x86_64", i64 4, i32 0, i32 0,
ptr null, ptr null)
  %7 = load i64, ptr %6
  call @newpc(ptr nonnull @"revng.const.0x40113c:Code_x86_64", i64 4, i32 0, i32 0,
ptr null, ptr null)
  %8 = add i64 %4, -16
  %9 = inttoptr i64 %8 to ptr
  store i64 %7, ptr %9
  call @newpc(ptr nonnull @"revng.const.0x401140:Code_x86_64", i64 4, i32 0, i32 0,
ptr null, ptr null)
  %10 = load i64, ptr %9
  call @newpc(ptr nonnull @"revng.const.0x401144:Code_x86_64", i64 4, i32 0, i32 0,
ptr null, ptr null)
  %11 = shl i64 %10, 1
  call @newpc(ptr nonnull @"revng.const.0x401148:Code_x86_64", i64 4, i32 0, i32 0,
ptr null, ptr null)
  store i64 %11, ptr %9
  call @newpc(ptr nonnull @"revng.const.0x40114c:Code_x86_64", i64 4, i32 0, i32 0,
ptr null, ptr null)
  %12 = load i64, ptr %9
  call @newpc(ptr nonnull @"revng.const.0x401150:Code_x86_64", i64 1, i32 0, i32 0,
ptr null, ptr null)
  %13 = load i64, ptr @rsp
  %14 = add i64 %13, 8
  store i64 %14, ptr @rsp
  call @newpc(ptr nonnull @"revng.const.0x401151:Code_x86_64", i64 1, i32 0, i32 0,
ptr null, ptr null)
  %15 = load i64, ptr @rsp
  %16 = add i64 %15, 8
  store i64 %16, ptr @rsp
  ret i64 %12
}

```

```
define i64 @local_myfunction(i64 %0) {
    %1 = call i64 @revng_stack_frame(i64 24)
    %2 = call i64 @AddressOf(ptr nonnull
@revng.const.c10d6afb753dc601da714646784a7e4040e86f7b, i64 %1)
    %3 = add i64 %2, 8
    %4 = inttoptr i64 %3 to ptr
    %5 = call ptr @stack_offset(ptr %4, i64 -16, i64 -7)
    store i64 %0, ptr %5
    %6 = inttoptr i64 %2 to ptr
    %7 = shl i64 %0, 1
    %8 = call ptr @stack_offset(ptr %6, i64 -24, i64 -15)
    store i64 %7, ptr %8
    ret i64 %7
}
```

```
define i64 @local_myfunction(i64 %0) {  
    %1 = shl i64 %0, 1  
    ret i64 %1  
}
```

A couple of demos

1. Write a small **taint analysis** (**try it!**)
2. Collect some information about **loops** (**try it!**)

- ENOTIME

You can check them out at

github.com/revng/demos

Let's find some bugs!

Both on **LLVM IR** and **decompiled C**

```
int main(int argc, char **argv) {
    return do_stuff(strtol(argv[1], NULL, 10));
}

void my_free(void *p) {
    free(p);
}

int do_stuff(int condition) {
    int *p = malloc(sizeof(int));

    if (condition > 4) {
        my_free(p);
        // programmer forgot to return 0;
    }

    *p = 3; // use-after-free
    int result = *p;

    my_free(p); // double free
    return result;
}
```

1. KLEE

A symbolic execution engine for LLVM IR

Try it!

newFuncRoot:

```
%1 = call i64 @local_malloc_(i64 4) #8  
%2 = trunc i64 %0 to i32  
%3 = icmp sgt i32 %2, 4  
br i1 %3, label %then, %else
```

T

F

then:

```
call void @local_my_free(i64 %1) #8  
br label %else
```

else:

```
%4 = inttoptr i64 %1 to ptr  
store i32 3, ptr %4, align 1  
call void @local_my_free(i64 %1) #8  
ret i64 3
```

CFG for 'local_do_stuff' function

Profit!

```
$ klee --posix-runtime --libc=klee lifted.bc --sym-args 1 1 2
[...]
KLEE: ERROR: revng.module:0: memory error: use after free
[...]
$ ktest-tool klee-last/test000016.ktest
ktest file : 'klee-last/test000016.ktest'
args       : ['use-after-free.bc', '--sym-args', '1', '1', '2']
num objects: 1
object 0: size: 3
object 0: data: b'\n9\n'
object 0: text: .9.

$ cat klee-last/test000016.ptr.err
Error: memory error: use after free
File: revng.module
assembly.ll line: 1233
State: 7
Stack:
    #000001233 in local_do_stuff(symbolic) at [...]
    #100001215 in local_main(2, 138819871686656) at [...]
```

2. clang static analyzer (try it!)

Original C

```
void my_free(void *p) {
    free(p);
}

int do_stuff(int condition) {
    int *p = malloc(sizeof(int));
    if (condition > 4)
        my_free(p);
    *p = 3;
    int result = *p;
    my_free(p);
    return result;
}
```

Clang Static Analyzer Report

```
79 ABI(SystemV_x86_64)
80 generic64_t do_stuff(generic64_t _argument0) {
81     void *_var_0;
82     _var_0 = malloc((size_t) 4);
83     if ((int32_t) (generic32_t) _argument0 > (int32_t) 2) {
84         my_free((generic32_t *) _var_0);
85     }
86     *(generic32_t *) _var_0 = 3;
87     my_free((generic32_t *) _var_0);
88     return 3;
89 }
```

2 ← Calling 'malloc.' →

4 ← Returned allocated memory →

5 ← Assuming '_argument0' is > 2 →

6 ← Taking true branch →

7 ← Calling 'my_free' →

11 ← Returning; memory was released via 1st parameter →

12 ← Use of memory after it is freed

3. CodeQL (try it!)

Original C

CodeQL Report

```
int do_stuff(int condition) {
    int *p = malloc(sizeof(int));
    if (condition > 4)
        free(p);
    *p = 3;
    int result = *p;
    free(p);
    return result;
}
```

```
_ABI(SystemV_x86_64)
generic64_t do_stuff(generic64_t _argument0) {
    void *_var_0;
    _var_0 = malloc((size_t) 4);
    if ((int32_t) (generic32_t) _argument0 > (int32_t) 4) {
        // BUG 1: call to free
        // BUG 2: call to free
        free((generic32_t *) _var_0); // line 69
    }
    // BUG 1: Potential use after free
    // An allocated memory block is used after it has been freed.
    // Behavior in such cases is undefined and can cause memory corruption.
    // Memory may have been previously freed by call to free at line 69.
    *(generic32_t *) _var_0 = 3;
    // BUG 2: Potential double free
    // Behavior in such cases is undefined and can cause memory corruption.
    // Memory may already have been freed by call to free at line 69.
    free((generic32_t *) _var_0);
    return 3;
}
```


tl;dr You can use tools
designed for **source code**
on binaries

Spoiler: their issues do not magically disappear.

The **status** of rev.ng

Supported CPUs

- x86
- x86-64
- ARM
- AArch64
- MIPS
- SystemZ

Importers

- ELF
- DWARF
- PE/COFF
- CodeView (.pdb)
- Mach-O
- IDA Pro

Where does rev.ng run?

- Daemon
 - Linux x86-64 natively
 - macOS via Docker
 - Windows via WSL
- Clients can run anywhere

Recently, we released the pipeline as open source.

We're now focusing

on **robustness** and **performance**.

Performance

We currently produce IR for GCC in 18 minutes.

- New argument detection analysis: 2.1x
- Reduce invalidation: ~1.7x (10min expected)
- Other low effort fixes: ~2x (5min expected)

Goal: **be on par** with
Ghidra and IDA

1min 30sec and 40sec

Next up: mass testing

- Decompile all binaries on:
 - Ubuntu x86-64
 - Windows x86-64
 - Android AArch64
 - Various malwares
- Focus on:
 - squashing bugs
 - optimizing performance

In short, rev.ng

- is FLOSS
- is declarative in user interactions
- has a modern design
- uses a “standard” IR and emits valid C
- interacts with existing tools
- has a nice (commercial) UI

rev.ng UI

- Based on **VSCode**
- Connects to **revng-daemon**
- **Collaboration** just works out of the box
- Runs as a **standalone** app or in the **cloud**
- Also, **hub.rev.ng** (think GitHub for reversers)
- Cloud version will be **free for public projects**

Final note

We're **hiring** and we do **consulting**

Questions?

EXPLORER

UNTITLED (WORKSPACE)

- linked_list (on rev.ng cloud)
 - binary
 - function
 - type
 - ! model.yml

Overview: linked_list (on rev.ng cloud) X

Upload Binary

Choose File No file chosen

Upload & Analyze Binary

Information

Architecture: x86_64
DefaultABI: SystemV_x86_64
Entrypoint: `unreserved__start` (0x4007c0:Code_x86_64)

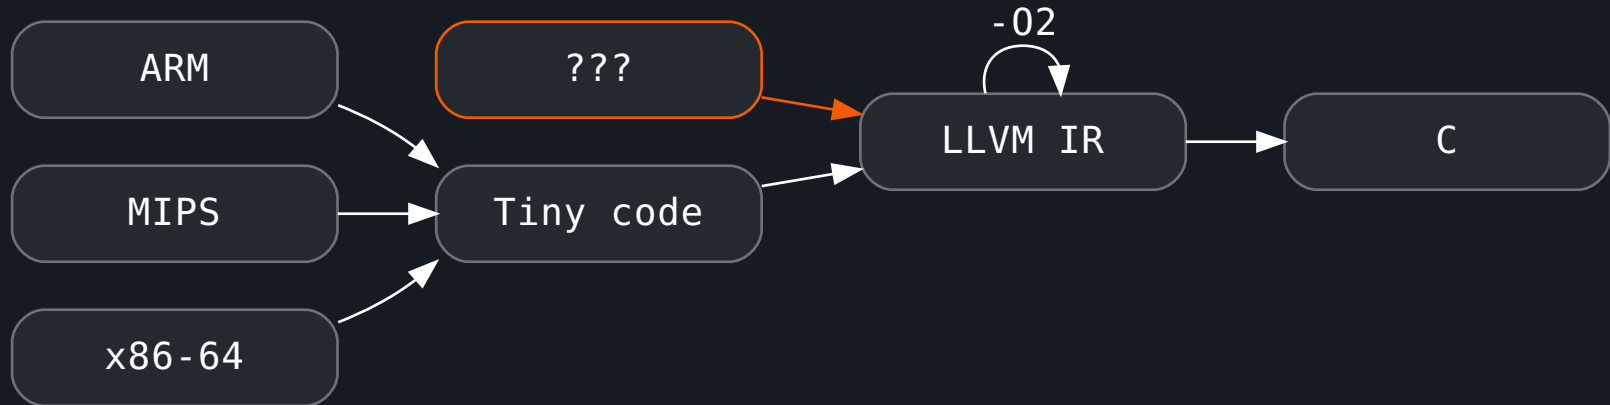
Segments:

	Start address	End address	Size	File Offset	Trailing Zeros Size	Permissions
>	0x400000	0x400b98	2968	0	0	r-x
>	0x401d98	0x402029	657	3480	1	rw-

0 0 rev.ng: no pending changes rev.ng projects status: Layout: us

Backup slides

non-QEMU frontends?



Idea

```
typedef struct {
    uint32_t rax;
    uint32_t rdi;
    // ...
} CPUState;

void add(CPUState *state) {
    state->rax = state->rax + state->rdi;
}
```

WIP, particularly interesting for WebAssembly

Automated type recovery

No type recovery

```
generic64_t sum(generic64_t _argument0) {
    generic64_t _var_0 = 0, _var_1 = 0;
    do {
        _var_1 = _var_1 + *(generic64_t *) ((_var_0 << 3) + _argument0);
        _var_0 = _var_0 + 1;
    } while (_var_0 != 5);
    return _var_1;
}

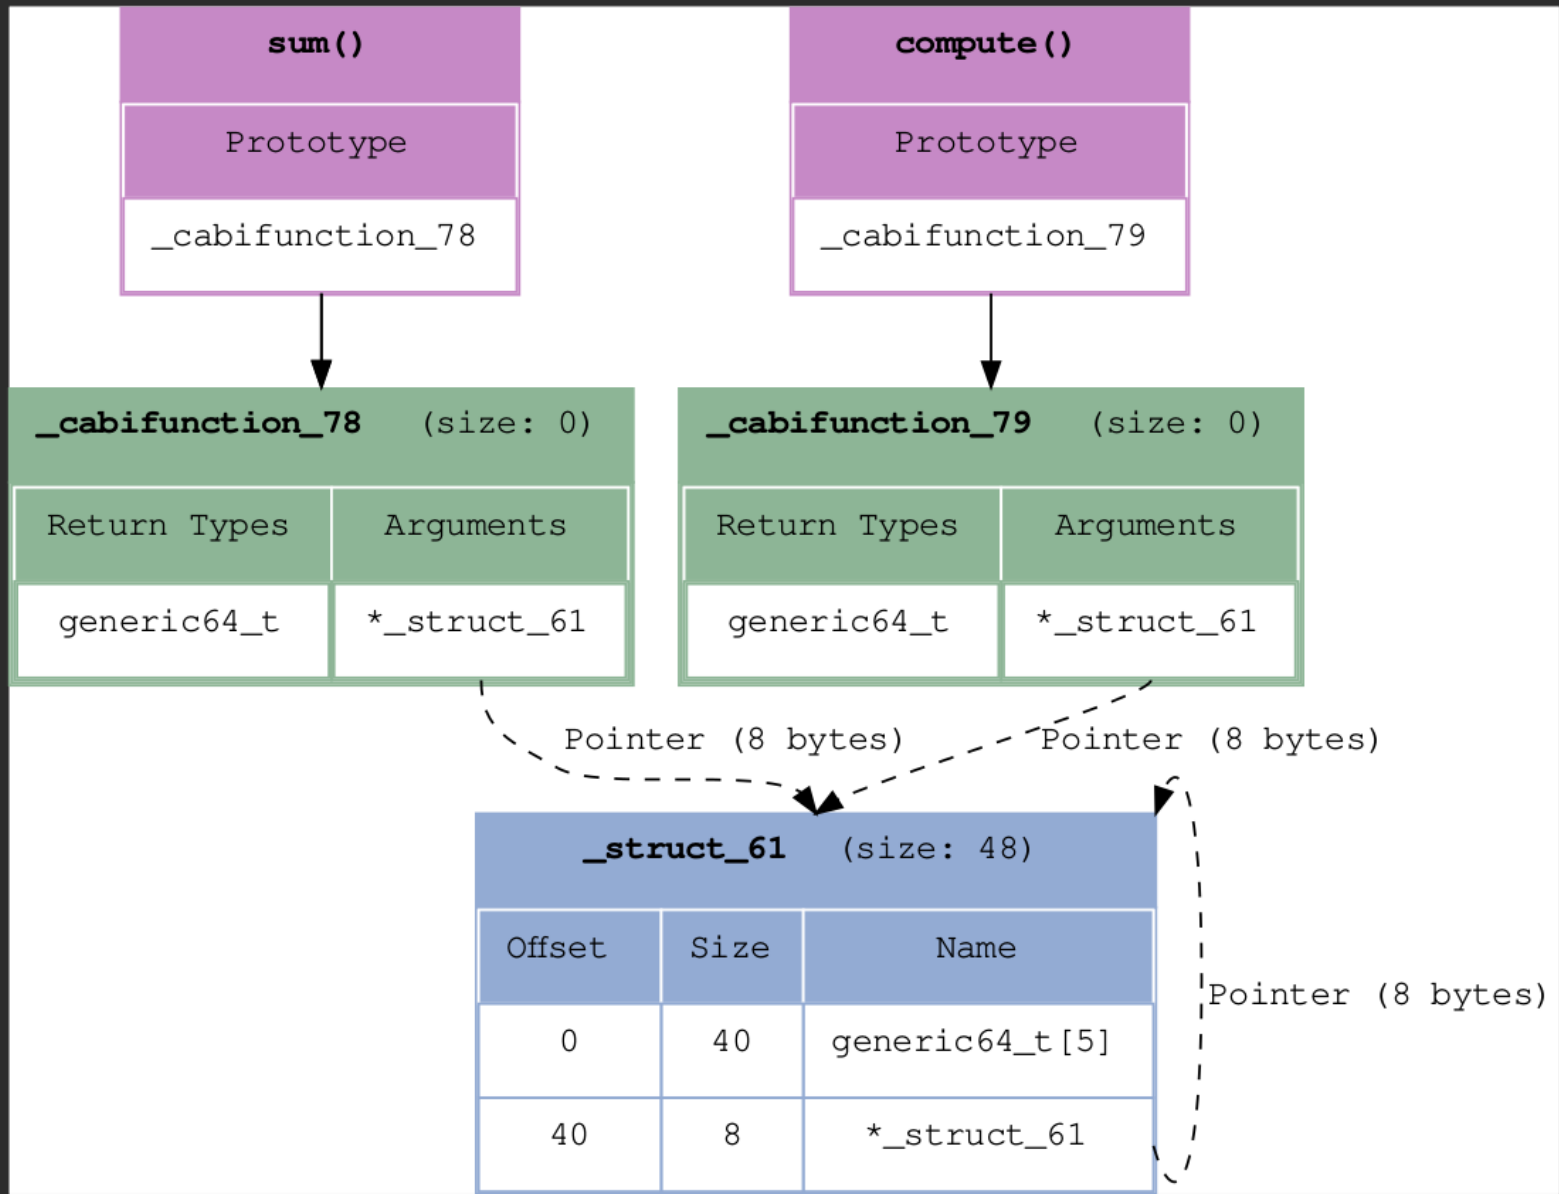
generic64_t compute(generic64_t _argument0) {
    generic64_t _var_0 = _argument0, _var_1 = 0;
    generic64_t _var_2;
    do {
        gen_var_2 = sum(_var_0);
        _var_1 = _var_1 + _var_2;
        _var_0 = *(generic64_t *) (_var_0 + 40);
    } while (_var_0);
    return _var_1;
}
```

With automated type recovery

```
typedef struct _PACKED _struct_61 {
    generic64_t _offset_0[5];
    _struct_61 *_offset_40;
} _struct_61;

generic64_t sum(_struct_61 *_argument0) {
    generic64_t _var_0 = 0, _var_1 = 0
    do {
        _var_0 = _var_0 + _argument0->_offset_0[_var_1];
        _var_1 = _var_1 + 1;
    } while (_var_1 != 5);
    return _var_0;
}

generic64_t compute(_struct_61 *_argument0) {
    _struct_61 *_var_0 = _argument0
    generic64_t _var_1 = 0, _var_2;
    do {
        _var_2 = sum(_var_0);
        _var_1 = _var_1 + _var_2;
        _var_0 = _var_0->_offset_40;
    } while (_var_0);
    return _var_1;
}
```



Things you can do in LLVM

- Graph theory
 - Build the **dominator tree**
 - Identify **strongly connected components**
 - Perform **visits** (depth first, topological...)
- Manipulate functions
 - Inline
 - Outline
 - Specialize