

Just-in-Time Recompilation and Optimization of Compiled Binaries

Student: Lucas Elisei

Supervisor: Alberto Dassatti

July 2018

Abstract

Over the past decades, computer science fields grew up exponentially. More and more computational power is needed to solve modern problems and this implies a lot of energy consumption.

A way to reduce the energy consumption of data centers and their applications could be to optimize the applications so that they require less computational power.

This document explores an approach to optimize applications by recompiling them without sources. Recompiling entire binaries would take a lot of time so we focus only on parts of the applications which require a lot of computational power.

Our method is based on finding a function that takes a lot of time to execute, recompiling it into more efficient native code to finally patch the binary so it executes the newly optimized function.

Early testing showed that our solution can enhance a binary performance up to 25%. Unfortunately, the decompilers that exist nowadays are still in an early stage of development, thus limiting the capacities of our method. In a near future, the decompilers technology will evolve to allow even more performance enhancement.

Abstract

Ces dernières années, le domaine de l'informatique a connu un grand essor. De plus en plus de puissance de calcul est requise pour résoudre des problèmes modernes et cela implique une grande consommation d'énergie.

Une façon de palier notamment à la consommation d'énergie des data centers et de leurs applications serait d'optimiser ces dernières afin qu'elles demandent moins de puissance de calcul.

Ce document explore une approche pour optimiser n'importe quelle application en la recompilant sans avoir accès au code source. Recompiler un programme dans son intégralité demanderait beaucoup de temps de ce fait nous nous concentrons uniquement sur les parties critiques de l'application.

Notre méthode se base sur trouver une fonction critique, la recompiler en du code natif plus efficace pour enfin modifier le programme afin qu'il exécute la fonction optimisée.

Des tests ont permis de montrer que notre solution améliore de la performance d'une application jusqu'à 25%. Malheureusement, les décompilateurs existant de nos jours sont encore en phase de développement prématuré, limitant les capacités de notre méthode. Dans un futur proche, la technologie de décompilation va évoluer pour proposer des améliorations encore plus marquées.

Table of contents

1	Introduction	1
1.1	Project aim and objectives	1
1.2	Disposition	1
1.3	Requirements	2
1.4	Theoretical overview	2
1.4.1	Executable and Linkable Format	2
1.4.2	Decompilation	3
2	Literature review	5
3	State of the art tools	6
3.1	Intermediate representations	6
3.2	Existing tools	6
4	Successful recompilation	9
5	Automatization tool	11
5.1	Design	11
5.2	elfparser	12
5.3	retdec	13
5.4	jit	13
5.5	live-patcher	14
5.6	liboptimizer	15
5.7	optimizer	16
6	Examples	18
6.1	Simple addition	18
6.1.1	32-bit	18
6.1.2	64-bit	19
6.2	Simple multiplication	21
6.2.1	32-bit	21
6.2.2	64-bit	22
6.3	Matrices multiplication	23
6.3.1	32-bit	23
6.3.2	64-bit	24
7	Conclusion	25
8	Acknowledgements	26
	References	28
	Appendices	29
A	Authentication	29
B	Requirements	30
C	Simple decompilation example	31
D	liboptimizer	33
D.1	elfparser	33
D.2	retdec	40
D.3	jit	42
D.4	live-patcher	45
D.5	liboptimizer	52
D.6	mmult	55

1 Introduction

Over the past decades, computer science fields grew up exponentially. More and more computational power is needed to solve actual problems (e.g. machine learning or artificial intelligence).

In 1965, Gordon Moore, Intel's co-founder, stated that the number of transistors¹ that can be placed on a integrated circuit doubles roughly every two years [7]. It is known as Moore's Law. Due to physics and economic limitations, this law is unfortunately coming to an end and that might be critical for technological progress.

A way to negate the end of Moore's Law could be to optimize existing applications so that they require less computational power. Most programs are proprietary software, which means we don't have access to the source code, leaving us this compiled binaries².

This document explores an approach to optimize applications by recompiling them without sources. Just-in-Time recompilation can make use of runtime to dynamically recompile parts of the executed applications to generate a more efficient native code.

Hence, as a consequence of this, executable programs need less computational power, negating the end of Moore's Law. Furthermore, it could allow any applications to run on any architecture, thus reducing development time. Indeed, nowadays, developers are forced to develop specific code for each architecture to produce high-performance applications.

1.1 Project aim and objectives

The aim of this project is to optimize any application binary by finding the potential bottlenecks at runtime. Those bottlenecks are then decompiled to an *Intermediate Representation* (IR), optimized and finally recompiled so that the running application uses the new version.

To achieve this aim, the following objectives have been identified:

1. Review several decompilation tools and find one that suits our needs.
2. Find a way to successfully recompile a given function from a simple program.
3. Develop a program to automatize the process.

A long-term vision could be to extend the Just-in-Time recompilation and optimization to heterogeneous systems³ by off-loading code to an FPGA or a GPU [39]. Doing so would significantly lower the CPU load and reduce application execution time.

1.2 Disposition

This document details every stage of the project. It follows a logical structure and outlines the major stages in chronological order. A brief summary of each section is presented in the list below.

1. Introduction

Presents the project, its aim and objectives and introduces some technical terms.

2. Literature review

Discusses papers and work that other teams have done and/or are doing with decompilers.

3. State of the art tools

Reviews several decompilation tools and evaluates the best one to use for this project.

4. Successful recompilation

Details how to achieve the second objective, the recompilation step.

5. Automatization tool

Details the implementation of the third objective, the automatization tool.

¹In a computer, a transistor is a component that represents the binary 0's and 1's (bits).

²Executable program.

³An heterogeneous system uses more than one kind of processor (e.g., CPU, GPU or FPGA) — Wikipedia

1.3 Requirements

The tools required to run and use the project are detailed in the appendix B.

1.4 Theoretical overview

Before diving into the core of the subject, it could be useful to review some technical terms that are necessary for the good understanding of this document.

1.4.1 Executable and Linkable Format

The Executable and Linkable Format (ELF) is common standard file format for executable files, object code and shared libraries. Since ELF is by design flexible, extensible and cross-platform, it has been adopted by a plethora of operating systems on many different platforms [8].

Each ELF file is made up of one ELF header. It starts with the four bytes magic number 0x7F, 'E', 'L', 'F'. The ELF file header contains general information about the executable, such as the addresses length (32- or 64-bit), the endianness (*big-endian* or *little-endian*), the object file type (executable, relocatable or shared object), the assembly architecture (e.g., x86 or ARM), the virtual address of its entry point (which indicates the starting point of the program execution) and the offsets to the program and sections headers.

The program header is meaningful to executables and shared objects only. It contains a description for each segment (which contains one or more sections) and other information the system needs to prepare the program for execution.

The section headers contain a description for each section like address location and access rights (i.e. read, write and execution).

The figure 1 shows an overview of how an ELF executable is structured.




Figure 1: Structure of an ELF executable [8]

1.4.2 Decompileation

A compiler is a software that translates human-readable programming languages (e.g. C or Java) into a machine-code language (e.g. assembly) to produce an executable program. A decompiler does the exact opposite: it lifts code from a low-level into a higher-level representation.

The figure 2 shows the difference between a simple C code and its translation into machine code.

```
int a = 5;  
int b = 4;  
int c = a + b;
```

```
movl    $5, %eax  
movl    $4, %ebx  
addl    %ebx, %eax  
ret
```

Figure 2: On the left, a simple C program that adds two integers a and b into an integer c. On the right, the resulting assembly code after compilation.

Compared to decompilers, compilers have a more privileged position: the input language is strictly defined and plenty of information is available for functions, variables, types, etc. If the compiler cannot generate a valid code because the input does not conform to the language standards, it is allowed to simply print an error message and stop. Decompilers do not benefit from anything similar. Just the contrary [6]:

- Compiled architecture instructions generally use variable length encoding ;
- The input binary is often obfuscated ;
- Many decompilation problems are unsolved or proven to be unsolvable in generic cases ;
- The output is examined in details by a human being and any sub-optimality is noticed.

In conclusion, robust machine code decompilation is impossible. A decompiler will always have some imperfections and eventually generate wrong output. Our best hope is to diminish the undesired effects as much as possible. To achieve this, here are some basic ideas [16]:

- Make some configurable assumptions about the input (e.g. calling conventions). The user will be able to control the decompiler by specifying the missing information. In simple cases, the decompiler will deduce or guess it.
- Use solid theoretical approach to solve problems (e.g. instruction simplification).
- Use heuristics for unsolvable problems (indirect jumps, call arguments).
- Prefer to generate ugly but correct output rather than nice but incorrect. Let the user embellish if he wants to do so.
- Let the user guide the decompilation in difficult cases (e.g. function prototypes).

Decompilation can be broken down into several phases:

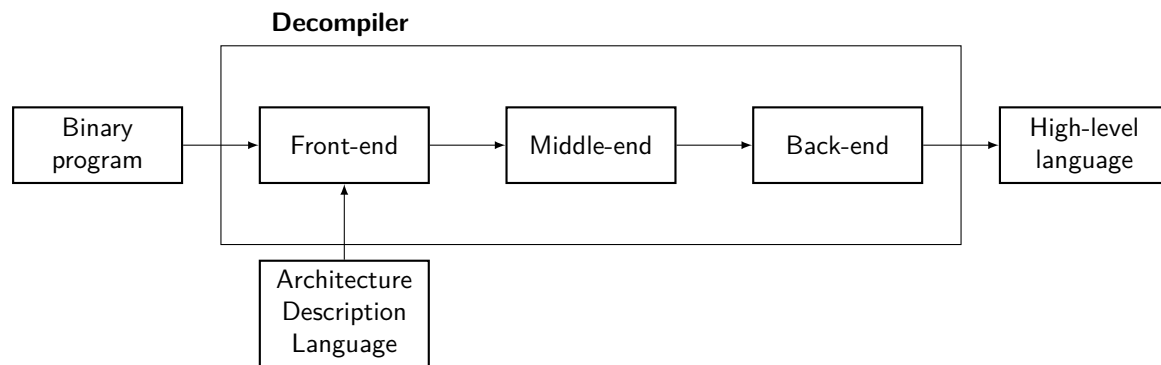


Figure 3: Decompiler structure.

1. Front-end

Parses a binary program (ELF) and translates an architecture-specific machine code into a sequence of low-level IR. To do so, the front-end uses an architecture description language (ADL). The ADL contains information such as the processor resources (i.e. available registers and memory) and instruction set (i.e. assembler language syntax, binary encoding and behaviour of each instruction). Each instruction of the ELF binary is then decoded into an intermediate representation, which describes the program behaviour in a platform-independent way.

2. Middle-end

Improves the properties of the previously generated low-level IR code and prepares it for the back-end phase. Improvements includes [5]:

- Search for idioms and other types of program analysis such as constant or expression propagation.
- Retrieval of high-level constructs, such as `if..else` statements or loops.
- Code optimization.

To do so, the decompiler generates a control-flow graph (CFG) and simplifies it to exclude useless instructions.

3. Back-end

Finally, the back-end converts the optimized IR into the target high-level language (e.g. C or Python). During this conversion, loops and conditional statements are identified and reconstructed into a human-readable way. An other optimization is done and the binary is emitted in the form of the target high-level language.

2 Literature review

As stated previously, a robust decompilation is impossible. But in most cases, we just need it to work. A lot of papers and books talk about the decompilation process and how to resolve undesired problems.

A list of documents related to decompilation may include:

- 13th chapter of the book *Reversing: Secrets of Reverse Engineering*, written by Eldad Eilam [6] ;
- Conference paper about *Reconstruction of instructions idioms in a retargetable decompiler*, by Jakub Křoustek and Fridolín Pokorný [18] ;
- All publications available on the RetDec website [34].

Also, some interesting projects have been achieved thanks to decompilation:

- **Mac 68k emulator**
Dynamic translating emulator for M68K code for Apple Macintoshes based on PowerPC [21].
- **StarCraft port to an ARM platform**
The Pandora console's community [26] generated an ARM version of the video game StartCraft thanks to static recompilation [29].
- **Dolphin emulator**
Dolphin emulates the GameCube and Wii consoles on PC thanks to Just-in-Time recompilation of PowerPC code to x86 and AArch64 [4].
- **x86 Intel CPUs**
Since their Pentium Pro CPU, CISC instructions are translated to more RISC-like internal micro-operations [28].

3 State of the art tools

This section reviews an intermediate representation named LLVM IR [20] and several decompilation tools and evaluates the best candidate to use for this project.

3.1 Intermediate representations

An intermediate representation is similar to a coding language. An IR is designed to be capable of representing the source code without loss of information and independent of any particular source and target. A compiler often translates a high-level code to an IR before compiling it into machine instructions.

Nowadays, there are two widely-used compilers: GCC [10] and LLVM [19]. Both of them offer an intermediate representation. Since most of the available decompilers – if not all – use the LLVM IR, we will not discuss GIMPLE [11], the GCC's IR.

The LLVM IR aims to be light-weight and low-level while being readable, typed and extensible. You can see it as a human-readable assembly language representation [20].

Let's pick up the same C and assembly codes shown in figure 2. The same program is translated as follow in LLVM IR:

```
%a = add i32 5, 0
%b = add i32 4, 0
%c = add i32 %a, %b
ret i32 %c
```

The first two lines store the values into two variables and the third line performs their sum. The last line returns the result. As you can see, the LLVM IR is easier to read as a human than the assembly code. Furthermore, it keeps track of the type of the variables, which is really harder in the assembly form.

3.2 Existing tools

The D-Neliac decompiler [13], built in the 1960s, was the first decompiler to prove that decompilation is feasible. Since this period, many projects have tried to offer correct decompilation.

Today, there are some interesting projects in the wild. Unfortunately, most of them are not open-source¹ or depend on proprietary software² (e.g. IDA-Pro [15]) thus they can't be used for our project.

rev.ng

rev.ng [37] is a suite of tools for binary analysis based on QEMU [33] and LLVM. It is (was) developed by Alessandro Di Federico, a former PhD student at Politecnico di Milano [30]. The project is open-source and licensed under GPLv2 [12]. Each individual file is released under the terms of the MIT License [22].

This project relies on a few components but the most interesting one is its static binary translator. Provided an input ELF binary, it will analyse it and emit an equivalent LLVM IR. The currently supported architectures are MIPS, ARM and x86-64.

The main issue with this project is that it was last updated more than ten months ago, so it seems like it has been abandoned. Furthermore, this project was maintained by only one developer, which is quite a small team.

We ran some tests to see how good it decompiles a binary. Unfortunately, most of the time rev.ng would crash or produce empty output.

¹Type of computer software whose source code is released under a license in which the copyright holder grants users the rights to study, change and distribute the software to anyone and for any purpose. — Wikipedia

²Proprietary software is non-free computer software for which the software's publisher or another person retains intellectual property rights—usually copyright of the source code, but sometimes patent rights. — Wikipedia

RetDec

RetDec [34], for *Retargetable Decompiler*, is an open-source machine-code decompiler based on LLVM [19]. It is being developed by the famous company Avast Software [1] and is licensed under the MIT License [22]. Its development was internal to Avast for several years but in February 2018, they decided to release the code publicly. It currently supports the following architectures: x86, ARM, MIPS, PIC32 and PowerPC. RetDec is also being actively developed by at least three people.

This project is composed of a plethora of libraries but we will mainly focus on the `bin2llvmmir` library which aims to translate binaries into LLVM IR modules. The project also includes a tool named `bin2llvmmirtool` which is a front-end for the `bin2llvmmir` library.

The main feature that offers RetDec against other decompilers is that it is retargetable. That means that – thanks to the ISAC architecture description language [14], also developed by the RetDec team – it is not necessary to manually reconfigure the decompiler for a new architecture, making it compatible with all machines (but not all architectures, as stated before).

BOLT

The BOLT project [25] was developed by a Facebook team and interns. It aims to boost the performance of 64-bit ELF applications by implementing a post-link optimizer. BOLT is built on top of LLVM and its optimization techniques are based on reorganizing code so caches suffer less fragmentation and at reordering basic blocks to relieve pressure from the branch predictor unit of a processor.

BOLT was deployed in Facebook data-centers and improvements ranging between 2% and 8% were observed, which is quite remarkable giving the fact that data-centers' applications are already highly optimized. These optimizations are really important since they reduce energy consumption thus reducing environmental impacts and costs.

Even if BOLT is a really promising and interesting project, it does not decompile code. It disassembles it and constructs a control-flow graph based on the disassembled code. However, this project might be helpful for inspiration about our project's architecture and the techniques that might be used to reach our aim.

Conclusion

After this analysis, we chose RetDec as the best candidate for this project since its development is active and led by a professional team. The MIT License also allows us to freely use their code which could be useful for the development of this project.

Some useful resources are available on their website *retdec.com* such as publications and presentations.

Previously, we stated that RetDec only supports 32-bit architectures. That's right for its official version but the 64-bit decompilation can be enabled by switching to another version of RetDec, even if it's not ready yet. For this document, we will use the RetDec version which is *capable* of decompiling both 32- and 64-bit architectures.

Since our project will mostly rely on the decompilation done by RetDec, we will analyse in details how it does decompile machine code.

RetDec is composed of two main parts: the *pre-processing* and the *core*.

The pre-processing part is responsible of analysing the binary program to produce an image that will be later used by the core.

The first step of the pre-processing is to discover the format of the binary program. RetDec supports the following formats: ELF, HEX, PE, COFF, Mach-O and raw binary. Then, a uniform binary representation is produced. Thanks to this uniform representation, the other parts of RetDec don't need to care about the original format of the program. The next step focuses on passing the uniform representation to an *image loader* library which will emulate a loader. This is important since depending on the loader used, the data loaded into memory can look different than the data in the original binary.

The image is now emitted. RetDec will perform an additional step by looking for available debugging

information. This will help to produce a better LLVM IR but most of proprietary programs ship with no debug information to avoid being analysed and possibility modified by hackers.

The purpose of the core is to lift LLVM IR code from the image produced by the pre-processing.

To do so, it starts by doing initialization passes to perform dead global elimination, constant propagation, inlining, loop optimization, etc... in the machine code. Then, RetDec calls an third-party framework called Capstone [3], which actually lifts instructions to the LLVM intermediate representation. Afterwards, some low-level passes are performed to identify global and local variables, functions' arguments and return type, data types and so on. Finally, some high-level LLVM passes are done on the LLVM IR already at disposition for final touches. Then, the LLVM IR is emitted by RetDec for future use.

We had the chance to attend the *Pass the SALT* 2018 [27] conference, where the RetDec team gave a talk. Thanks to this talk, we had more overview about the RetDec architecture and how it internally works.

4 Successful recompilation

Now that we have reached the first objective of this project by identifying a decompilation tool that suits our needs, we can focus on the second objective: *find a way to successfully recompile a given function from a simple program*.

First of all, we code a function which simply adds two unsigned integers and returns the result:

```
simple.c
uint32_t simple_add(uint32_t *a, uint32_t *b) {
    return *a + *b;
}
```

Then, we create a simple main function which takes two integers as arguments, calls the `simple_add()` method and displays the result.

```
main.c
int main(int argc, char** argv) {
    uint32_t a, b;

    a = atoi(argv[1]);
    b = atoi(argv[2]);

    fprintf(stdout, "Result: %u\n", simple_add(&a, &b));

    return EXIT_SUCCESS;
}
```

Now that we have a fully functional program, we compile it:

```
gcc -std=c99 -Wall -Werror -pedantic -Iinclude -m32 main.c simple.c -o main
```

We now have a 32-bit binary. The most interesting part is not the compilation, but the recompilation. So let's get started. We invoke RetDec decompilation script to get the LLVM IR of the `simple_add()` method:

```
retdec-decompiler.sh --stop-after bin2llvmir --select-functions simple_add ./main
```

The command above calls the RetDec decompilation script and asks it to stop after it has translated the binary (`./main`) into the LLVM IR (`--stop-after bin2llvmir`). We also ask it to only decompile the `simple_add()` method (`--select-functions simple_add`). The following code is the resulting LLVM IR:

```
main.c.backend.ll
define i32 @simple_add(i32* %arg1, i32 %arg2) local_unnamed_addr {
entry:
    %v2_57d = load i32, i32* %arg1, align 4
    %v1_582 = inttoptr i32 %arg2 to i32*
    %v2_582 = load i32, i32* %v1_582, align 4
    %v2_584 = add i32 %v2_582, %v2_57d
    ret i32 %v2_584
}
```

You might have noted that the method's signature is not the same as the one defined in `simple.c`. It is an acknowledged bug that should be fixed in the future [35]. Apart from the signature, everything appears to be in order, which is already quite motivating. We will now recompile the LLVM IR into a new object file and link it with the previously created object file `main.o`.

```
llc-5.0 -march=x86 main.c.backend.ll -o main.c.backend.s
gcc -m32 -c main.c.backend.s -o main.c.backend.o
gcc -m32 main.o main.c.backend.o -o ./main.translated
```

The first line invokes the LLVM static compiler, which can compile LLVM IR code. We pass the option `-march=x86` to tell the compiler to compile into a 32-bit version. This is required since RetDec can only decompile 32-bit code at the moment. The resulting code is assembly code.

The second line calls GCC, a widely-used compiler, to compile the assembly code into an object file. Again, the `-m32` flag asks the compiler to produce a 32-bit version of the binary.

The last line links the existing `main.o` and the newly created `main.c.backend.o` object files and creates a new binary named `main.translated`. By running it, we acknowledge that it still does its job:

```
./main.translated 5 4
Result: 9
```

This result demonstrates that a simple function can successfully be recompiled and linked with an existing file object, which is exactly the second objective of this project. Quite exciting!

Since the recompilation process requires quite a few commands, we created a `Makefile` to automatize the process. The source of the `Makefile` and of the `main.c` and `simple.c` files are available in appendix C.

5 Automatization tool

This section reviews the development of the third objective: *develop a program to automatize the recompilation process*. The tool is composed of a library, `liboptimizer` and a program to call the library, called `optimizer`.

5.1 Design

Above all, we have to define what the tool will do and how we will achieve it. The program will take several arguments for execution:

1. The function name to optimize or its offset ;
2. The target binary to execute ;
3. The arguments of the target binary.

The first argument is either the function name or its offset. When the function name is passed to the tool, we have to discover the offset of the function in the target binary and, reciprocally, when the function offset is passed to the tool, we must compute its symbol. So the first step of the library is to implement an ELF parser. The part of the library that will achieve this task will be referred as `elfparser`.

Once the target ELF has been parsed and that we have enough information, the next step is to call the RetDec's decompilation scripts to retrieve the LLVM IR of the function to optimize. This part is called `retdec`.

Then we must recompile it with a JIT compiler. The part of the library in charge of doing the JIT recompilation is named `jit`.

Finally, once we have the machine code of the recompiled function, the final step is to patch the target binary memory space so that it executes the new optimized function instead of the old one. To do that, we implemented a third part named `live-patcher`.

Each of these parts will be further discussed later. The figure 4 shows a simple overview of how the library is structured.

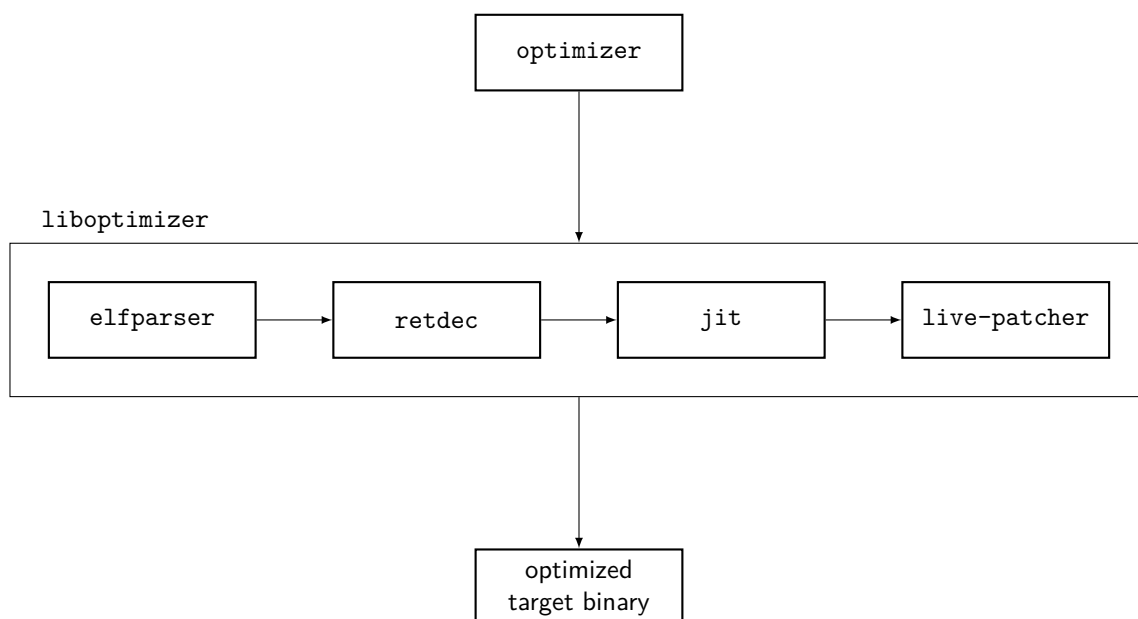


Figure 4: `liboptimizer` pipeline

5.2 elfparser

As said above, this part of the library is responsible for parsing the ELF binary passed to the library. Its purpose is to offer the possibility to retrieve the offset of a given symbol and vice-versa, given an offset, retrieve the corresponding symbol.

To implement this, some previous work is required: the ELF header has to be parsed so that we have information about the binary. The most wanted information to know is the class of the ELF file: is it 32-bit or 64-bit? Indeed, this information is the most important since the ELF is not structured the same way for the two classes. The first bytes of the ELF header are the same for 32- and 64-bit so we assume the ELF is 32-bit and then we adapt. The following `structs` describe the structure of an ELF header for 32-bit (left) and 64-bit (right):

```

struct elf32_hdr
typedef struct elf32_hdr {
    unsigned char e_ident[16];
    Elf32_Half    e_type;
    Elf32_Half    e_machine;
    Elf32_Word    e_version;
    Elf32_Addr    e_entry;
    Elf32_Off     e_phoff;
    Elf32_Off     e_shoff;
    Elf32_Word    e_flags;
    Elf32_Half    e_ehsize;
    Elf32_Half    e_phentsize;
    Elf32_Half    e_phnum;
    Elf32_Half    e_shentsize;
    Elf32_Half    e_shnum;
    Elf32_Half    e_shstrndx;
} Elf32_Ehdr;

```

```

struct elf64_hdr
typedef struct elf64_hdr {
    unsigned char e_ident[16];
    Elf64_Half    e_type;
    Elf64_Half    e_machine;
    Elf64_Word    e_version;
    Elf64_Addr    e_entry;
    Elf64_Off     e_phoff;
    Elf64_Off     e_shoff;
    Elf64_Word    e_flags;
    Elf64_Half    e_ehsize;
    Elf64_Half    e_phentsize;
    Elf64_Half    e_phnum;
    Elf64_Half    e_shentsize;
    Elf64_Half    e_shnum;
    Elf64_Half    e_shstrndx;
} Elf64_Ehdr;

```

And the following `typedefs` define the types used for the above `structs`:

```

Base types for 32-bit
typedef uint32_t    Elf32_Addr;
typedef uint16_t    Elf32_Half;
typedef uint32_t    Elf32_Off;
typedef int32_t     Elf32_Sword;
typedef uint32_t     Elf32_Word;

```

```

Base types for 64-bit
typedef uint64_t    Elf64_Addr;
typedef uint16_t    Elf64_Half;
typedef int16_t     Elf64_SHalf;
typedef uint64_t    Elf64_Off;
typedef int32_t     Elf64_Sword;
typedef uint32_t     Elf64_Word;
typedef uint64_t    Elf64_Xword;
typedef int64_t     Elf64_Sxword;

```

Thanks to the fields `e_shoff`, `e_shnum`, `e_shentsize`, we obtain information about the location of *Section Header Table*, the number of entries it contains and their size.

Additionally, the index of the *String Table* in the *Section Header Table* is stored into the field `e_shstrndx`. Thanks to this index, we already know where the *String Table* is located. The last information we need is the location of the *Symbol Table*.

The *Section Header Table* contains all the information necessary to locate each ELF section. But we only really need one section: the *Symbol Table*. An ELF file contains only one *Symbol Table* and the section has a unique type to identify it. This type has the value 2.

```

32-bit Section Header
typedef struct elf32_shdr {
    Elf32_Word    sh_name;
    Elf32_Word    sh_type;
    Elf32_Word    sh_flags;
    Elf32_Addr    sh_addr;
    Elf32_Off     sh_offset;
    Elf32_Word    sh_size;
    Elf32_Word    sh_link;
    Elf32_Word    sh_info;
    Elf32_Word    sh_addralign;
    Elf32_Word    sh_entsize;
} Elf32_Shdr;

```

```

64-bit Section Header
typedef struct elf64_shdr {
    Elf64_Word    sh_name;
    Elf64_Word    sh_type;
    Elf64_Xword   sh_flags;
    Elf64_Addr    sh_addr;
    Elf64_Off     sh_offset;
    Elf64_Xword   sh_size;
    Elf64_Word    sh_link;
    Elf64_Word    sh_info;
    Elf64_Xword   sh_addralign;
    Elf64_Xword   sh_entsize;
} Elf64_Shdr;

```

The field `sh_type` contains the type of the section. So we iterate over the table until we find a section with the type 2. The location of the symbols is stored by the field `sh_offset`. So what is left to do is to look at

the symbols location, iterate over the list until we find the symbol we are interested in. A symbol is structured as follow:

```
32-bit symbol
typedef struct elf32_sym{
    Elf32_Word    st_name;
    Elf32_Addr    st_value;
    Elf32_Word    st_size;
    unsigned char st_info;
    unsigned char st_other;
    Elf32_Half    st_shndx;
} Elf32_Sym;
```

```
64-bit symbol
typedef struct elf64_sym {
    Elf64_Word    st_name;
    unsigned char st_info;
    unsigned char st_other;
    Elf64_Half    st_shndx;
    Elf64_Addr    st_value;
    Elf64_Xword   st_size;
} Elf64_Sym;
```

The field `st_name` contains an offset from the *String Table* where the symbol character string is located. The field `st_value` contains the offset of the symbol in the ELF file. So if we want to retrieve an offset from a given symbol, we iterate over all symbols until we find the corresponding character string and return the offset.

On the contrary, if we want to retrieve a symbol from a given offset, we iterate over all symbols until we find the corresponding offset and return the character string.

Summary

Developing this part of the library was not really difficult. The main *issue* was that there is a lot of duplicate code since 32-bit and 64-bit ELF have different structures.

The code relative to `elfparser` is available in appendix D.1.

5.3 retdec

`retdec` is responsible for calling the RetDec's decompilation scripts and retrieve the resulting LLVM IR file.

Since RetDec's scripts generate several files, we chose to create a temporary directory which is deleted once `liboptimizer` is done running. Then, the RetDec's scripts are called with corresponding arguments: stopping after the generation of the LLVM IR, the temporary directory as working directory, the name of the function to decompile and the target binary.

Summary

This part of the library is very lightweight so no problem were encountered during its development.

The code relative to `retdec` is available in appendix D.2.

5.4 jit

The purpose of the library `jit` part is to implement a JIT compiler so we can recompile the target function. It uses LLVM's On-Request-Compilation (ORC) APIs [2].

Since the offered APIs are really simple to use, the implementation is straight-forward. We simply initialize an `ExecutionEngine` [9], parse the LLVM IR file containing the function to optimize and tell the engine to compile it. It returns a pointer on the optimized function so we can use it later.

We have the optimized function machine code at our disposal and that's great. But we need one more information: the size (in bytes) of the optimized function. Unfortunately, the `ExecutionEngine` class doesn't provide a *simple* way to get the size.

The solution is to implement a subclass of the `JITEventListener` [17] class and register this new listener so that every time the JIT compiles a function, it notifies its registered listeners with more information than just a pointer to the new function's machine code.

Summary

The code for this part of `liboptimizer` is really simple thanks to the LLVM APIs. The main problem we came across was to get the size of the compiled function but after some research we achieve to overcome the issue.

The code relative to the JIT front-end and the `JITEventListener` subclass is available in appendix D.3.

5.5 live-patcher

The aim of the `live-patcher` is to modify the target process memory space so that it calls the optimized function instead of the old one during its execution.

It mostly relies on the `ptrace` [32] system call [38] which allows a *tracer* process (in our case, `liboptimizer`) to observe and control the execution of a *tracee* (the target binary).

First things first, we have to attach the tracer to the target process and stop its execution so that we can modify its memory.

The next step is to allocate a new memory segment into the target process memory space. This segment is used to store the machine code of the optimized function. To do so, we must inject a `mmap2` [23] system call. The injection consists of modifying some tracee registers with pre-defined values to ensure that we have the correct access rights, enough allocated space, etc... After the system call is injected, we retrieve the address of the newly allocated memory segment by reading back the value of the `RAX` register.

Now that we have a memory segment that we can execute and write into as we wish, we can write the optimized function's machine code into it.

The last but not least step is to hook the old function so that the process executes the optimized one. The hook consists of replacing the first bytes of the old function with an unconditional jump to the location of the optimized function, which is stored at the address of the memory segment we got before.

For the sake of clarity, we created two simple macros that allow to create the assembly code for 32- and 64-bit programs. The details of these macros are available in appendix D.4.

```
32-bit hook
unsigned char jump_32[] =
↳ MAKE_JUMP32(process->freeselement_address);
```

```
64-bit hook
unsigned char jump_64[] =
↳ MAKE_JUMP64(process->freeselement_address);
```

These macros take as argument the address of the memory segment and then append machine instructions to create a hook.

There are only two instructions: the first one loads the address of the memory segment into a register and the last one tells the processor to jump unconditionally to the location stored into the same register. Since we support both 32- and 64-bit binaries, we have two different jumps because the address length and the instruction sets are different. After all these memory replacements, we can finally let the tracee process continue its execution with the optimized function.

Summary

This part of the library was the most difficult but also the most interesting one to develop. It required to dig at a very low level into the tracee memory space, reading instructions byte per byte, find a way to inject a system call, etc... To ease debugging, we had to develop our very own debugging function. We learned a lot developing this part and the effort was rewarding.

The code relative to the `live-patcher` part is available in appendix D.4.

5.6 liboptimizer

All those parts work great individually but we need to wire them up so we can expose them to any user that wants to use the library. To keep things simple, we offer no more than four methods and a `struct` that is meant to contain all the information that is needed.

```
process_info_t
typedef struct {
    const char *path;
    int argc;
    char **argv;
    pid_t pid;
    const char *function_name;
    uint64_t function_offset;
    uint64_t codesegment_address;
    uint64_t freesegment_address;
    uint8_t *optimized_function;
    size_t optimized_function_size;
    uint8_t is64;
} process_info_t;
```

- `*path`: Path of the target binary.
- `argc`: Arguments count of the target binary.
- `**argv`: Array of the arguments of the target binary.
- `pid`: The Process ID of the target binary.
- `*function_name`: The name of the function to optimize.
- `function_offset`: The offset of the function to optimize.
- `codesegment_address`: Address of the Code Segment of the target process.
- `freesegment_address`: Address of the newly allocated memory chunk.
- `*optimized_function`: Pointer to the optimized function machine code.
- `optimized_function_size`: Size in bytes of the optimized function.
- `is64`: Is the target process 32- or 64-bit?

Library methods:

- `char *symbol_at_address(const char *path, uint64_t address)`
Resolves the symbol of the function at the given address in the binary located at `path`. Allocates and returns a null-terminated string containing the symbol.
Returns `NULL` on error.
- `process_info_t *init_process(int argc, char **argv, const char *function_name)`
Given `argc` and `argv` of the binary and the name of the function to optimize, allocates and returns a pointer to a `process_info_t` that contains basic information about the process.
Returns `NULL` on error.
This method calls the RetDec's scripts to decompile the target function, pass the resulting LLVM IR file to the JIT, retrieve the optimized function and attach the target binary so its memory space can be modified.
- `int modify_process(process_info_t *process)`
Modifies the memory of the associated process of the argument. Basically hooks the function to optimize with the optimized one.
Returns 0 on success, 1 otherwise.
- `int execute_process(process_info_t *process, bool wait_for_exit)`
Starts the execution of the process. If `wait_for_exit` is `true`, waits for the process to exit and returns its exit status. If `wait_for_exit` is `false`, doesn't wait for the process to exit and returns 0 on success, 1 otherwise.

5.7 optimizer

Now that we have a library that works, we need a small program to show how to call it. The following code shows a simple use case of liboptimizer.

```
optimizer.c

/*
 * File: optimizer.c
 *
 * Created by: Lucas Elisei <lucas.elisei@heig-vd.ch>
 *
 * Shows how to use the liboptimizer library.
 */

#include <stdio.h>
#include <stdlib.h>
#include <string.h>

#include "debug.h"
#include "liboptimizer.h"

static void print_help(char **argv) {
    fprintf(stderr, "Usage: %s {-f <function_name>, -a <function_offset>} <target_bin>
    ↪ [target_args]\n", argv[0]);
}

int main(int argc, char **argv) {
    process_info_t *process;
    char *function_name;
    uint64_t function_offset;
    int rc;

    if (argc < 4) {
        print_help(argv);

        return EXIT_FAILURE;
    }

    if (strcmp(argv[1], "-a") == 0) {
        function_offset = strtoull(argv[2], NULL, 16);
        function_name = symbol_at_address(argv[3], function_offset);
    }
    else if (strcmp(argv[1], "-f") == 0) {
        function_name = argv[2];
    }
    else {
        fprintf(stderr, "Unrecognized option: %s\n", argv[1]);
        print_help(argv);

        return EXIT_FAILURE;
    }

    process = init_process(argc - 3, argv + 3, function_name);
    if (process == NULL) {
        fprintf(stderr, "[optimizer] ERROR: Error during initialization.\n");

        return EXIT_FAILURE;
    }

#ifdef LIBOPTIMIZER_DEBUG
    print_process_info(process);
#endif

    rc = modify_process(process);
    if (rc < 0) {
        fprintf(stderr, "[optimizer] ERROR: Error while modifying process.\n");

        return EXIT_FAILURE;
    }

    rc = execute_process(process, true);
    if (rc < 0) {
        fprintf(stderr, "[optimizer] ERROR: Error while executing process.\n");
    }
}
```

```
        return EXIT_FAILURE;
    }

    return EXIT_SUCCESS;
}
```

This simple program works as follow: the first argument is either the name of the function to decompile (specified by `-f`) or its offset (specified by `-a`). The rest of the arguments are the binary name and its arguments.

Remember the small program we decompiled in section 4? If we pass it to `optimizer` to decompile the `simple_add` method, it should look like something like the following:

```
./optimizer -f simple_add main 5 4
```

6 Examples

This section shows examples of using the `liboptimizer` and limitations of the `tool` (SPOILER: its dependencies).

Because we support both 32- and 64-bit binaries, we will perform those examples for both.

6.1 Simple addition

The first example focuses on a simple function that sums two integers and prints the result. The code of the target program is the following:

```
simple_add.c
/*
 * File: simple_add.c
 *
 * Created by: Lucas Elisei <lucas.elisei@heig-ud.ch>
 *
 * Sums two integers and prints the result.
 */

#include <stdio.h>
#include <stdlib.h>
#include <unistd.h>

int simple_add(int a, int b) {
    return a + b;
}

int main(int argc, char **argv) {
    int a, b, rc;

    if (argc != 3) {
        fprintf(stderr, "[simple_add] Usage: %s <a> <b>\n", argv[0]);

        return EXIT_FAILURE;
    }

    // Retrieve arguments.
    a = atoi(argv[1]);
    b = atoi(argv[2]);

    // Sum the arguments.
    rc = simple_add(a, b);

    // Print the result.
    fprintf(stdout, "[simple_add] Result: %d\n", rc);

    return EXIT_SUCCESS;
}
```

For this first example, we will not recompile the `simple_add` function to optimize the program but to see if `liboptimizer` does its job. The function is too simple to measure its impact when optimized.

The 32- and 64-bit versions of this program are respectively named `simple_add32` and `simple_add64`.

6.1.1 32-bit

To compile the `simple_add32` program, we use the following command:

```
gcc -Wall -Werror -O0 -m32 simple_add.c -o simple_add32
```

The option `-O0` tells the compiler to avoid doing optimization. This way, we will clearly see if the function has been optimized after being processed by `liboptimizer`.

Below, the machine code corresponding to the non-optimized `simple_add` function:


```

000011bd <simple_add>:
 11bd: 55                push    %ebp
 11be: 89 e5            mov     %esp,%ebp
 11c0: e8 cc 00 00 00   call    1291 <__x86.get_pc_thunk.ax>
 11c5: 05 3b 2e 00 00   add     $0x2e3b,%eax
 11ca: 8b 55 08         mov     0x8(%ebp),%edx
 11cd: 8b 45 0c         mov     0xc(%ebp),%eax
 11d0: 01 d0           add     %edx,%eax
 11d2: 5d              pop     %ebp
 11d3: c3              ret

```

We see that the size of the function is 22 bytes and is composed of 9 instructions (one of which is a call to another function – `__x86.get_pc_thunk.ax`). Now, we call the optimizer program to optimize the function.

```
./optimizer -f simple_add simple_add32 3 2
```

We are greeted with the message `[simple_add] Result: 5` which means that it returned the correct result, and that's great! Now, let's take a look to the machine code of the optimized function:

```

f7f70000 <simple_add>:
 0: 8b 44 24 08      mov     0x8(%esp),%eax
 4: 03 44 24 04      add     0x4(%esp),%eax
 8: c3              ret

```

We clearly see that the function has been optimized! The size of the `simple_add` function is now of 9 bytes and is composed of only 3 instructions.

6.1.2 64-bit

To compile the `simple_add64` program, we use the following command:

```
gcc -Wall -Werror -O0 simple_add.c -o simple_add64
```

Below, the machine code corresponding to the non-optimized `simple_add` function:

```

000000000000116a <simple_add>:
 116a: 55                push    %rbp
 116b: 48 89 e5          mov     %rsp,%rbp
 116e: 89 7d fc          mov     %edi,-0x4(%rbp)
 1171: 89 75 f8          mov     %esi,-0x8(%rbp)
 1174: 8b 55 fc          mov     -0x4(%rbp),%edx
 1177: 8b 45 f8          mov     -0x8(%rbp),%eax
 117a: 01 d0           add     %edx,%eax
 117c: 5d              pop     %rbp
 117d: c3              retq

```

Before optimization, the size of the function is 19 bytes and is composed of 9 instructions. Now, we call the optimizer program with the same arguments as for the 32-bit version (only replacing `simple_add32` with `simple_add64`).

This time, the optimized program doesn't print any message... That's not good. Let's analyse the resulting machine code of the optimized function:

```

00007f0c7a90b000 <simple_add>:
 0: 48 c7 44 24 e8 00 00 movq    $0x0,-0x18(%rsp)
 7: 00 00
 9: 48 8b 44 24 e0      mov     -0x20(%rsp),%rax
 e: 48 89 04 25 f8 ff ff mov     %rax,0xfffffffffffff8
15: ff
16: 48 c7 44 24 e0 f8 ff movq    $0xfffffffffffff8,-0x20(%rsp)
1d: ff ff
1f: 8b 44 24 f8         mov     -0x8(%rsp),%eax

```

```

23: 89 04 25 f4 ff ff ff    mov    %eax,0xfffffffffffffff4
2a: 8b 44 24 f0             mov    -0x10(%rsp),%eax
2e: 48 8b 4c 24 e0          mov    -0x20(%rsp),%rcx
33: 89 41 f8                mov    %eax,-0x8(%rcx)
36: 48 8b 4c 24 e0          mov    -0x20(%rsp),%rcx
3b: 8b 41 f8                mov    -0x8(%rcx),%eax
3e: 03 41 fc                add    -0x4(%rcx),%eax
41: 48 8b 4c 24 e8          mov    -0x18(%rsp),%rcx
46: 48 8b 09                mov    (%rcx),%rcx
49: 48 89 4c 24 e0          mov    %rcx,-0x20(%rsp)
4e: c3                     retq

```

First of all, the *optimized* function is bigger than the original one, quite odd. Second, the second instruction (at the offset 0x7), is invalid. Its opcode [24] doesn't correspond to any instruction that the processor supports... We can also see that at the offsets 0x15 and 0x1d, the opcode are not recognized.

To be sure of what causes the problem, we look at the optimized process' execution instruction per instruction. When the program reach the instruction at the offset 0x9, it crashes. So the problem was effectively this instruction (and the other two might cause a problem too).

The culprit is RetDec. As said in the section 3, the 64-bit version of RetDec has been used but it doesn't work as great as 32-bit decompilation. If we take a look at the LLVM IR generated by RetDec, we clearly see that there was some misunderstanding:

```

define i64 @simple_add() local_unnamed_addr {
dec_label_pc_116a:
    %rbp.global-to-local = alloca i64, align 8
    %rdi.global-to-local = alloca i64, align 8
    %rsi.global-to-local = alloca i64, align 8
    %rsp.global-to-local = alloca i64, align 8
    store i64 0, i64* %rsp.global-to-local, align 8
    %v0_116a = load i64, i64* %rbp.global-to-local, align 8
    %v1_116a = load i64, i64* %rsp.global-to-local, align 8
    %v2_116a = add i64 %v1_116a, -8
    %v3_116a = inttoptr i64 %v2_116a to i64*
    store i64 %v0_116a, i64* %v3_116a, align 8
    store i64 %v2_116a, i64* %rbp.global-to-local, align 8
    %v0_116e = load i64, i64* %rdi.global-to-local, align 8
    %v1_116e = trunc i64 %v0_116e to i32
    %v3_116e = add i64 %v1_116e, -12
    %v4_116e = inttoptr i64 %v3_116e to i32*
    store i32 %v1_116e, i32* %v4_116e, align 4
    %v0_1171 = load i64, i64* %rsi.global-to-local, align 8
    %v1_1171 = trunc i64 %v0_1171 to i32
    %v2_1171 = load i64, i64* %rbp.global-to-local, align 8
    %v3_1171 = add i64 %v2_1171, -8
    %v4_1171 = inttoptr i64 %v3_1171 to i32*
    store i32 %v1_1171, i32* %v4_1171, align 4
    %v0_1174 = load i64, i64* %rbp.global-to-local, align 8
    %v1_1174 = add i64 %v0_1174, -4
    %v2_1174 = inttoptr i64 %v1_1174 to i32*
    %v3_1174 = load i32, i32* %v2_1174, align 4
    %v1_1177 = add i64 %v0_1174, -8
    %v2_1177 = inttoptr i64 %v1_1177 to i32*
    %v3_1177 = load i32, i32* %v2_1177, align 4
    %v4_117a = add i32 %v3_1177, %v3_1174
    %v20_117a = zext i32 %v4_117a to i64
    %v0_117c = load i64, i64* %rsp.global-to-local, align 8
    %v1_117c = inttoptr i64 %v0_117c to i64*
    %v2_117c = load i64, i64* %v1_117c, align 8
    store i64 %v2_117c, i64* %rbp.global-to-local, align 8
    ret i64 %v20_117a

;    uselistorder directives
    uselistorder i64* %rbp.global-to-local, { 0, 2, 3, 4, 1 }
    uselistorder i64 -8, { 1, 2, 0 }
    uselistorder i32 1, { 0, 2, 1, 3 }
}

```

The LLVM IR is really big for a simple addition and the function signature is not interpreted correctly: RetDec seems to think that the function requires no arguments.

Conclusion

In conclusion for this simple addition function, the 32-bit version works great but the 64-bit doesn't because of RetDec. We consider this a normal behaviour since we use a modified version of RetDec with enabled 64-bit decompilation that is not officially supported. There is no solution but waiting for an official 64-bit support or implementing it, but that's not the scope of our work.

The 32-bit version could be optimized even more by implementing constant propagation, which means that the optimized function would pre-calculate the result and simply return the 5 instead of doing the operation itself.

6.2 Simple multiplication

This example is quite the same as the first one but this time, the target program multiplies two integers. The code is the following:

```

                                simple_mul.c
/*
 * File: simple_mul.c
 *
 * Created by: Lucas Elisei <lucas.elisei@heig-vd.ch>
 *
 * Multiplies two integers and prints the result.
 */

#include <stdio.h>
#include <stdlib.h>
#include <unistd.h>

int simple_mul(int a, int b) {
    return a * b;
}

int main(int argc, char **argv) {
    int a, b, rc;

    if (argc != 3) {
        fprintf(stderr, "[simple_mul] Usage: %s <a> <b>\n", argv[0]);

        return EXIT_FAILURE;
    }

    // Retrieve the arguments.
    a = atoi(argv[1]);
    b = atoi(argv[2]);

    // Multiply the arguments.
    rc = simple_mul(a, b);

    // Print the result.
    fprintf(stdout, "[simple_mul] Result: %d\n", rc);

    return EXIT_SUCCESS;
}

```

The 32- and 64-bit versions of this program are respectively named `simple_mul32` and `simple_mul64`.

6.2.1 32-bit

For compilation, the command is the same as the one used in the first example. Below, the machine code corresponding to the non-optimized `simple_mul` function:

```

000011bd <simple_mul>:
    11bd: 55                push    %ebp
    11be: 89 e5            mov     %esp,%ebp
    11c0: e8 cb 00 00 00   call   1290 <__x86.get_pc_thunk.ax>

```

```

11c5: 05 3b 2e 00 00      add    $0x2e3b,%eax
11ca: 8b 45 08             mov    0x8(%ebp),%eax
11cd: 0f af 45 0c          imul   0xc(%ebp),%eax
11d1: 5d                   pop    %ebp
11d2: c3                   ret

```

We see that the size of the function is 21 bytes and is composed of 8 instructions. Now, we call the optimizer program to optimize the function:

```
./optimizer -f simple_mul simple_mul32 3 2
```

The program prints [simple_mul] Result: 0, which is not really the result we expected. Let's take a look at the optimized function's machine code:

```

f7f66000 <simple_mul>:
0: 8b 44 24 04          mov    0x4(%esp),%eax
4: 0f af 44 24 0c        imul   0xc(%esp),%eax
9: c3                   ret

```

At first glance, the machine code seems to not contain any error. But if we look closely at the second instruction, which loads an argument on the stack, the offset is wrong. It is 0xc and should be 0x8 because the function arguments are 4-byte long, and not 8. There are two possibilities:

1. The compiler thinks that there are two arguments: one is 4-byte long and the other is 8-byte long ;
2. The compiler thinks that there are three arguments, each 4-byte long, and the second one is useless.

Taking a look at the generated LLVM IR from RetDec might help us solve the mystery. Here the LLVM IR representing the simple_mul function:

```

1  define i32 @simple_mul(i64 %arg1, i32 %arg2) local_unnamed_addr {
2  dec_label_pc_11bd:
3      %v4_11ca = trunc i64 %arg1 to i32
4      %v7_11cd = mul i32 %v4_11ca, %arg2
5      ret i32 %v7_11cd
6  }

```

If we look at the function's signature, we can see that it is wrong. The first argument is interpreted as 64-bit integer, which is false. At line 3, we see that the decompiler casts the first argument to a 32-bit integer. So RetDec gets the final type right but not the arguments.

The result printed by the optimized binary can be explained as follow: the compiler retrieves one of the two arguments 4 bytes farther than expected, resulting in an unexpected value.

6.2.2 64-bit

To compile the simple_mul64 program, we use the following command:

```
gcc -Wall -Werror -O0 simple_mul.c -o simple_mul64
```

Below, the machine code corresponding to the non-optimized simple_mul function:

```

000000000000116a <simple_mul>:
116a: 55                   push   %rbp
116b: 48 89 e5             mov    %rsp,%rbp
116e: 89 7d fc             mov    %edi,-0x4(%rbp)
1171: 89 75 f8             mov    %esi,-0x8(%rbp)
1174: 8b 45 fc             mov    -0x4(%rbp),%eax
1177: 0f af 45 f8          imul   -0x8(%rbp),%eax
117b: 5d                   pop    %rbp
117c: c3                   retq

```

Before optimization, the size of the function is 18 bytes and is composed of 8 instructions. The next step is to call the optimizer program with the same arguments as for the 32-bit version:

```
./optimizer -f simple_mul simple_mul64 3 2
```

As for the 64-bit version of the first example, the optimized doesn't print any message. Let's take a look at the optimized function's machine code:

```
7f807a326000 <simple_mul>:
 0: 48 c7 44 24 e8 00 00    movq    $0x0,-0x18(%rsp)
 7: 00 00
 9: 48 8b 44 24 e0          mov     -0x20(%rsp),%rax
 e: 48 89 04 25 f8 ff ff    mov     %rax,0xffffffffffffff8
15: ff
16: 48 c7 44 24 e0 f8 ff    movq    $0xffffffffffffff8,-0x20(%rsp)
1d: ff ff
1f: 8b 44 24 f8            mov     -0x8(%rsp),%eax
23: 89 04 25 f4 ff ff ff    mov     %eax,0xffffffffffffff4
2a: 8b 44 24 f0            mov     -0x10(%rsp),%eax
2e: 48 8b 4c 24 e0          mov     -0x20(%rsp),%rcx
33: 89 41 f8                mov     %eax,-0x8(%rcx)
36: 48 8b 44 24 e0          mov     -0x20(%rsp),%rax
3b: 8b 48 fc                mov     -0x4(%rax),%ecx
3e: 48 63 40 f8            movslq  -0x8(%rax),%rax
42: 48 0f af c1            imul    %rcx,%rax
46: 48 8b 4c 24 e8          mov     -0x18(%rsp),%rcx
4b: 48 8b 09                mov     (%rcx),%rcx
4e: 48 89 4c 24 e0          mov     %rcx,-0x20(%rsp)
53: c3                      retq
```

We see the same behaviour as for the first example: one or more instructions are not recognized by the processor, resulting in the crash of the program.

Conclusion

The 32-bit version of the optimized binary does the good operation but with the wrong values since RetDec wrongly decompiles the function¹. For 64-bit, the behaviour is the same as the first example.

The implementation of constant propagation might fix the problem encountered for the 32-bit version.

6.3 Matrices multiplication

For the third and last example, we will focus on a bigger program. This program performs the multiplication of two randomly generated matrices. It takes their dimensions as arguments. At the end of its execution, it prints the time taken by the target function to be executed.

The target function is named `matrix_mult` and its complexity is $\mathcal{O}(n^3)$, which means that the running time cubic grows as the input size grows.

The 32- and 64-bit versions of this program are respectively named `mmult32` and `mmult64`. The code is available at appendix D.5.

The correctness of the resulting matrix has been verified every time the program has been run.

6.3.1 32-bit

To compile the `mmult32` program, we use the following command:

```
gcc -Wall -Werror -O0 -m32 mmult.c main.c -o mmult32
```

Because the program measures the time the target function took to execute, we first do a run of the non-optimized program:

¹We reported this bug on the RetDec's Github repository (issue #269).

```
./mmult32 1000 1000  
Elapsed time for matrix multiplication (1000x1000): 9s 357689098ns
```

As said above, the function is quite complex so its machine code is long and hard to decrypt. However, it is available at appendix D.6.

The non-optimized function is 298 bytes long and is composed of 96 instructions. Now, we call the optimizer program to optimize the function with matrices of dimension 1000x1000 (which means 1 billion operations).

```
./optimizer -f matrix_mult mmult32 1000 1000  
Elapsed time for matrix multiplication (1000x1000): 6s 260002895ns
```

Interesting: the target program doesn't crash, it prints a reasonable measurement and its value is less than the non-optimized run. Let's check if the resulting matrix is correct. And... yes, it is. Is this our first win?

Let's *joyfully* take a look at the optimized function machine code, available at the appendix D.6. It is 339 bytes long and is composed of 144 instructions. How come the optimized function is bigger than the normal one and still it is faster?

The instructions used might be the answer. Indeed, depending on their complexity, some instructions take more clock cycle to execute than other. For example, the optimized machine code contains 39 NOP instructions, which take only one clock cycle to execute.

Some measurements have been done on this example. The non-optimized binary and the optimized one have been run 30 times each and an average of the execution time for the `matrix_mult` function (on 1000x1000 matrices) has been calculated:

On average, the non-optimized function took 10.04 seconds to execute. On average, the optimized function took 7.46 seconds to execute. Hence, the optimized function shows a 25% increase in performance.

However, it is interesting to measure the total time that the optimizer took to optimize the target binary. On average, it took 1.66 to optimize. Hence, counting in the time for optimizer to execute, the runs still show a 10% increase in performance.

6.3.2 64-bit

As for the previous examples, the 64-bit *optimized* machine code contains undefined instructions that make the program crash.

Conclusion

Thanks to this last example, we tested `liboptimizer` on a target program that does a lot of calculation. No surprises for the 64-bit version: the decompilation is in its early stages so it doesn't give satisfactory results but the 32-bit version shows a 25% increase in performance which is a really promising result.

More optimization could be done thanks to an LLVM pass named Polly [31].

7 Conclusion

The Moore's law is coming to an end and emerging computer science fields like machine learning or artificial intelligence require a lot of computational power.

To tackle such problems, we propose a method to optimize the bottlenecks of any binaries by identifying them, recompiling them specifically for the platform they are being run on and do it without the target program knowing it.

Through examples, we saw that our solution can enhance a binary performance up to 25%. Unfortunately, the decompiler that exist nowadays are still in an early stage of development, thus limiting the capacities of such a method.

Some improvements can be made on the library we developed for this project. For instance, it could be possible to merge the solicited parts of RetDec and our project to produce a standalone application that would not depend on RetDec being installed on the machine to run the library.

Another feature could be to save the patched binary to avoid calling the library each time the target process is being run.

With the constant development of RetDec, it might be possible that the 64-bit support of RetDec could come in next months. Testing the library with this hypothetically new version of RetDec might demonstrate better results on 64-bit binaries.

8 Acknowledgements

I would like to thank Alberto Dassatti for supervising this project and introducing me to the wonderful possibilities of recompilation.

I also would like to thank the HEIG-VD for giving me the opportunity to attend the Pass the SALT conference to meet and discuss with the RetDec team.

A last thank to all my friends and my family for being supportive during this project, taking the time to give me feedback about my work and being awesome.

References

- [1] Avast Software, Inc. URL: <https://www.avast.com/>.
- [2] Building a JIT. URL: <https://llvm.org/docs/tutorial/BuildingAJIT1.html>.
- [3] Capstone Framework. URL: <https://www.capstone-engine.org>.
- [4] Dolphin Emulator. URL: <https://dolphin-emu.org/>.
- [5] L. Ďurfina et al. *Design of a Retargetable Decompiler for a Static Platform-Independent Malware Analysis*. URL: http://www.sersc.org/journals/IJSIA/vol5_no4_2011/8.pdf.
- [6] Eldad Eilam. *Reversing: Secrets of Reverse Engineering*. John Wiley & Sons, 2005. URL: https://archive.org/details/reversing-secrets-of-reverse-engineering_2.
- [7] *End of Moore's Law: It's not just about physics*. URL: <https://www.scientificamerican.com/article/end-of-moores-law-its-not-just-about-physics/>.
- [8] *Executable and Linkable Format*. In: Wikipedia. URL: https://en.wikipedia.org/w/index.php?title=Executable_and_Linkable_Format&oldid=844622391.
- [9] *ExecutionEngine class*. URL: http://llvm.org/doxygen/classllvm_1_1ExecutionEngine.html.
- [10] GCC, the GNU Compiler Collection. URL: <https://gcc.gnu.org/>.
- [11] GIMPLE. URL: <https://gcc.gnu.org/onlinedocs/gccint/GIMPLE.html>.
- [12] GPLv2. URL: <https://opensource.org/licenses/GPL-2.0>.
- [13] Maurice H. Halstead. *Machine-Independent Computer Programming*. Spartan Books, 1962.
- [14] Adam Husár et al. *Automatic C Compiler Generation from Architecture Description Language ISAC*. URL: <http://drops.dagstuhl.de/opus/volltexte/2011/3065/pdf/9.pdf>.
- [15] IDA. URL: <https://www.hex-rays.com/products/ida/>.
- [16] Hex-Rays SA Ilfak Guilfanov. *Decompilers and beyond*. 2008. URL: https://www.hex-rays.com/products/ida/support/ppt/decompilers_and_beyond_white_paper.pdf.
- [17] *JITEventListener class*. URL: http://llvm.org/doxygen/classllvm_1_1JITEventListener.html.
- [18] J. Křoustek and F. Pokorný. "Reconstruction of instruction idioms in a retargetable decompiler". In: *2013 Federated Conference on Computer Science and Information Systems*. Sept. 2013, pp. 1519–1526.
- [19] LLVM. URL: <https://llvm.org>.
- [20] *LLVM Intermediate Representation*. URL: <https://llvm.org/docs/LangRef.html>.
- [21] *Mac 68k emulator*. URL: https://en.wikipedia.org/wiki/Mac_68k_emulator.
- [22] MIT License. URL: <https://opensource.org/licenses/MIT>.
- [23] *mmap2*. URL: <http://man7.org/linux/man-pages/man2/mmap2.2.html>.
- [24] *Opcode*. URL: <https://en.wikipedia.org/wiki/Opcode>.
- [25] Maksim Panchenko et al. *BOLT: A Practical Binary Optimizer for Data Centers and Beyond*. URL: <https://arxiv.org/abs/1807.06735>.
- [26] *Pandora console*. URL: [https://en.wikipedia.org/wiki/Pandora_\(console\)](https://en.wikipedia.org/wiki/Pandora_(console)).
- [27] *Pass the SALT 2018*. URL: <https://2018.pass-the-salt.org/>.
- [28] *Pentium Pro*. URL: https://en.wikipedia.org/wiki/Pentium_Pro.
- [29] *Playing StarCraft on an ARM*. URL: <https://hackaday.com/2014/07/31/playing-starcraft-on-an-arm/>.
- [30] *Politecnico di Milano*. URL: <https://www.polimi.it/>.
- [31] *Polly - Polyhedral optimizations for LLVM*. URL: <https://polly.llvm.org/>.
- [32] *ptrace*. URL: <http://man7.org/linux/man-pages/man2/ptrace.2.html>.
- [33] QEMU. URL: <https://www.qemu.org>.
- [34] RetDec. URL: <https://retdec.com>.
- [35] *RetDec Github Issue 269*. URL: <https://github.com/avast-tl/retdec/issues/269>.
- [36] *RetDec Github Repository*. URL: <https://github.com/avast-tl/retdec>.

- [37] *rev.ng*. URL: <https://rev.ng>.
- [38] *System call*. URL: https://en.wikipedia.org/wiki/System_call.
- [39] *TFA - Transparent Live Code Offloading on FPGA*. URL: <http://reds.heig-vd.ch/en/rad/projects/tfa>.

Appendices

A Authentication

I, Lucas Elisei, hereby declare having realized this work alone and not having used any other resources than those quoted in the bibliography

Par la présente, je soussigné, Lucas Elisei, déclare avoir réalisé seul ce travail et ne pas avoir utilisé d'autres sources que celles citées dans la bibliographie.

Date

Signature

Lucas Elisei

B Requirements

Since we mainly work on ELF binaries, it is highly recommended to use a Linux distribution to build and run the project.

The second step is to download a 32-bit toolchain and the LLVM compiler. You can install them with your preferred package manager. We recommend using the LLVM version 5.

Finally, the most important step, you must download and build RetDec. To do so, clone the RetDec's Github repository [36] and follow the build instructions.

C Simple decompilation example

include/simple.h

```
1  #ifndef LIB_SIMPLE_H
2  #define LIB_SIMPLE_H
3
4  #include <stdint.h>
5
6  /*
7   * Simple addition of two integers.
8   *
9   * Returns the result of the addition.
10  */
11  uint32_t simple_add(uint32_t *a, uint32_t *b);
12
13  /*
14   * Simple addition of two integers. The result is stored at the address of the
15   * third parameter.
16   */
17  void simple_add_ref(uint32_t *a, uint32_t *b, uint32_t *result);
18
19  #endif
```

src/simple.c

```
1  #include <stdint.h>
2
3  #include "simple.h"
4
5  uint32_t simple_add(uint32_t *a, uint32_t *b) {
6      return *a + *b;
7  }
8
9  void simple_add_ref(uint32_t *a, uint32_t *b, uint32_t *result) {
10     *result = *a + *b;
11 }
```

src/main.c

```
1  #include <stdint.h>
2  #include <stdio.h>
3  #include <stdlib.h>
4
5  #include "simple.h"
6
7  int main(int argc, char** argv) {
8      uint32_t a, b;
9
10     a = atoi(argv[1]);
11     b = atoi(argv[2]);
12
13     fprintf(stdout, "Result: %u\n", simple_add(&a, &b));
14
15     return EXIT_SUCCESS;
16 }
```

Makefile

```

1 SHELL := /bin/bash
2
3 CC = gcc
4 override CFLAGS += -std=c99 -Wall -Werror -pedantic -Iinclude -m32
5
6 LDFLAGS = -m32
7
8 CLANG = clang-5.0
9 LLC = llc-5.0
10 LLC_FLAGS = -march=x86
11
12 SRC_DIR = src
13
14 # Logs directory.
15 LOGS_DIR = logs
16 # Rule to create logs.
17 LOGS_RULE = $(shell date "+%Y%m%d-%H%M%S")
18 LOGS_PATH := $(LOGS_DIR)/$(LOGS_RULE)
19
20 # Binary options.
21 BIN = main
22 BIN_SRC = $(wildcard $(SRC_DIR)/*.c)
23 BIN_OBJ = $(patsubst %.c,%.o,$(BIN_SRC))
24 # Arguments to pass to the translated binary.
25 BIN_ARGS = 4 5
26
27 # RetDec options.
28 RETDEC_DIR = $(HOME)/opt/retdec
29 RETDEC_BIN = $(RETDEC_DIR)/bin/retdec-decompiler.sh
30 RETDEC_FLAGS = --stop-after bin2llvmir
31 # Functions to decompile (temporary).
32 RETDEC_FUNCS = simple_add
33 # Only select some functions if asked to.
34 ifdef RETDEC_FUNCS
35 RETDEC_FLAGS += --select-functions $(RETDEC_FUNCS)
36 endif
37
38 .PHONY: all clean decompile recompile
39
40 all: recompile
41
42 decompile: $(BIN)
43     $(RETDEC_BIN) $(RETDEC_FLAGS) $(BIN)
44     @mkdir -p $(LOGS_PATH)
45     @mv -f $(BIN)* $(LOGS_PATH)/
46
47 recompile: decompile
48     @cp $(SRC_DIR)/$(BIN).o $(LOGS_PATH)/$(BIN).o
49     @sed -i '/@_x86.get_pc_thunk.ax()/d' $(LOGS_PATH)/$(BIN).c.backend.ll
50     $(LLC) $(LLC_FLAGS) $(LOGS_PATH)/$(BIN).c.backend.ll -o $(LOGS_PATH)/$(BIN).c.backend.s
51     $(CC) $(CFLAGS) -c $(LOGS_PATH)/$(BIN).c.backend.s -o $(LOGS_PATH)/$(BIN).c.backend.o
52     $(CC) $(LDFLAGS) $(LOGS_PATH)/$(BIN).o $(LOGS_PATH)/$(BIN).c.backend.o -o
53     ↪ $(LOGS_PATH)/$(BIN).translated
54     @echo ---
55     @echo Testing translated binary with parameters: $(BIN_ARGS)
56     @./$(LOGS_PATH)/$(BIN).translated $(BIN_ARGS)
57
58 $(BIN): $(BIN_OBJ)
59     $(CROSS_COMPILE)$(CC) $(CFLAGS) $^ -o $@
60
61 %.o: %.c
62     $(CROSS_COMPILE)$(CC) $(CFLAGS) -c $< -o $@
63
64 clean:
65     rm -rf $(BIN_OBJ)
66     rm -rf $(BIN)*

```

D liboptimizer

D.1 elfparser

elfparser.c

```

1  /*
2   * File: elfparser.c
3   *
4   * Created by: Lucas Elisei <lucas.elisei@heig-ud.ch>
5   */
6
7  #include <elf.h>
8  #include <inttypes.h>
9  #include <stdio.h>
10 #include <stdlib.h>
11 #include <string.h>
12
13 #include <sys/mman.h>
14
15 #include "debug.h"
16 #include "liboptimizer.h"
17 #include "elfparser.h"
18
19 typedef struct {
20     uint8_t    *mem;        // Binary mapped into memory
21     long        file_size;  // Binary size.
22     uint8_t     is64;       // 32- or 64-bit
23     uint64_t    sh_off;     // Points to the start of the section header table
24     uint16_t    sh_num;     // Number of entries in the section header table
25     uint16_t    sh_entsize; // Section header size
26     uint16_t    sh_strndx;  // Index of SH String table.
27     uint64_t    sym_off;    // Symbol table offset
28     uint64_t    sym_size;   // Symbol table size
29     uint64_t    str_off;    // String table size
30 } elf_file_t;
31
32 static void _elf_dump_file(elf_file_t *elf) {
33     DBG("=== Printing elf_file_t at 0x%" PRIXPTR "\n", (uintptr_t)elf);
34     DBG(" is64.....: %" PRIu8 " \n", elf->is64);
35     DBG(" sh_off....: 0x%" PRIx64 " \n", elf->sh_off);
36     DBG(" sh_num....: %" PRIu16 " \n", elf->sh_num);
37     DBG(" sh_entsize: %" PRIu16 " \n", elf->sh_entsize);
38     DBG(" sh_strndx.: %" PRIu16 " \n", elf->sh_strndx);
39     DBG(" sym_off....: 0x%" PRIx64 " \n", elf->sym_off);
40     DBG(" sym_size...: %" PRIu64 " \n", elf->sym_size);
41     DBG(" str_off....: 0x%" PRIx64 " \n", elf->str_off);
42     DBG("=====\n\n");
43 }
44
45 static char *_elf_resolve_symbol(elf_file_t *elf, uint64_t address) {
46     int i;
47     size_t offset;
48     // Symbol header for 32- and 64-bit.
49     Elf32_Sym sym32;
50     Elf64_Sym sym64;
51     char *symbol;
52     size_t sym_len;
53
54     // ELF32
55     if (elf->is64 == 0) {
56         // Iterate over all symbols.
57         for (i = 0; i * sizeof(Elf32_Sym) < elf->sym_size; ++i) {
58             // Calculate offset and get symbol.
59             offset = elf->sym_off + (sizeof(Elf32_Sym) * i);
60             memmove(&sym32, elf->mem + offset, sizeof(Elf32_Sym));
61
62             // If the symbol is at the address we are looking for.
63             if (sym32.st_value == address) {
64                 // Sanity check.
65                 if (sym32.st_name == 0) {

```

```
66         return NULL;
67     }
68
69     // Get symbol size for memory allocation.
70     offset = elf->str_off + sym32.st_name;
71     sym_len = strlen((char *) (elf->mem + offset));
72
73     // Allocate memory for symbol.
74     symbol = (char *) calloc(1, sizeof(char) * (sym_len + 1));
75     if (symbol == NULL) {
76         return NULL;
77     }
78
79     // Copy symbol.
80     sprintf(symbol, "%s", (char *) (elf->mem + offset));
81 }
82 }
83 }
84 // ELF64
85 else {
86     // Iterate over all symbols.
87     for (i = 0; i * sizeof(Elf64_Sym) < elf->sym_size; ++i) {
88         // Calculate offset and get symbol.
89         offset = elf->sym_off + (sizeof(Elf64_Sym) * i);
90         memmove(&sym64, elf->mem + offset, sizeof(Elf64_Sym));
91
92         if (sym64.st_value == address) {
93             // Sanity check.
94             if (sym64.st_name == 0) {
95                 return NULL;
96             }
97
98             // Get symbol size for memory allocation.
99             offset = elf->str_off + sym64.st_name;
100             sym_len = strlen((char *) (elf->mem + offset));
101
102             // Allocate memory for symbol.
103             symbol = (char *) calloc(1, sizeof(char) * (sym_len + 1));
104             if (symbol == NULL) {
105                 return NULL;
106             }
107
108             // Copy symbol.
109             sprintf(symbol, "%s", (char *) (elf->mem + offset));
110         }
111     }
112 }
113
114     return symbol;
115 }
116
117 static uint64_t _elf_resolve_address(elf_file_t *elf, const char *symbol) {
118     int i;
119     size_t offset;
120     // Symbol header for 32- and 64-bit.
121     Elf32_Sym sym32;
122     Elf64_Sym sym64;
123
124     // ELF32
125     if (elf->is64 == 0) {
126         // Iterate over all symbols.
127         for (i = 0; i * sizeof(Elf32_Sym) < elf->sym_size; ++i) {
128             // Calculate offset and get symbol struct.
129             offset = elf->sym_off + (sizeof(Elf32_Sym) * i);
130             memmove(&sym32, elf->mem + offset, sizeof(Elf32_Sym));
131
132             // Sanity check.
133             if (sym32.st_name == 0) {
134                 continue;
135             }
136
137             // Calculate symbol offset.
138             offset = elf->str_off + sym32.st_name;
```



```

139
140     // If the symbol is equal to the one we are looking for, return
141     // its value.
142     if (!strcmp((char *) (elf->mem + offset), symbol)) {
143         return (uint64_t) sym32.st_value;
144     }
145 }
146 }
147 // ELF64
148 else {
149     // Iterate over all symbols.
150     for (i = 0; i * sizeof(Elf64_Sym) < elf->sym_size; ++i) {
151         // Calculate offset and get symbol.
152         offset = elf->sym_off + (sizeof(Elf64_Sym) * i);
153         memmove(&sym64, elf->mem + offset, sizeof(Elf64_Sym));
154
155         // Sanity check.
156         if (sym64.st_name == 0) {
157             continue;
158         }
159
160         // Calculate symbol offset.
161         offset = elf->str_off + sym64.st_name;
162
163         // If the symbol is equal to the one we are looking for, return
164         // its value.
165         if (!strcmp((char *) (elf->mem + offset), symbol)) {
166             return (uint64_t) sym64.st_value;
167         }
168     }
169 }
170
171 return 0;
172 }
173
174 static int _elf_resolve_sections(elf_file_t *elf) {
175     int i;
176     size_t offset;
177     uint64_t shstrtab_off;
178     // Section header for 32- and 64-bit.
179     Elf32_Shdr sec32;
180     Elf64_Shdr sec64;
181
182     // ELF32
183     if (elf->is64 == 0) {
184         // We need to get the Section Header STRING TABLE offset before others.
185         offset = elf->sh_off + (elf->sh_entsize * elf->sh_strndx);
186         memmove(&sec32, elf->mem + offset, sizeof(Elf32_Shdr));
187         shstrtab_off = sec32.sh_offset;
188
189         // Iterate over all section headers.
190         for (i = 0; i < elf->sh_num; ++i) {
191             // Calculate offset and get section header.
192             offset = elf->sh_off + (elf->sh_entsize * i);
193             memmove(&sec32, elf->mem + offset, sizeof(Elf32_Shdr));
194
195             switch (sec32.sh_type) {
196                 // Static symbols table.
197                 case SHT_SYMTAB:
198                     elf->sym_off = sec32.sh_offset;
199                     elf->sym_size = sec32.sh_size;
200                     break;
201
202                 // String table. Since there are more than one string table, we
203                 // have to be sure to get the .strtab one.
204                 case SHT_STRTAB:
205                     if (!strcmp((char *) (elf->mem + shstrtab_off + sec32.sh_name), ".strtab")) {
206                         elf->str_off = sec32.sh_offset;
207                     }
208                     break;
209
210                 default:
211                     break;

```

```

212     }
213 }
214 }
215 // ELF64
216 else {
217     // We need to get the Section Header STRing TABle offset before others.
218     offset = elf->sh_off + (elf->sh_entsize * elf->sh_strndx);
219     memmove(&sec64, elf->mem + offset, sizeof(Elf64_Shdr));
220     shstrtab_off = sec64.sh_offset;
221
222     // Iterate over all section headers.
223     for (i = 0; i < elf->sh_num; ++i) {
224         // Calculate offset and get section header.
225         offset = elf->sh_off + (elf->sh_entsize * i);
226         memmove(&sec64, elf->mem + offset, sizeof(Elf64_Shdr));
227
228         switch (sec64.sh_type) {
229             // Static symbols table.
230             case SHT_SYMTAB:
231                 elf->sym_off = sec64.sh_offset;
232                 elf->sym_size = sec64.sh_size;
233                 break;
234
235             // String table.
236             case SHT_STRTAB:
237                 if (!strcmp((char *) (elf->mem + shstrtab_off + sec64.sh_name), ".strtab")) {
238                     elf->str_off = sec64.sh_offset;
239                 }
240                 break;
241
242             default:
243                 break;
244         }
245     }
246 }
247
248 return 0;
249 }
250
251 static int _elf_read_header(elf_file_t *elf) {
252     // The first bytes of the header are same-sized for 32- and 64-bit archs.
253     // To identify the file's class and magic number, we assume it's 32-bit.
254     Elf32_Ehdr hdr32;
255     Elf64_Ehdr hdr64;
256     int rc;
257
258     // Retrieve ELF header.
259     memmove(&hdr32, elf->mem, sizeof(Elf32_Ehdr));
260
261     // Check that the file is a valid ELF.
262     rc = (hdr32.e_ident[EI_MAG0] == ELFMAG0 &&
263          hdr32.e_ident[EI_MAG1] == ELFMAG1 &&
264          hdr32.e_ident[EI_MAG2] == ELFMAG2 &&
265          hdr32.e_ident[EI_MAG3] == ELFMAG3);
266     if (rc == 0) {
267         fprintf(stderr, "[liboptimizer] ERROR: File is not a valid ELF\n");
268
269         rc = -2;
270         goto _elf_not_valid;
271     }
272
273     // Check ELF class.
274     switch (hdr32.e_ident[EI_CLASS]) {
275         case ELFCLASS32:
276             elf->is64 = 0;
277             elf->sh_off = (uint64_t)hdr32.e_shoff;
278             elf->sh_num = hdr32.e_shnum;
279             elf->sh_entsize = hdr32.e_shentsize;
280             elf->sh_strndx = hdr32.e_shstrndx;
281             break;
282         case ELFCLASS64:
283             memmove(&hdr64, elf->mem, sizeof(Elf64_Ehdr));
284             elf->is64 = 1;

```

```
285         elf->sh_off = hdr64.e_shoff;
286         elf->sh_num = hdr64.e_shnum;
287         elf->sh_entsize = hdr64.e_shentsize;
288         elf->sh_strndx = hdr64.e_shstrndx;
289         break;
290     default:
291         fprintf(stderr, "[liboptimizer] ERROR: Invalid ELF class\n");
292         rc = -2;
293         goto _elf_not_valid;
294     }
295
296     rc = 0;
297
298 _elf_not_valid:
299     return rc;
300 }
301
302 static elf_file_t *_elf_init(const char *path) {
303     elf_file_t *elf;
304     FILE *file;
305
306     // Allocate memory.
307     elf = (elf_file_t *)calloc(1, sizeof(elf_file_t));
308     if (elf == NULL) {
309         perror("[liboptimizer] calloc()");
310
311         elf = NULL;
312         goto _failed_calloc;
313     }
314
315     // Open binary.
316     file = fopen(path, "rb");
317     if (file == NULL) {
318         perror("[liboptimizer] fopen()");
319
320         elf = NULL;
321         goto _failed_fopen;
322     }
323
324     // Get size.
325     fseek(file, 0L, SEEK_END);
326     elf->file_size = ftell(file);
327
328     // Map binary into memory.
329     elf->mem = mmap(NULL, elf->file_size, PROT_READ, MAP_PRIVATE, fileno(file), 0);
330     if (elf->mem == NULL) {
331         perror("[liboptimizer] mmap()");
332
333         elf = NULL;
334         goto _failed_mmap;
335     }
336
337     fclose(file);
338
339     return elf;
340
341 _failed_mmap:
342     fclose(file);
343 _failed_fopen:
344     free(elf);
345 _failed_calloc:
346     return NULL;
347 }
348
349 static elf_file_t *_parse_elf(const char *path) {
350     elf_file_t *elf;
351     int rc;
352
353     elf = _elf_init(path);
354     if (elf == NULL) {
355         fprintf(stderr, "[liboptimizer] ERROR: Failed to allocate memory\n");
356
357         goto _failed_init;
```

```
358     }
359
360     rc = _elf_read_header(elf);
361     if (rc < 0) {
362         fprintf(stderr, "[liboptimizer] ERROR: Error while parsing ELF\n");
363
364         goto _failed_read_header;
365     }
366
367     rc = _elf_resolve_sections(elf);
368     if (rc < 0) {
369         fprintf(stderr, "[liboptimizer] ERROR: Error while resolving sections\n");
370
371         goto _failed_resolve_sections;
372     }
373
374     return elf;
375
376 _failed_resolve_sections:
377     munmap(elf->mem, elf->file_size);
378 _failed_read_header:
379     free(elf);
380 _failed_init:
381     return NULL;
382 }
383
384 char *get_symbol_at_address(const char *path, uint64_t address) {
385     elf_file_t *elf;
386     char *symbol;
387
388     elf = _parse_elf(path);
389     if (elf == NULL) {
390         symbol = NULL;
391
392         goto _failed_parse;
393     }
394
395     _elf_dump_file(elf);
396
397     symbol = _elf_resolve_symbol(elf, address);
398     if (symbol == NULL) {
399         fprintf(stderr, "[liboptimizer] ERROR: Error while retrieving symbols info\n");
400
401         goto _failed_sym_info;
402     }
403
404 _failed_sym_info:
405     munmap(elf->mem, elf->file_size);
406     free(elf);
407 _failed_parse:
408     return symbol;
409 }
410
411 uint64_t get_symbol_address(const char *path, const char *symbol) {
412     elf_file_t *elf;
413     uint64_t address;
414
415     elf = _parse_elf(path);
416     if (elf == NULL) {
417         address = 0;
418
419         goto _failed_parse;
420     }
421
422     _elf_dump_file(elf);
423
424     address = _elf_resolve_address(elf, symbol);
425     if (address == 0) {
426         fprintf(stderr, "[liboptimizer] ERROR: Error while retrieving address\n");
427
428         goto _failed_resolve_addr;
429     }
430 }
```

```
431     _failed_resolve_addr:
432         munmap(elf->mem, elf->file_size);
433         free(elf);
434     _failed_parse:
435         return address;
436 }
437
438 int8_t is64bit(const char *path) {
439     elf_file_t *elf;
440     int8_t is64;
441
442     elf = _parse_elf(path);
443     if (elf == NULL) {
444         fprintf(stderr, "[liboptimizer] ERROR: Failed to allocate memory\n");
445
446         return -1;
447     }
448
449     is64 = (int8_t)elf->is64;
450
451     munmap(elf->mem, elf->file_size);
452     free(elf);
453
454     return is64;
455 }
```

D.2 retdec

retdec.c

```
1  /*
2   * File: retdec.c
3   *
4   * Created by: Lucas Elisei <lucas.elisei@heig-vd.ch>
5   */
6
7  #define _POSIX_C_SOURCE 200809L
8  #define _XOPEN_SOURCE 500
9
10 #include <fcntl.h>
11 #include <ftw.h>
12 #include <libgen.h>
13 #include <limits.h>
14 #include <stdlib.h>
15 #include <string.h>
16 #include <unistd.h>
17
18 #include <sys/stat.h>
19 #include <sys/types.h>
20 #include <sys/wait.h>
21
22 #include "debug.h"
23 #include "liboptimizer.h"
24 #include "retdec.h"
25 #include "TFAJITWrapper.h"
26
27 int _remove_file(const char *fpath, const struct stat *sb, int typeflag, struct FTW *ftwbuf) {
28     int rv = remove(fpath);
29
30     if (rv) {
31         perror((char*)fpath);
32     }
33
34     return rv;
35 }
36
37 static int _compile_llvmir(process_info_t *process, char *tmp_directory) {
38     int rc;
39
40     rc = compile_llvmir_file(process, tmp_directory);
41
42     return rc;
43 }
44
45 static int _execute_script(process_info_t *process, char *tmp_directory) {
46     int log_fd;
47     char log_path[PATH_MAX];
48     char tmp_binary[PATH_MAX];
49     pid_t pid;
50
51     sprintf(log_path, "%s/decompile.log", tmp_directory);
52     sprintf(tmp_binary, "%s/%s", tmp_directory, basename((char *)process->path));
53
54     log_fd = open(log_path, O_RDWR | O_CREAT, 0666);
55     if (log_fd < 0) {
56         perror("[liboptimizer] open()");
57
58         return -1;
59     }
60
61     pid = fork();
62     if (pid == 0) {
63         dup2(log_fd, STDOUT_FILENO);
64         dup2(log_fd, STDERR_FILENO);
65         close(log_fd);
66
67         execl("/bin/bash", "sh", RETDEC_DECOMPILER,
68             "--stop-after", "bin2llvmir",
```

```
69         "--select-functions", process->function_name,
70         "--output", tmp_binary,
71         process->path, NULL);
72
73     exit(0);
74 }
75
76 waitpid(pid, NULL, 0);
77
78 return 0;
79 }
80
81 int retdec_recompile(process_info_t *process) {
82     // Store return codes.
83     int rc;
84     // Temporary directory name.
85     char *tmp_dir;
86     // Buffer to store temporary directory template.
87     char template_buffer[strlen(TMP_DIR_TEMPLATE) + 1];
88
89     sprintf(template_buffer, TMP_DIR_TEMPLATE);
90
91     // Create a temporary directory to store temporary files.
92     tmp_dir = mkdtemp(template_buffer);
93     if (tmp_dir == NULL) {
94         perror("[liboptimizer] mkdtemp");
95
96         return -1;
97     }
98
99     DBG("Created %s temporary directory\n", tmp_dir);
100
101     // Execute RetDec's script into temporary directory.
102     rc = _execute_script(process, tmp_dir);
103     if (rc != 0) {
104         fprintf(stderr, "[liboptimizer] ERROR: An error occurred while executing decompilation
105             ↪ script\n");
106
107         goto _delete_tmp_dir;
108     }
109
110     // Call JIT compiler.
111     rc = _compile_llvmir(process, tmp_dir);
112     if (rc != 0) {
113         fprintf(stderr, "[liboptimizer] ERROR: An error occurred during compilation of the optimized
114             ↪ function\n");
115
116         goto _delete_tmp_dir;
117     }
118
119     rc = 0;
120
121     _delete_tmp_dir:
122     #ifndef LIBOPTIMIZER_DEBUG
123     if (nftw(tmp_dir, _remove_file, 64, FTW_DEPTH | FTW_PHYS)) {
124         perror("[liboptimizer] nftw()");
125     }
126     #endif
127
128     return rc;
129 }
```

D.3 jit

TFAJITEventListener.hpp

```
1  /*
2   * File: TFAJITEventListener.hpp
3   *
4   * Created by: Lucas Elisei <lucas.elisei@heig-vd.ch>
5   */
6
7  #ifndef __LIBOPTIMIZER_INCLUDE_TFAJITEVENTLISTENER_H__
8  #define __LIBOPTIMIZER_INCLUDE_TFAJITEVENTLISTENER_H__
9
10 #include "llvm/ExecutionEngine/JITEventListener.h"
11 #include "llvm/Object/SymbolSize.h"
12
13 #include <iostream>
14 #include <map>
15
16 using namespace llvm;
17 using namespace llvm::object;
18
19 typedef struct {
20     uint64_t address;
21     uint64_t size;
22 } symbol_info_t;
23
24 class TFAJITEventListener : public JITEventListener {
25
26 private:
27     // Map used to store symbols information.
28     std::map<std::string, symbol_info_t *> symbolsMap;
29
30 public:
31     // Default constructor.
32     TFAJITEventListener() {}
33     // Default destructor.
34     ~TFAJITEventListener() {
35         for (auto it = symbolsMap.begin(); it != symbolsMap.end(); ++it) {
36             free(it->second);
37         }
38     }
39
40     // Function called when the JIT has emitted an object file.
41     virtual void NotifyObjectEmitted(const ObjectFile &obj, const RuntimeDyld::LoadedObjectInfo &L) {
42         OwningBinary<ObjectFile> OWOF = L.GetObjectForDebug(obj);
43         ObjectFile &OF = *OWOF.getBinary();
44
45         // Iterate over symbols and their respective size.
46         for (const std::pair<SymbolRef, uint64_t> &pair : computeSymbolSizes(OF)) {
47             SymbolRef symbolRef = std::get<0>(pair);
48             uint64_t size = std::get<1>(pair);
49
50             // Symbol with an empty size aren't interesting.
51             if (size > 0) {
52                 symbol_info_t *symbol_info;
53
54                 // Allocate memory for symbol_info.
55                 symbol_info = (symbol_info_t *)calloc(1, sizeof(symbol_info_t));
56                 if (symbol_info == NULL) {
57                     std::cerr << "ERROR: Could not allocate memory" << std::endl;
58
59                     continue;
60                 }
61
62                 // Retrieve address of symbol.
63                 Expected<uint64_t> eAddr = symbolRef.getAddress();
64                 if (!eAddr) {
65                     continue;
66                 }
67
68                 // Assign fields.
```



```
69         symbol_info->size = size;
70         symbol_info->address = *eAddr;
71
72         // Insert symbol info into map.
73         symbolsMap[symbolRef.getName().get().str()] = symbol_info;
74     }
75 }
76 }
77
78 // Returns a pointer on a symbol_info struct corresponding to the argument.
79 // Returns NULL if the symbol could not be found.
80 symbol_info_t *GetSymbolInfo(const std::string symbol) {
81     return symbolsMap[symbol];
82 }
83 };
84
85 #endif
```

TFAJITWrapper.cpp

```
1  /*
2   * File: TFAJITWrapper.cpp
3   *
4   * Created by: Lucas Elisei <lucas.elisei@heig-vd.ch>
5   */
6
7  #include "llvm/ExecutionEngine/ExecutionEngine.h"
8  #include "llvm/IR/Module.h"
9  #include "llvm/IR/Verifier.h"
10 #include "llvm/IRReader/IRReader.h"
11 #include "llvm/Support/CodeGen.h"
12 #include "llvm/Support/SourceMgr.h"
13
14 #include <climits>
15 #include <cstdlib>
16 #include <iostream>
17
18 #include "debug.h"
19 #include "liboptimizer.h"
20 #include "TFAJITEventListener.hpp"
21
22 using namespace llvm;
23
24 extern "C" {
25
26 int compile_llvmir_file(process_info_t *process, char *tmp_path) {
27     SMDiagnostic err;
28     static LLVMContext context;
29     char path[PATH_MAX];
30     char command[PATH_MAX];
31
32     // Initialize targets.
33     LLVMInitializeAllTargets();
34     LLVMInitializeAllTargetMCs();
35     LLVMInitializeAllTargetInfos();
36     LLVMInitializeAllAsmPrinters();
37     LLVMInitializeAllAsmParsers();
38     LLVMInitializeAllDisassemblers();
39
40     sprintf(path, "%s/%s.backend.ll", tmp_path, basename(process->path));
41
42     // Remove occurrences of __x86.get_pc_thunk.ax() function into LLVM IR file.
43     sprintf(command, "sed -i '/@__x86.get_pc_thunk.ax/d' %s", path);
44     system(command);
45
46     // Parse LLVM IR file.
47     std::unique_ptr<Module> module = parseIRFile(path, err, context);
48
49     if (!module) {
50         err.print(process->argv[0], llvm::errs());
```

```

51     return -1;
52 }
53
54 // Verify that the module is valid.
55 if (verifyModule(*module)) {
56     std::cerr << "ERROR: The LLVM IR module is not valid" << std::endl;
57
58     return -2;
59 }
60
61 // Initialize ExecutionEngine as a JIT Compiler.
62 StringRef *arch;
63 if (process->is64 == 0) {
64     arch = new StringRef("x86");
65 } else {
66     arch = new StringRef("x86-64");
67 }
68
69 ExecutionEngine *EE = EngineBuilder(std::move(module))
70     .setEngineKind(EngineKind::JIT)
71     .setMARCH(*arch)
72     .setOptLevel(CoGenOpt::Level::Aggressive)
73     .setVerifyModules(true)
74     .create();
75
76 // Initialize JIT Event Listener.
77 TFAJITEventListener *EL = new TFAJITEventListener();
78
79 // Register JITEventListener.
80 EE->RegisterJITEventListener(EL);
81
82 // Compile module.
83 EE->finalizeObject();
84
85 DBG("=== ExecutionEngine %s\n", "dump");
86 DBG("  triple: %s\n", EE->getTargetMachine()->getTargetTriple().str().c_str());
87 DBG("  cpu...: %s\n", EE->getTargetMachine()->getTargetCPU().str().c_str());
88 DBG("  layout: %s\n", EE->getDataLayout().getStringRepresentation().c_str());
89
90 // Retrieve symbol info.
91 std::string symbol_str(process->function_name);
92 symbol_info_t *symbol_info = EL->GetSymbolInfo(symbol_str);
93 if (symbol_info == NULL) {
94     std::cerr << "ERROR: Could not retrieve recompiled symbol info" << std::endl;
95
96     return -3;
97 }
98
99 symbol_info->address = EE->getFunctionAddress(process->function_name);
100
101 // Assign fields.
102 process->optimized_function_size = symbol_info->size;
103 process->optimized_function = (uint8_t *)calloc(process->optimized_function_size, sizeof(uint8_t));
104 if (process->optimized_function == NULL) {
105     std::cerr << "ERROR: Could not allocate memory for optimized function" << std::endl;
106 }
107 memcpy(process->optimized_function, (void *)symbol_info->address,
108     ↪ process->optimized_function_size);
109
110 return 0;
111 }
112 // extern "C"

```

D.4 live-patcher

live-patcher.c

```
1  /*
2   * File: ptrace.c
3   *
4   * Created by: Lucas Elisei <lucas.elisei@heig-vd.ch>
5   */
6
7  #define _GNU_SOURCE
8
9  #include <fcntl.h>
10 #include <inttypes.h>
11 #include <stdio.h>
12 #include <stdlib.h>
13 #include <string.h>
14 #include <unistd.h>
15
16 #include <sys/mman.h>
17 #include <sys/ptrace.h>
18 #include <sys/types.h>
19 #include <sys/uio.h>
20 #include <sys/user.h>
21 #include <sys/wait.h>
22
23 #include "debug.h"
24 #include "liboptimizer.h"
25
26 #define MMAP2_SYSCALL_X32 192
27 #define MMAP2_SYSCALL_X64 9
28
29 /*
30  * Creates an assembly unconditional jump for 32-bit binaries.
31  *
32  * MOV _A_, %eax
33  * JMP *%eax
34  */
35 #define MAKE_JUMP32(_A_) { \
36     0xB8, \
37     ((unsigned char *)&(_A_))[0], \
38     ((unsigned char *)&(_A_))[1], \
39     ((unsigned char *)&(_A_))[2], \
40     ((unsigned char *)&(_A_))[3], \
41     0xFF, 0xE0 \
42 }
43
44 /*
45  * Creates an assembly unconditional jump for 64-bit binaries.
46  *
47  * MOVABS _A_, %rax
48  * JMP *%rax
49  */
50 #define MAKE_JUMP64(_A_) { \
51     0x48, 0xB8, \
52     ((unsigned char *)&(_A_))[0], \
53     ((unsigned char *)&(_A_))[1], \
54     ((unsigned char *)&(_A_))[2], \
55     ((unsigned char *)&(_A_))[3], \
56     ((unsigned char *)&(_A_))[4], \
57     ((unsigned char *)&(_A_))[5], \
58     ((unsigned char *)&(_A_))[6], \
59     ((unsigned char *)&(_A_))[7], \
60     0xFF, 0xE0 \
61 }
62
63 /*
64  * Methods used for debugging purposes.
65  */
66 #ifdef LIBOPTIMIZER_DEBUG
67 void print_regs(process_info_t *process) {
68     struct user_regs_struct regs;
```

```

69     ptrace(PTRACE_GETREGS, process->pid, NULL, &regs);
70
71     DBG("Registers:\n");
72     DBG("  ebp: 0x%" PRIx64 " (0x%11X)\n", ptrace(PTRACE_PEEKDATA, process->pid, regs.rbp, NULL),
73         ↪ regs.rbp);
74     DBG("  esp: 0x%" PRIx64 " (0x%11X)\n", ptrace(PTRACE_PEEKDATA, process->pid, regs.rsp, NULL),
75         ↪ regs.rsp);
76 }
77
78 void print_stack(process_info_t *process) {
79     long addr, value;
80     int size;
81     struct user_regs_struct regs;
82
83     ptrace(PTRACE_GETREGS, process->pid, NULL, &regs);
84
85     fprintf(stdout, "Stack:");
86
87     if (process->is64 == 0) {
88         size = 4;
89     } else {
90         size = 8;
91     }
92
93     for (addr = regs.rbp; addr > ((regs.rsp - 0x30) & ~0xFFu); addr -= size) {
94         if (((addr ^ regs.rbp) & 0xFFu) == 0) {
95             fprintf(stdout, "\n");
96             DBG("0x%0*1X: ", size * 2, addr);
97         }
98         value = ptrace(PTRACE_PEEKDATA, process->pid, addr, NULL);
99         if (addr == regs.rsp) {
100             fprintf(stdout, "*");
101         }
102         fprintf(stdout, "0x%0*1X ", size * 2, value);
103     }
104     fprintf(stdout, "\n");
105 }
106
107 void debug(process_info_t *process) {
108     struct user_regs_struct regs;
109     char c;
110     uint64_t address;
111
112     address = process->codesegment_address + process->function_offset;
113
114     ptrace(PTRACE_GETREGS, process->pid, NULL, &regs);
115     while (regs.rip != address) {
116         ptrace(PTRACE_SINGLESTEP, process->pid, NULL, NULL);
117         waitpid(process->pid, NULL, 0);
118         ptrace(PTRACE_GETREGS, process->pid, NULL, &regs);
119     }
120
121     DBG("=== DEBUG COMMANDS ===\n");
122     DBG("  p : print EBP and ESP registers\n");
123     DBG("  s : print stack\n");
124     DBG("  n : execute next instruction\n");
125     DBG("  q : continue execution\n");
126     DBG("=====\n");
127
128     DBG("Got at 0x%" PRIx64 "\n", address);
129
130     do {
131         ptrace(PTRACE_GETREGS, process->pid, NULL, &regs);
132
133         DBG("=== 0x%0811X (eip)\n", regs.rip);
134         DBG("> ");
135         c = getchar();
136         while(getchar() != '\n');
137
138         switch (c) {
139             case 'p':
140                 print_regs(process);

```

```
140         break;
141
142         case 's':
143             print_stack(process);
144             break;
145
146         case 'n':
147             ptrace(PTRACE_SINGLESTEP, process->pid, NULL, NULL);
148             waitpid(process->pid, NULL, 0);
149             break;
150
151         default:
152             break;
153     }
154     } while (c != 'q');
155 }
156 #endif
157
158 /*
159  * Waits for a syscall.
160  *
161  * Returns '0' if the child stopped on a syscall. Returns '1' if the child
162  * exited.
163  */
164 static int _wait_for_syscall(pid_t pid) {
165     int status;
166
167     while (1) {
168         // Tell ptrace to wait for a syscall.
169         ptrace(PTRACE_SYSCALL, pid, 0, 0);
170
171         // Wait for child to be stopped.
172         waitpid(pid, &status, 0);
173
174         // Check status.
175         if (WIFSTOPPED(status) && WSTOPSIG(status) & 0x80) {
176             return 0;
177         }
178
179         if (WIFEXITED(status)) {
180             return 1;
181         }
182     }
183 }
184
185 /*
186  * Update the address of the code segment of the given process in its struct.
187  *
188  * On success, returns 0. Otherwise, prints an error and returns '-1'.
189  */
190 static int _get_codesegment_address(process_info_t *process) {
191     int rc;
192     char file_path[255];
193     char seg_addr_str[17];
194     long seg_addr;
195     FILE *file;
196
197     // Get path of the `maps` file for the given pid.
198     snprintf(file_path, sizeof(file_path), "/proc/%d/maps", process->pid);
199
200     file = fopen(file_path, "r");
201     if (file == NULL) {
202         perror("[liboptimizer] fopen()");
203
204         return -1;
205     }
206
207     // For now, we only retrieve the first 16 characters. Might need to check
208     // permissions and all to be sure we parse the code segment.
209     rc = fread(seg_addr_str, 1, sizeof(seg_addr_str) - 1, file);
210     if (rc < 0) {
211         perror("[liboptimizer] fread()");
212     }
```

```
213     return -1;
214 }
215 // Append end-of-string character.
216 seg_addr_str[sizeof(seg_addr_str) - 1] = 0;
217
218 fclose(file);
219
220 // Convert string to long.
221 seg_addr = strtol(seg_addr_str, NULL, 16);
222
223 // Assign value.
224 process->codesegment_address = seg_addr;
225
226 return 0;
227 }
228
229 /*
230  * Replaces the memory of the child process to hook the `func_addr` function
231  * with the `free_addr` function.
232  *
233  * Returns '0' on success. Otherwise, prints an error and returns '-1'.
234  */
235 static int _replace_mem(process_info_t *process) {
236     unsigned char jump_32[] = MAKE_JUMP32(process->freeselement_address);
237     unsigned char jump_64[] = MAKE_JUMP64(process->freeselement_address);
238
239     int i, rc;
240     uint64_t old_func_addr = process->function_offset + process->codesegment_address;
241
242     // Put the new function into free allocated space.
243     for (i = 0; i < process->optimized_function_size; i += 4) {
244         rc = ptrace(PTRACE_POKEDATA,
245                     process->pid,
246                     (void *)process->freeselement_address + i,
247                     ((unsigned int *)process->optimized_function)[i / 4]);
248         if (rc < 0) {
249             perror("[liboptimizer] PTRACE_POKEDATA");
250
251             return -1;
252         }
253     }
254
255     // 32-bit
256     if (process->is64 == 0) {
257         for (i = 0; i < sizeof(jump_32); i += 4) {
258             rc = ptrace(PTRACE_POKEDATA,
259                         process->pid,
260                         (void *)old_func_addr + i, ((unsigned int *)jump_32)[i / 4]);
261             if (rc < 0) {
262                 perror("[liboptimizer] PTRACE_POKEDATA");
263
264                 return -1;
265             }
266         }
267     }
268     // 64-bit
269     else {
270         for (i = 0; i < sizeof(jump_64); i += 4) {
271             rc = ptrace(PTRACE_POKEDATA,
272                         process->pid,
273                         (void *)old_func_addr + i, ((unsigned int *)jump_64)[i / 4]);
274             if (rc < 0) {
275                 perror("[liboptimizer] PTRACE_POKEDATA");
276
277                 return -1;
278             }
279         }
280     }
281
282     return 0;
283 }
284
285 /*
```

```
286  * Injects a 'mmap2' syscall into a child process.
287  *
288  * On success, returns the address of the newly allocated memory segment.
289  * If the child process unexpectedly stopped, returns '-1'.
290  */
291  static int _inject_mmap(process_info_t *process) {
292      struct user_regs_struct old_regs, new_regs;
293      int rc;
294
295      if (_wait_for_syscall(process->pid) != 0) {
296          return -1;
297      }
298
299      rc = ptrace(PTRACE_GETREGS, process->pid, 0, &old_regs);
300      if (rc < 0) {
301          perror("[liboptimizer] PTRACE_GETREGS");
302
303          return -1;
304      }
305
306      memcpy(&new_regs, &old_regs, sizeof(struct user_regs_struct));
307
308      // 32-bit.
309      if (process->is64 == 0) {
310          // Registers for 32-bit binary on 64-bit machine.
311          new_regs.rax = MMAP2_SYSCALL_X32;
312          new_regs.rbx = 0;
313          new_regs.rcx = process->optimized_function_size;
314          new_regs.rdx = PROT_READ | PROT_WRITE | PROT_EXEC;
315          new_regs.rsi = MAP_PRIVATE | MAP_ANONYMOUS;
316          new_regs.rdi = -1;
317          new_regs.rbp = 0;
318          new_regs.orig_rax = MMAP2_SYSCALL_X32;
319      }
320      // 64-bit.
321      else {
322          // Registers for 64-bit binary on 64-bit machine.
323          new_regs.rax = MMAP2_SYSCALL_X64;
324          new_regs.rdi = 0;
325          new_regs.rsi = process->optimized_function_size;
326          new_regs.rdx = PROT_READ | PROT_WRITE | PROT_EXEC;
327          new_regs.r10 = MAP_PRIVATE | MAP_ANONYMOUS;
328          new_regs.r8 = -1;
329          new_regs.r9 = 0;
330          new_regs.orig_rax = MMAP2_SYSCALL_X64;
331      }
332
333      rc = ptrace(PTRACE_SETREGS, process->pid, NULL, &new_regs);
334      if (rc < 0) {
335          perror("[liboptimizer] PTRACE_SETREGS");
336
337          return -1;
338      }
339
340      rc = ptrace(PTRACE_SINGLESTEP, process->pid, NULL, NULL);
341      if (rc < 0) {
342          perror("[liboptimizer] PTRACE_SINGLESTEP");
343
344          return -1;
345      }
346
347      waitpid(process->pid, NULL, 0);
348
349      rc = ptrace(PTRACE_GETREGS, process->pid, NULL, &new_regs);
350      if (rc < 0) {
351          perror("[liboptimizer] PTRACE_GETREGS");
352
353          return -1;
354      }
355
356      rc = ptrace(PTRACE_SETREGS, process->pid, NULL, &old_regs);
357      if (rc < 0) {
358          perror("[liboptimizer] PTRACE_SETREGS");
```

```
359         return -1;
360     }
361
362     process->freeselement_address = new_regs.rax;
363
364     return 0;
365 }
366
367 /*
368  * Main function of parent process.
369  *
370  * On success, returns 0. Otherwise, prints an error message and returns the
371  * corresponding value.
372  */
373
374 static int _do_parent(process_info_t *process) {
375     int rc;
376
377     // Wait for child to be stopped.
378     waitpid(process->pid, NULL, 0);
379
380     // Set options for ptrace.
381     ptrace(PTRACE_SETOPTIONS, process->pid, 0, PTRACE_O_TRACESYSGOOD);
382
383     // Get address of the function to replace in the code segment.
384     rc = _get_codesegment_address(process);
385     if (rc < 0) {
386         fprintf(stderr, "[liboptimizer] ERROR: Failed to get code segment address\n");
387
388         return -1;
389     }
390
391     // Allocate new memory segment.
392     rc = _inject_mmap(process);
393     if (rc < 0) {
394         fprintf(stderr, "[liboptimizer] ERROR: Failed to allocate memory segment\n");
395
396         return -2;
397     }
398
399     return 0;
400 }
401
402 /*
403  * Launches the child process.
404  */
405 static int _do_child(process_info_t *process) {
406     ptrace(PTRACE_TRACEME, 0, NULL, NULL);
407
408     return execvp(process->path, process->argv);
409 }
410
411 int patcher_attach_process(process_info_t *process) {
412     process->pid = fork();
413
414     if (process->pid != 0) {
415         DBG("child pid: %d\n", process->pid);
416         return _do_parent(process);
417     } else {
418         return _do_child(process);
419     }
420 }
421
422 int patcher_modify_process(process_info_t *process) {
423     int rc;
424
425     rc = _replace_mem(process);
426     if (rc < 0) {
427         fprintf(stderr, "[liboptimizer] ERROR: Failed to replace process memory\n");
428
429         return -1;
430     }
431 }
```



```
432     // NOTE: Uncomment if you want to debug the child process execution step
433     // by step.
434     #ifdef LIBOPTIMIZER_DEBUG
435         debug(process);
436     #endif
437
438     return 0;
439 }
440
441 int patcher_continue_exec(process_info_t *process, bool wait_for_exit) {
442     int rc;
443
444     // Continue child execution.
445     rc = ptrace(PTRACE_CONT, process->pid, NULL, NULL);
446     if (rc < 0) {
447         perror("[liboptimizer] PTRACE_CONT");
448
449         return -1;
450     }
451
452     // Wait for child to quit if asked to.
453     if (wait_for_exit) {
454         waitpid(process->pid, NULL, 0);
455     }
456
457     return 0;
458 }
```

D.5 liboptimizer

liboptimizer.c

```
1  /*
2   * File: liboptimizer.c
3   *
4   * Created by: Lucas Elisei <lucas.elisei@heig-vd.ch>
5   *
6   * Main entry point of the library.
7   */
8
9  #include <inttypes.h>
10 #include <limits.h>
11 #include <stdlib.h>
12
13 #include "debug.h"
14 #include "liboptimizer.h"
15 #include "live-patcher.h"
16 #include "elfparser.h"
17 #include "retdec.h"
18 #include "utils.h"
19
20 static int _recompile_function(process_info_t *process) {
21     return retdec_recompile(process);
22 }
23
24 char *symbol_at_address(const char *path, uint64_t address) {
25     return get_symbol_at_address(path, address);
26 }
27
28 process_info_t *init_process(int argc, char **argv, const char *function_name) {
29     // Pointer to process information.
30     process_info_t *process;
31     // Store return codes.
32     int rc;
33     // Absolute path of the binary.
34     char absolute_path[PATH_MAX];
35
36     // Allocate memory for `process_info_t` struct and set the memory to zero.
37     process = (process_info_t *)calloc(1, sizeof(process_info_t));
38     if (process == NULL) {
39         perror("[liboptimizer] calloc");
40
41         return NULL;
42     }
43
44     // Retrieve absolute path of the binary.
45     if (get_absolute_path(absolute_path, argv[0]) == NULL) {
46         perror("[liboptimizer] get_absolute_path");
47
48         goto _failed_abs_path;
49     }
50
51     DBG("absolute_path: %s\n", absolute_path);
52
53     // Initialize the fields we already have information about.
54     process->argc = argc;
55     process->argv = argv;
56     process->path = absolute_path;
57     process->function_name = function_name;
58     process->is64 = is64bit(process->path);
59
60     // Get function offset from function_name.
61     process->function_offset = get_symbol_address(process->path, process->function_name);
62     if (process->function_offset == 0) {
63         fprintf(stderr, "[liboptimizer] ERROR: Could not get function offset\n");
64
65         goto _failed_func_offset;
66     }
67
68     // Call the RetDec's script to generate LLVM IR of the function passed as
```

```
69     // argument.
70     rc = _recompile_function(process);
71     if (rc < 0) {
72         fprintf(stderr, "[liboptimizer] ERROR: Error while recompiling function\n");
73
74         goto _failed_recompilation;
75     }
76
77     // Attach the child process.
78     rc = patcher_attach_process(process);
79     if (rc < 0) {
80         fprintf(stderr, "[liboptimizer] ERROR: Error while attaching process\n");
81
82         goto _failed_attach;
83     }
84
85     return process;
86
87     // _not64bit:
88     _failed_func_offset:
89     _failed_abs_path:
90     _failed_recompilation:
91     _failed_attach:
92         free(process);
93
94         return NULL;
95     }
96
97     int modify_process(process_info_t *process) {
98         int rc;
99
100         rc = patcher_modify_process(process);
101         if (rc < 0) {
102             fprintf(stderr, "[liboptimizer] ERROR: Error while modifying process memory\n");
103
104             return -1;
105         }
106
107         return 0;
108     }
109
110     int execute_process(process_info_t *process, bool wait_for_exit) {
111         int rc;
112
113         rc = patcher_continue_exec(process, wait_for_exit);
114         if (rc < 0) {
115             fprintf(stderr, "[liboptimizer] ERROR: Error while continuing process execution\n");
116
117             return -1;
118         }
119
120         return 0;
121     }
122
123     #ifdef LIBOPTIMIZER_DEBUG
124     void print_process_info(process_info_t *process) {
125         int i;
126         FILE *fp;
127         char filename[PATH_MAX];
128
129         sprintf(filename, "%s%d.bin", process->function_name, process->is64 ? 64 : 32);
130
131         fp = fopen(filename, "wb");
132
133         DBG("== Printing process_info_t at 0x%" PRIXPTR "\n", (uintptr_t)process);
134         DBG("  path.....: %s\n", process->path);
135         DBG("  argc.....: %d\n", process->argc);
136         DBG("  argv.....");
137         for (i = 0; i < process->argc; ++i) {
138             fprintf(stderr, " %s", process->argv[i]);
139         }
140         fprintf(stderr, "\n");
141         DBG("  is64.....: %d\n", process->is64);
```

```
142     DBG(" pid.....: %d\n", process->pid);
143     DBG(" function_name.....: %s\n", process->function_name);
144     DBG(" function_offset.....: 0x%" PRIx64 "\n", process->function_offset);
145     DBG(" codesegment_address....: 0x%" PRIx64 "\n", process->codesegment_address);
146     DBG(" freesegment_address....: 0x%" PRIx64 "\n", process->freesegment_address);
147     DBG(" optimized_function_size: %zu\n", process->optimized_function_size);
148     DBG(" optimized_function.....: 0x%" PRIXPTR "\n", (uintptr_t)process->optimized_function);
149     if (process->optimized_function_size > 0) {
150         DBG(" ");
151         for (i = 0; i < process->optimized_function_size; ++i) {
152             uint8_t byte = process->optimized_function[i];
153
154             fputc(byte, fp);
155
156             fprintf(stderr, " %02X", byte);
157             if (i % 8 == 7) {
158                 fprintf(stderr, "\n");
159                 DBG(" ");
160             }
161         }
162         fprintf(stderr, "\n");
163     }
164     DBG("=====\n\n");
165
166     fclose(fp);
167 }
168 #endif // LIBOPTIMIZER_DEBUG
```

utils.c

```
1  /*
2   * File: utils.c
3   *
4   * Created by: Lucas Elisei <lucas.elisei@heig-vd.ch>
5   */
6
7  #include <limits.h>
8  #include <stdio.h>
9  #include <stdlib.h>
10 #include <string.h>
11 #include <unistd.h>
12
13 #include "utils.h"
14
15 char *get_absolute_path(char *dest, const char *path) {
16     char buffer[PATH_MAX];
17
18     if (*path == '/') {
19         strcpy(buffer, path);
20     }
21     else {
22         if (getcwd(buffer, PATH_MAX) == NULL) {
23             perror("[liboptimizer] getcwd error");
24
25             return NULL;
26         }
27
28         strcat(buffer, "/");
29         strcat(buffer, path);
30     }
31
32     if (realpath(buffer, dest) == NULL) {
33         perror("[liboptimizer] realpath error");
34
35         return NULL;
36     }
37
38     return dest;
39 }
```

D.6 mmult

mmult.c

```
1  /*
2   * File: mmult.c
3   *
4   * Created by: Lucas Elisei <lucas.elisei@heig-vd.ch>
5   */
6
7  #include <stdint.h>
8  #include <stdio.h>
9  #include <stdlib.h>
10 #include <string.h>
11
12 #include "mmult.h"
13
14 #define FILENAME    "/dev/urandom"
15
16 void matrix_init_data(matrix_t *m) {
17     FILE* fp;
18     size_t i;
19
20     fp = fopen(FILENAME, "rb");
21     if (fp == NULL) {
22         fprintf(stderr, "Could not open %s\n", FILENAME);
23
24         exit(EXIT_FAILURE);
25     }
26
27     for (i = 0; i < m->row * m->col; ++i) {
28         uint8_t data;
29
30         if (fread(&data, sizeof(uint8_t), 1, fp) != 1) {
31             fprintf(stderr, "fread error\n");
32             fclose(fp);
33
34             exit(EXIT_FAILURE);
35         }
36
37         m->data[i] = data % 8;
38     }
39
40     fclose(fp);
41 }
42
43 matrix_t *matrix_init(size_t row, size_t col) {
44     matrix_t *m;
45
46     m = (matrix_t *)malloc(sizeof(matrix_t));
47     if (m == NULL) {
48         fprintf(stderr, "Could not allocate memory\n");
49
50         return NULL;
51     }
52
53     m->data = (uint32_t *)malloc(row * col * sizeof(uint32_t));
54     if (m->data == NULL) {
55         fprintf(stderr, "Could not allocate memory\n");
56         free(m);
57
58         return NULL;
59     }
60
61     m->col = col;
62     m->row = row;
63
64     return m;
65 }
66
67 int matrix_mult(matrix_t *a, matrix_t *b, matrix_t *res) {
68     size_t i, j, k;
```

```
69
70     for (i = 0; i < res->row * res->col; ++i) {
71         res->data[i] = 0;
72     }
73
74     for (i = 0; i < a->row; ++i) {
75         for (j = 0; j < a->col; ++j) {
76             for (k = 0; k < b->col; ++k) {
77                 res->data[i*res->col + j] += a->data[i*res->col + k] * b->data[k*res->col + j];
78             }
79         }
80     }
81
82     return 0;
83 }
84
85 void matrix_save(matrix_t *m, const char *filename) {
86     size_t i, j;
87     FILE *fp;
88
89     fp = fopen(filename, "w");
90     if (fp == NULL) {
91         perror("fopen()");
92
93         return;
94     }
95
96     for (i = 0; i < m->row; ++i) {
97         for (j = 0; j < m->col; ++j) {
98             fprintf(fp, "%3d ", m->data[i*m->col + j]);
99         }
100        fprintf(fp, "\n");
101    }
102
103    fclose(fp);
104 }
```

main.c

```
1  /*
2   * File: main.c
3   *
4   * Created by: Lucas Elisei <lucas.elisei@heig-vd.ch>
5   */
6
7  #define _POSIX_C_SOURCE 199309L
8
9  #include <stdio.h>
10 #include <stdlib.h>
11 #include <time.h>
12 #include <unistd.h>
13
14 #include "mmult.h"
15
16 int main(int argc, char **argv) {
17     matrix_t *m1, *m2, *res;
18     size_t col, row;
19     struct timespec start, stop;
20
21     if (argc != 3) {
22         fprintf(stdout, "usage: %s <row> <col>\n", argv[0]);
23
24         return EXIT_FAILURE;
25     }
26
27     row = (size_t)atoi(argv[1]);
28     col = (size_t)atoi(argv[2]);
29
30     m1 = matrix_init(row, col);
31     m2 = matrix_init(col, row);
```

```

32     res = matrix_init(row, row);
33
34     matrix_init_data(m1);
35     matrix_init_data(m2);
36
37     matrix_save(m1, "m1.mat");
38     matrix_save(m2, "m2.mat");
39
40     if (clock_gettime(CLOCK_REALTIME, &start) == -1) {
41         perror("clock_gettime()");
42         exit(EXIT_FAILURE);
43     }
44     matrix_mult(m1, m2, res);
45     if (clock_gettime(CLOCK_REALTIME, &stop) == -1) {
46         perror("clock_gettime()");
47         exit(EXIT_FAILURE);
48     }
49
50     printf("Elapsed time for matrix multiplication (%zu%zu): %lus %luns\n",
51           row, col, (stop.tv_sec - start.tv_sec),
52           (stop.tv_nsec - start.tv_nsec));
53
54     matrix_save(res, "res.mat");
55
56     return EXIT_SUCCESS;
57 }

```

Machine code of non-optimized 32-bit matrix_mult function

```

1  00001424 <matrix_mult>:
2      1424:    55                push    %ebp
3      1425:    89 e5             mov     %esp,%ebp
4      1427:    56                push    %esi
5      1428:    53                push    %ebx
6      1429:    83 ec 10          sub     $0x10,%esp
7      142c:    e8 eb 01 00 00    call   161c <_x86.get_pc_thunk.ax>
8      1431:    05 cf 2b 00 00    add     $0x2bcf,%eax
9      1436:    c7 45 f4 00 00 00 movl    $0x0,-0xc(%ebp)
10     143d:    eb 18             jmp     1457 <matrix_mult+0x33>
11     143f:    8b 45 10          mov     0x10(%ebp),%eax
12     1442:    8b 40 08          mov     0x8(%eax),%eax
13     1445:    8b 55 f4          mov     -0xc(%ebp),%edx
14     1448:    c1 e2 02          shl     $0x2,%edx
15     144b:    01 d0             add     %edx,%eax
16     144d:    c7 00 00 00 00 00 movl    $0x0,(%eax)
17     1453:    83 45 f4 01       addl    $0x1,-0xc(%ebp)
18     1457:    8b 45 10          mov     0x10(%ebp),%eax
19     145a:    8b 10             mov     (%eax),%edx
20     145c:    8b 45 10          mov     0x10(%ebp),%eax
21     145f:    8b 40 04          mov     0x4(%eax),%eax
22     1462:    0f af c2          imul    %edx,%eax
23     1465:    39 45 f4          cmp     %eax,-0xc(%ebp)
24     1468:    72 d5             jb      143f <matrix_mult+0x1b>
25     146a:    c7 45 f4 00 00 00 movl    $0x0,-0xc(%ebp)
26     1471:    e9 bf 00 00 00    jmp     1535 <matrix_mult+0x111>
27     1476:    c7 45 f0 00 00 00 movl    $0x0,-0x10(%ebp)
28     147d:    e9 a0 00 00 00    jmp     1522 <matrix_mult+0xfe>
29     1482:    c7 45 ec 00 00 00 movl    $0x0,-0x14(%ebp)
30     1489:    e9 81 00 00 00    jmp     150f <matrix_mult+0xeb>
31     148e:    8b 45 10          mov     0x10(%ebp),%eax
32     1491:    8b 50 08          mov     0x8(%eax),%edx
33     1494:    8b 45 10          mov     0x10(%ebp),%eax
34     1497:    8b 40 04          mov     0x4(%eax),%eax
35     149a:    0f af 45 f4       imul    -0xc(%ebp),%eax
36     149e:    89 c1             mov     %eax,%ecx
37     14a0:    8b 45 f0          mov     -0x10(%ebp),%eax
38     14a3:    01 c8             add     %ecx,%eax
39     14a5:    c1 e0 02          shl     $0x2,%eax
40     14a8:    01 d0             add     %edx,%eax
41     14aa:    8b 08             mov     (%eax),%ecx

```

```

42      14ac: 8b 45 08      mov     0x8(%ebp),%eax
43      14af: 8b 50 08      mov     0x8(%eax),%edx
44      14b2: 8b 45 10      mov     0x10(%ebp),%eax
45      14b5: 8b 40 04      mov     0x4(%eax),%eax
46      14b8: 0f af 45 f4    imul    -0xc(%ebp),%eax
47      14bc: 89 c3         mov     %eax,%ebx
48      14be: 8b 45 ec      mov     -0x14(%ebp),%eax
49      14c1: 01 d8         add     %ebx,%eax
50      14c3: c1 e0 02      shl     $0x2,%eax
51      14c6: 01 d0         add     %edx,%eax
52      14c8: 8b 10         mov     (%eax),%edx
53      14ca: 8b 45 0c      mov     0xc(%ebp),%eax
54      14cd: 8b 58 08      mov     0x8(%eax),%ebx
55      14d0: 8b 45 10      mov     0x10(%ebp),%eax
56      14d3: 8b 40 04      mov     0x4(%eax),%eax
57      14d6: 0f af 45 ec    imul    -0x14(%ebp),%eax
58      14da: 89 c6         mov     %eax,%esi
59      14dc: 8b 45 f0      mov     -0x10(%ebp),%eax
60      14df: 01 f0         add     %esi,%eax
61      14e1: c1 e0 02      shl     $0x2,%eax
62      14e4: 01 d8         add     %ebx,%eax
63      14e6: 8b 00         mov     (%eax),%eax
64      14e8: 0f af d0      imul    %eax,%edx
65      14eb: 8b 45 10      mov     0x10(%ebp),%eax
66      14ee: 8b 58 08      mov     0x8(%eax),%ebx
67      14f1: 8b 45 10      mov     0x10(%ebp),%eax
68      14f4: 8b 40 04      mov     0x4(%eax),%eax
69      14f7: 0f af 45 f4    imul    -0xc(%ebp),%eax
70      14fb: 89 c6         mov     %eax,%esi
71      14fd: 8b 45 f0      mov     -0x10(%ebp),%eax
72      1500: 01 f0         add     %esi,%eax
73      1502: c1 e0 02      shl     $0x2,%eax
74      1505: 01 d8         add     %ebx,%eax
75      1507: 01 ca         add     %ecx,%edx
76      1509: 89 10         mov     %edx,%eax
77      150b: 83 45 ec 01    addl    $0x1,-0x14(%ebp)
78      150f: 8b 45 0c      mov     0xc(%ebp),%eax
79      1512: 8b 40 04      mov     0x4(%eax),%eax
80      1515: 39 45 ec      cmp     %eax,-0x14(%ebp)
81      1518: 0f 82 70 ff ff  jb      148e <matrix_mult+0x6a>
82      151e: 83 45 f0 01    addl    $0x1,-0x10(%ebp)
83      1522: 8b 45 08      mov     0x8(%ebp),%eax
84      1525: 8b 40 04      mov     0x4(%eax),%eax
85      1528: 39 45 f0      cmp     %eax,-0x10(%ebp)
86      152b: 0f 82 51 ff ff  jb      1482 <matrix_mult+0x5e>
87      1531: 83 45 f4 01    addl    $0x1,-0xc(%ebp)
88      1535: 8b 45 08      mov     0x8(%ebp),%eax
89      1538: 8b 00         mov     (%eax),%eax
90      153a: 39 45 f4      cmp     %eax,-0xc(%ebp)
91      153d: 0f 82 33 ff ff  jb      1476 <matrix_mult+0x52>
92      1543: b8 00 00 00 00  mov     $0x0,%eax
93      1548: 83 c4 10      add     $0x10,%esp
94      154b: 5b           pop     %ebx
95      154c: 5e           pop     %esi
96      154d: 5d           pop     %ebp
97      154e: c3           ret

```

Machine code of optimized 32-bit matrix_mult function

```

1      00000000 <.data>:
2      0: 55           push    %ebp
3      1: 53           push    %ebx
4      2: 57           push    %edi
5      3: 56           push    %esi
6      4: 83 ec 18     sub     $0x18,%esp
7      7: 8b 4c 24 34  mov     0x34(%esp),%ecx
8      b: 8b 44 24 04  mov     0x4(%esp),%eax
9      f: 89 44 24 0c  mov     %eax,0xc(%esp)
10     13: 8b 04 24     mov     (%esp),%eax
11     16: 89 44 24 14  mov     %eax,0x14(%esp)

```


12	1a:	8b 41 04	mov	0x4(%ecx),%eax
13	1d:	0f af 01	imul	(%ecx),%eax
14	20:	85 c0	test	%eax,%eax
15	22:	74 24	je	0x48
16	24:	31 c0	xor	%eax,%eax
17	26:	31 f6	xor	%esi,%esi
18	28:	90	nop	
19	29:	90	nop	
20	2a:	90	nop	
21	2b:	90	nop	
22	2c:	90	nop	
23	2d:	90	nop	
24	2e:	90	nop	
25	2f:	90	nop	
26	30:	8b 79 08	mov	0x8(%ecx),%edi
27	33:	c7 04 38 00 00 00 00	movl	\$0x0,(%eax,%edi,1)
28	3a:	46	inc	%esi
29	3b:	8b 79 04	mov	0x4(%ecx),%edi
30	3e:	0f af 39	imul	(%ecx),%edi
31	41:	83 c0 04	add	\$0x4,%eax
32	44:	39 fe	cmp	%edi,%esi
33	46:	72 e8	jb	0x30
34	48:	8b 4c 24 2c	mov	0x2c(%esp),%ecx
35	4c:	8b 19	mov	(%ecx),%ebx
36	4e:	85 db	test	%ebx,%ebx
37	50:	0f 84 e5 00 00 00	je	0x13b
38	56:	8b 54 24 30	mov	0x30(%esp),%edx
39	5a:	8b 69 04	mov	0x4(%ecx),%ebp
40	5d:	31 f6	xor	%esi,%esi
41	5f:	89 e8	mov	%ebp,%eax
42	61:	90	nop	
43	62:	90	nop	
44	63:	90	nop	
45	64:	90	nop	
46	65:	90	nop	
47	66:	90	nop	
48	67:	90	nop	
49	68:	90	nop	
50	69:	90	nop	
51	6a:	90	nop	
52	6b:	90	nop	
53	6c:	90	nop	
54	6d:	90	nop	
55	6e:	90	nop	
56	6f:	90	nop	
57	70:	85 c0	test	%eax,%eax
58	72:	0f 84 b8 00 00 00	je	0x130
59	78:	8b 42 04	mov	0x4(%edx),%eax
60	7b:	31 ff	xor	%edi,%edi
61	7d:	89 74 24 10	mov	%esi,0x10(%esp)
62	81:	90	nop	
63	82:	90	nop	
64	83:	90	nop	
65	84:	90	nop	
66	85:	90	nop	
67	86:	90	nop	
68	87:	90	nop	
69	88:	90	nop	
70	89:	90	nop	
71	8a:	90	nop	
72	8b:	90	nop	
73	8c:	90	nop	
74	8d:	90	nop	
75	8e:	90	nop	
76	8f:	90	nop	
77	90:	85 c0	test	%eax,%eax
78	92:	b8 00 00 00 00	mov	\$0x0,%eax
79	97:	74 7f	je	0x118
80	99:	31 ed	xor	%ebp,%ebp
81	9b:	89 7c 24 08	mov	%edi,0x8(%esp)
82	9f:	90	nop	
83	a0:	8b 44 24 34	mov	0x34(%esp),%eax
84	a4:	89 c2	mov	%eax,%edx

85	a6:	8b 42 08	mov	0x8(%edx),%eax
86	a9:	8b 52 04	mov	0x4(%edx),%edx
87	ac:	0f af d6	imul	%esi,%edx
88	af:	8b 7c 24 08	mov	0x8(%esp),%edi
89	b3:	8d 34 3a	lea	(%edx,%edi,1),%esi
90	b6:	8b 04 b0	mov	(%eax,%esi,4),%eax
91	b9:	8b 71 08	mov	0x8(%ecx),%esi
92	bc:	89 14 24	mov	%edx,(%esp)
93	bf:	01 ea	add	%ebp,%edx
94	c1:	8b 14 96	mov	(%esi,%edx,4),%edx
95	c4:	8b 5c 24 30	mov	0x30(%esp),%ebx
96	c8:	8b 73 08	mov	0x8(%ebx),%esi
97	cb:	89 34 24	mov	%esi,(%esp)
98	ce:	8b 4c 24 34	mov	0x34(%esp),%ecx
99	d2:	8b 49 04	mov	0x4(%ecx),%ecx
100	d5:	0f af cd	imul	%ebp,%ecx
101	d8:	89 4c 24 04	mov	%ecx,0x4(%esp)
102	dc:	01 f9	add	%edi,%ecx
103	de:	0f af 14 8e	imul	(%esi,%ecx,4),%edx
104	e2:	8b 74 24 10	mov	0x10(%esp),%esi
105	e6:	01 c2	add	%eax,%edx
106	e8:	8b 4c 24 34	mov	0x34(%esp),%ecx
107	ec:	8b 41 08	mov	0x8(%ecx),%eax
108	ef:	89 04 24	mov	%eax,(%esp)
109	f2:	8b 49 04	mov	0x4(%ecx),%ecx
110	f5:	0f af ce	imul	%esi,%ecx
111	f8:	89 4c 24 04	mov	%ecx,0x4(%esp)
112	fc:	01 f9	add	%edi,%ecx
113	fe:	89 14 88	mov	%edx,(%eax,%ecx,4)
114	101:	8b 4c 24 2c	mov	0x2c(%esp),%ecx
115	105:	45	inc	%ebp
116	106:	8b 43 04	mov	0x4(%ebx),%eax
117	109:	39 c5	cmp	%eax,%ebp
118	10b:	72 93	jb	0xa0
119	10d:	8b 69 04	mov	0x4(%ecx),%ebp
120	110:	8b 54 24 30	mov	0x30(%esp),%edx
121	114:	8b 7c 24 08	mov	0x8(%esp),%edi
122	118:	47	inc	%edi
123	119:	39 ef	cmp	%ebp,%edi
124	11b:	0f 82 6f ff ff ff	jb	0x90
125	121:	8b 19	mov	(%ecx),%ebx
126	123:	89 e8	mov	%ebp,%eax
127	125:	46	inc	%esi
128	126:	39 de	cmp	%ebx,%esi
129	128:	0f 82 42 ff ff ff	jb	0x70
130	12e:	eb 0b	jmp	0x13b
131	130:	31 c0	xor	%eax,%eax
132	132:	46	inc	%esi
133	133:	39 de	cmp	%ebx,%esi
134	135:	0f 82 35 ff ff ff	jb	0x70
135	13b:	8b 44 24 14	mov	0x14(%esp),%eax
136	13f:	89 04 24	mov	%eax,(%esp)
137	142:	8b 44 24 0c	mov	0xc(%esp),%eax
138	146:	89 44 24 04	mov	%eax,0x4(%esp)
139	14a:	31 c0	xor	%eax,%eax
140	14c:	83 c4 18	add	\$0x18,%esp
141	14f:	5e	pop	%esi
142	150:	5f	pop	%edi
143	151:	5b	pop	%ebx
144	152:	5d	pop	%ebp
145	153:	c3	ret	