Transforming Malicious Code to ROP Gadgets for Antivirus Evasion

Giorgos Poulios [1], Christoforos Ntantogian [1], Giorgos Karopoulos [2], Christos Xenakis [1*]

[1] Department of Digital Systems, University of Piraeus, Greece
[2] Department of Informatics and Telecommunications, University of Athens, Greece
[*] xenakis@unipi.gr

Abstract: The downside of current polymorphism techniques lies to the fact that they require a writeable code section, either marked as such in the corresponding Portable Executable (PE) section header, or by changing permissions during runtime. Both approaches are identified by AV software as alarming characteristics and/or behavior, since they are rarely found in benign PEs unless they are packed. In this paper we propose the use of Return-Oriented Programming (ROP) as a new way to achieve polymorphism and evade AV software. To this end, we have developed a tool named ROPInjector which, given any piece of shellcode and any non-packed Portable Executable (PE) file, it transforms the shellcode to its ROP equivalent and patches it into (i.e. infects) the PE file. After trying various combinations of evasion techniques, the results show that ROPInjector can evade nearly and completely all antivirus software employed in the online VirusTotal service. The main outcome of this research is the developed algorithms for: a) analysis and manipulation of assembly code on the x86 instruction set, and b) the automatic chaining of gadgets by ROPInjector to form safe, and functional ROP code that is equivalent to a given shellcode.

1. Introduction

Return Oriented Programming (ROP) gained increased attention during the late 2000's [3] as an advanced stack smashing method that could bypass Data Execution Prevention (DEP) mechanisms. ROP is a rediscovery of threaded code in which programs typically consist of a chain of addresses in the stack pointing to code chunks in the attacked executable (or its loaded libraries) each of them ending with a return instruction (commonly ret but not only). These borrowed code chunks are called *gadgets* and their "return" is in fact a call to the next gadget in the chain. As an analogy to regular code, in ROP, gadgets are the "instructions" and esp is the program counter.

Polymorphism is a technique for AV bypass in which the code changes itself each time it runs, but the function of the code (its semantics) does not change at all. In this way, AVs cannot create a signature for detection of the shellcode. However, the downside of current polymorphism techniques lies to the fact that they require a writeable code section, either marked as such in the corresponding PE section header, or by changing permissions during runtime. Both approaches are identified by AV software as alarming characteristics and/or behavior, since they are rarely found in benign PEs unless they are packed.

In this work, we claim Return-Oriented Programming (ROP) to be a strong polymorphism alternative that eliminate the need of writable code section. More specifically, the first and most important benefit of using ROP for AV evasion is that such borrowed code (that of gadgets) is always benign and tested against false positives. Evidently, the return address chain has to be built somehow onto the stack and that would leave a footprint subject to signing. The process involves either pushing the return addresses to the stack or just copying the whole chain from another memory location (possibly some .data segment) and adjusting the stack pointer. However, we argue that: i) the code required for such operations is very common and seemingly benign, ii) can be randomized or encoded in many and trivial ways, iii) it largely depends on the attacked PE and its image base since in the worst case it is a series of push <VAi> operations. This holds because gadget addresses change for different PEs and different image bases, hence changing the footprint and statistics of the chain building instructions even if they originate from the same source shellcode. Given these features, ROP enables polymorphism **without requiring a writeable code section in memory**. Encoding/decoding can be applied on the gadget chain in memory (i.e. in the stack and not in the code section) and/or different gadgets can be randomly chosen for the same operation hence altering the shellcode's footprint.

Based on the above observations, in this paper we present ROPInjector, a tool which, given any piece of shellcode (hereafter, also referred to as *source (shell)code*) and any non-packed executable file, it transforms the shellcode into its ROP equivalent and patches it into (i.e. infects) the PE file. ROPInjector, which is written in C programming language, infects Portable Executables (PEs) for Windows OS (a previous version of the tool has been presented in Blackhat [17]). Since it is very common for AVs to detect minor deviations from the typical arrangement of the file sections and their characteristics (e.g. a second executable section with RWX permissions), besides the transformation of the code into a non-recognizable, non-recurrent form, the developed tool addresses several additional issues to achieve evasion, such as the positioning of the shellcode in the carrier executable and the way of transferring control to the shellcode. Moreover, we have performed several experiments to evaluate the effectiveness of the proposed tool by injecting shellcodes to well-known executable files including acrobat reader, firefox, Java, etc. Quantitative results show that our proposed technique, if combined with simple behavioral anti-profiling techniques may render AV detection infeasible.

The rest of the paper is organized as follows. Section 2 presents the related work, while section 3 provides the required background for ROP. Section 4 elaborates on the architecture of the proposed ROPInjector and its functionality details. In section 5, we analyze experimental results and in section 6 we provide a discussion of possible mitigation techniques. Finally, section 7 concludes the article.

2. Related Work

While the traditional use of ROP is software exploitation (i.e., bypass non-executable stack and heap), there are some previous works that have proposed alternative uses of ROP. More specifically, in [8] the authors propose ROP for benign purposes; specifically, they use ROP for software watermarking. The proposed ROP-based watermarking is able to transform watermarking code into ROP gadgets and build them in the data region. Once triggered using a secret message, the pre-constructed ROP execution will recover the hidden watermark message. The proposed method ensures that the watermarked program does not have an explicit code stream that belongs exclusively to watermarking. Instead the authors use operating system libraries to borrow the ROP gadgets, preventing detection by software analysis. Towards this direction, RopSteg [9] has been proposed for program steganography. The latter is a variation of software obfuscation but it differs from it, since in program steganography the instructions are hidden instead of being transformed. RopSteg achieves to hide selected code protection by generating equivalent ROP gadgets and blending them into the executable. Finally, in [10], the authors propose ROPOB, a code transformation technique to obfuscate control flow using ROP. The main contribution of ROPOB is that due to the use of ROP to complete control flow transfer, static reverse engineering methods cannot discover the real control flow, even though they can disassemble software correctly. However, the main limitation of ROPOB is that through dynamic analysis the obfuscation trivially breaks. Although all the aforementioned works have implemented the proposed tools to evaluate their effectiveness, none of them are available in the internet (i.e., source code or in the form of an executable).

The work closest to ours is presented by Mohan et al. [15]. The authors have developed a metamorphic obfuscator called Frankenstein which is able to reassemble a given malware with code fragments entirely from other benign programs. Authors' motivation was the creation of malware variations from benign pieces of entirely randomly selected binaries residing in a system. Their goal is to avoid Signature Matching (i.e., syntax based heuristics) detection. They deduce the problem of generating mutations into a searching problem. Their proposed method is able to search for segments of code found in benign binaries and evaluate them semantically with the given malicious instructions. The evaluation is based on a symbolic machine state. Finally, it performs the suitable code arrangements to construct the final payload. Frankenstein has several limitations compared to our proposed ROPInjector. First, the authors consider a relaxed version of a gadget. While ROP considers that a gadget ends with the `ret` instruction, Frankestain definition of a gadget is any sequence of bytes that are interpretable as valid x86 instructions, since it statically stitch gadgets together. On the contrary, ROPInjector considers is a metamorphic malware generator based on the pure definition of ROP and gadgets. Second, Frankenstein purpose is to modify only code snippets which may look suspicious in a malware. On the contrary, ROPInjector the whole binary to its (one of the many) ROP counterpart. Third, authors have implemented a prototype which is not available and therefore, we could not repeat the experiments for a quantitative comparative analysis.

To the authors' best knowledge this is the first practical work that infects PEs with pure ROP-encoded payload. Nevertheless, in this section we examine two tools having the same purpose with ROPInjector, that is, to infect PE files with common (possibly encrypted) shellcode in order to bypass AV software.

The first, Shellter [1], focuses on maintaining the original structure of the PE file, by avoiding injection of the shellcode into predefined locations or changing the characteristics of the existing sections. It achieves so by overwriting existing code for which it is certain that will be given control during execution of the program. The latter is deduced by tracing the executable file and analyzing its execution flow. Shellter is also capable of reusing imports of the original PE file to change the writing permissions of the section containing the shellcode so that encrypted and self-modifying code can be used. It is also capable of injecting "junk code" before the shellcode that delays execution as a means to anti-emulation. Shellter is advanced in terms of dynamically selecting the location of the patch in the shellcode (as opposed to extending the .text section). However, while it features a patching method that introduces variability (as to where in the file is the shellcode injected), it relies on traditional polymorphism methods, that are still subject to signature generation and detection of write permissions or modifications of the .text section in memory. Moreover, our proposed approach introduces variability too, due to the transformation to ROP (which is dependent on the PE file).

PEinject [2] is mostly a method (and referenced as such) rather than a full-featured tool. It injects the shellcode in the (first sufficiently large) padding space found in the .text section (either 0xCC caves or section padding) and does not encode or modify the payload in any way, neither does it anticipate for self-modifying or encrypted payloads. Control is passed to the injected shellcode by modifying the address of entry point of the PE file's NT_HEADER. The evasion ratios of both methods are compared with our proposed approach in Section 5.

3. Return Oriented Programming

ROP gained increased attention [14] as an advanced stack smashing attack that could bypass Data Execution Prevention (DEP) mechanisms. It is a rediscovery of threaded code in which programs typically consist of a chain of addresses in the stack pointing to code chunks in the attacked executable (or its loaded libraries) each of them ending with a return instruction (commonly `ret`, 0xC3, but not only). These borrowed code chunks are called gadgets and their "return" is in fact a call to the next gadget in the chain. As an analogy to regular code, in ROP, gadgets are the "instructions" and esp is the program counter.

The first and most important benefit of using ROP for AV evasion is that such borrowed code (that of gadgets) is always benign and tested against false positives. Of course, the return address chain has to be built somehow in the stack and that would leave a footprint subject to signing. The process involves either pushing the return addresses to the stack or just copying the whole chain from another memory location (possibly some .data segment) and adjusting the stack pointer. However, a) the kind of code required for such operations is very common and seemingly benign, b) it largely depends on the attacked PE and its image base since in the worst case it is a series of push <VA$_i$> operations, and c) can be randomized or encrypted in many and trivial ways. Especially, what we mentioned for reason (b) holds because gadget addresses change for different PEs and different image bases, hence changing the footprint and statistics of the chain building instructions even if they are for the same source shellcode.

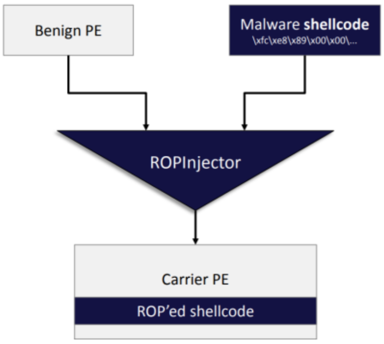

Fig. 1. *ROPInjector functionality*

Given these features, ROP enables polymorphism without requiring a writeable code section in memory (which is very rare in benign PEs unless they are packed, as well as a typical heuristic for detection). Encryption/decryption can be applied on the gadget chain in memory. (i.e. in the stack and not in the code section) and/or different gadgets can be randomly chosen for the same operation hence altering the malware's footprint.

4. ROPInjector

ROPInjector takes as in input a PE together with a malicious piece of code and outputs the PE which is infected with the malicious piece of code in a ROP form (see figure 1).

In general, ROPInjector approach can be divided into 6 distinct phases as follows:

1) Reverse analysis of machine code
2) Finding ROP gadgets in PE
3) Transform instructions to ROP equivalents using an intermediate representation language
4) Inject gadgets
5) Create chain from gadgets
6) Patch PE

In the next sections, we are going to analyse each of the above phases providing detailed examples to gain better understanding of the presented notions.

4.1 Reverse analysis of machine code

Reverse analysis of machine code into data structures that are easy to handle is crucial to perform any kind of patching, modifications, re-assembly, and any transformation to ROP. Two are the most important pieces of information required: i) the origin and destination of all relative references (e.g. a relative jump and its target) and ii) which registers are being written or read during each instruction, as well as which registers are free to modify. The former is required for injecting or removing instructions from a code segment without breaking its functioning. The latter is particularly useful to enhance gadget matching, either by performing permutations, or by using gadgets that contain redundant but safe instructions (in this case, *unsafe* are branch, privileged, or indirect addressing mode instructions because they risk raising errors such as access violation).

4.1.1: MOD/REG/RM and SIB unrolling: Instructions using the MOD/REG/RM indirect addressing mode with displacement or the Scaled Index Byte (SIB) addressing scheme in the shellcode are treated specially before the transformation to ROP. Such instructions are unwanted for the following reasons:

i) They are long (in the best and not so likely case 3 bytes long: 1 for opcode, 1 for MOD/REG/RM and 1 for SIB) hence unlikely to be found in gadgets;
ii) They often read many general purpose registers at once, thus reserving them while as mentioned earlier, the more the free registers the better;

Table 1 List of PE files used as carriers in the experiments

Executable	Version
AcroRd32.exe	Version DC of Adobe Acrobat Reader
Acrobat.exe	Version DC of Adobe Acrobat Pro
cmd.exe	Version 10.0.17134.165 Windows Command Prompt
Rainmeter.exe	Version 4.2.0 Build 3111 of Rainmeter
firefox.exe	Version 61.0.1 of Mozilla Firefox
java.exe	Version 10.0.1 of Oracle Java
wmplayer.exe	Version 12.0.9600.17415 of Microsoft Windows Media Player
nam.exe	Version 1.11 of "The Network Animator"
notepad++.exe	Version 7.5.5 of the GNU text editor for Windows

Table 2: List of patching scenarios tested against VirusTotal

Patching Scenario	**Description**
Original	The executable file is not patched at all
ROP-Exit	This is the executable file generated by the ROPInjector. The executable file is patched with the shellcode unrolled, converted to ROP, and entry point before the original program's exit (hook ExitProcess or exit)
Exit	In this scenario, the executable file is patched with the shellcode intact and entry point before the original program's exit (hook ExitProcess or exit)
Shellcode	The executable file is patched with the shellcode intact, and entry point before the original program.

iii) Their respective gadgets (should they be found or injected) will probably not be reusable, due to the use of displacement and index constants (e.g. `mov edx, [esi*2+16]`).

In order to circumvent this kind of situations, we reduce such instructions to their arithmetic equivalents one-by-one. We call this process *unrolling* and it is performed to the shellcode before any transformation to ROP. For instance, `[1] mov eax, [ebx+ecx*2]` may be replaced by:

```
[a']      Mo    v eax, ecx
[b']      sal eax, 1
[c']      add eax, ebx
[d']      mov eax, [eax]
```

If the register eax is not free to use for the arithmetic operations, another temporary register that is free may be used.

Noteworthy is how unrolling unlocks register access from one atomic instruction to many. For instance, in the latter example, ecx is freed at [a'] and ebx at [c']. If for example eax were to be freed at the preceding 10 instructions, then instructions [a'] to [c'] could be moved 10 instructions behind, thus resulting in an additional free register (i.e., ecx and ebx, but not eax which will not be free) in that preceding code chunk.

4.2 Finding gadgets

Candidate gadgets in the executable sections of the given PE file must end in one of :

- `ret,`
- `retn,`
- `pop regX;`
- `jmp regX,` or `jmp regX`.

Exceptionally for the latter, the gadget in question must be first paired with a *loader gadget* that loads the required return address into `regX`. The process begins by finding all gadget endings and temporarily storing them to a list. For each of those endings, *n* bytes of preceding machine code is disassembled for each *n* up to maximum depth *N* (typically 20 bytes). If such disassembly aligns with the ending (not guaranteed since x86 instructions are of variable length) a candidate gadget has been found. Candidate gadgets containing any illegal, privileged (e.g. `sysenter`, `int`, `iret`), branch or esp modifying instructions are filtered out.

4.3 Parsing gadgets into Intermediate Representation and One to One Permutations Between Source Code and Gadgets (ROP Transformation)

The gadgets found in the aforementioned process are first analyzed instruction-by-instruction to infer register access. Since gadgets are allowed to contain safe but redundant instructions, their register access is tested for modifications to the register in question (e.g. a `mov ecx, eax; pop ecx; ret;` gadget cannot be used for moving eax to ecx) as well as the non-free registers of the source instruction to be encoded.

Following that, they are parsed into an Intermediate Representation (IR) consisting of an operation-type, and 3 operands with different meaning depending on the type. If a multi-instruction gadget contains more than one representable instructions, only the first is considered. However, the following ones have also been considered in other gadgets with the same ending, because of the backwards gadget finding process described in the previous paragraph. Noteworthy is the fact that by parsing into this higher level IR, one-to-one permutations are automatically performed. That is because both gadgets and instructions

```
mov esp, ebp
pop ebp
ret(n)
CCCCCCCCCCCCCCCCCC
```

```
        jmp epilogue; normal flow avoiding gadget
        mov ecx, eax; the injected gadget
        jmp return; gadget flow avoiding std. epilogue
epilogue:
        mov esp, ebp
        pop ebp
return:
        ret(n)
        CCCCCCCC
```

Fig.2. *Injection of gadget (right) in 0xCC cave preceded by standard function epilogue (left)*

are classified into one of these types, based on which the encoding is then performed, rather than on the instructions per se. The IR is also useful for selecting the encoder function accompanying every gadget. Encoders are responsible to answer "*whether their assigned gadget can encode a given instruction*", as well as to encode it into a list of stack operations if requested to.

Predefined, one-to-one permutations (i.e. one instruction to one gadget) are achieved through the IR and encoder functions. Encoders will also perform basic algebraic permutations based on the properties of addition, subtraction multiplication and division. For instance, if the instruction to be encoded is of type ADD_IMM (`add reg, imm`), an encoder will repeat anything `add reg, x` with `x` being an integer divisor of `imm`, `imm/x` times. Addition and subtraction with constants will also be swapped if the signs of the constants are flipped. M-to-N permutations quickly scale to exponentially growing space and are out of the scope of this work.

4.4 Injecting Gadgets

In order to enhance transformation of the source shellcode, and since not all required gadgets are always found in the PE file, new ones are also injected as needed. Firstly, the 0xCC caves are used for this injection, and if they are filled, the .text section is extended before the actual patch. The injection is performed in the least noticeable way to avoid alarms. If a standard epilogue (`mov esp, ebp; pop ebp; ret`) is found right before the 0xCC cave, the gadget is injected in-between the preceding code and the epilogue. Figure 2 depicts such an example gadget injection of a mov ecx, eax gadget. In the case that no epilogue is found at the boundary with the 0xCC cave, a pseudo-function with standard prologue and epilogue is injected to avoid heuristics or n-grams that might raise suspicion due to non-ordinary returns. This pseudo-function has the following form shown in figure 3.

Following gadget insertions will then reuse this pseudo-epilogue as stated above, by injecting before the standard epilogue, thus making it look more like a real function.

```
        push ebp
        mov ebp, esp
        <gadget code>
        jmp return
        mov esp, ebp
        pop ebp
return:
        ret
```

Fig. 3. *Pseudo-function ending used during gadget injection*

4.5 Chaining gadgets

The return address chain can be built either during runtime or during compile-time and saved to the initialized data section of the file (to be then copied at runtime to the stack). The most alarming option would be the first (during runtime) and we choose this to evaluate our evasion ratio (also chosen as an implementation option). During this process, besides the pushing of the VAs onto the stack, the ROP compiler must consider pushing immediate constants, adjustments for stack pointer modifications in the gadget (e.g. redundant `pops`, `retns`) and gadgets with loader gadgets. For this purpose, the following types of *stack operations* are defined:

PUSH_VA ; *push a (loader) gadget VA onto the stack*
PUSH_IMM ; *push an immediate constant onto the stack*
ADVANCE ; *advance (subtract from) the stack pointer a number of bytes*
CHAIN ; *pseudo operation denoting a placeholder for the next gadget's VA*

The result of the encoding process of a given instruction by a given gadget is a series of stack operations for the invocation of the gadget. The list of such operations for all gadget calls describes the assembly instructions that if executed, will build the chain in the stack. Alternatively, such operations may be used to create the required stack frame during compile-time, save it as initialized data and copy it over from the data section during runtime. The latter process allows also for encoding/decoding of the stack frame. In the former case, and when multiple calls are made

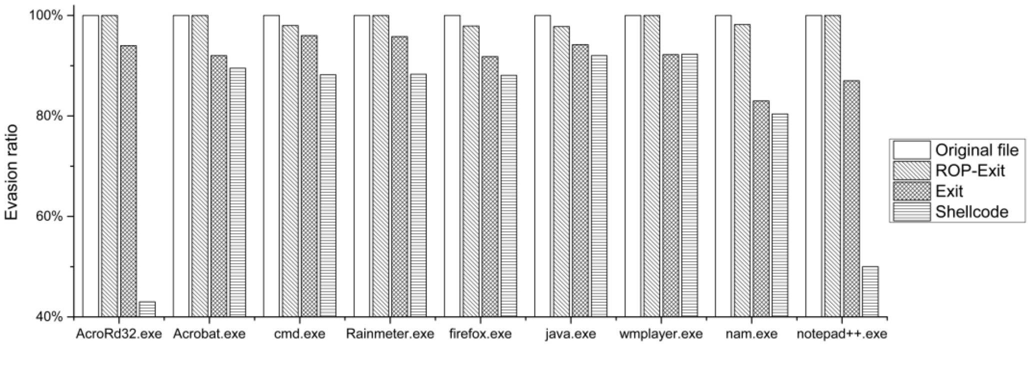

Fig. 4. *Evasion ratio of ROPInjector for the Meterpreter Shellcode*

to the same gadget (e.g. as in using `inc eax` to achieve `add eax, X`) the compiler wraps the call with a conditional jump loop using a free register.

However, not all types of instructions can be easily encoded into ROP. In this work we do not consider the encoding of branches (jumps, calls, loops, interrupts), privileged instructions and pops. Hence, the return-oriented code chunks must finally return back to the source shellcode. This is achieved by wrapping the chain building instructions in the following:

```
    [1] call build_chain
    [2] jmp past_the_chain
 build_chain:
    [3] push <VA of gadget N>
    [4] ....
    [5] push <VA of gadget 1>
    [6] ret
 past_the_chain:
    [7] <other instructions/chains>
```

In this way, the last gadget (N) will return to instruction [2] jumping past the chain building instructions and continuing normal execution flow.

4.6 PE Patching and Passing Control to the Shellcode

First of all, the patching of the PE file and the passing of control to the shellcode must be done in the least noticeable way. A second executable section hosting the shellcode would be too alarming, since the vast majority of executables has only one. The next least disruptive and easy to implement option would be to inject the shellcode in the 0xCC padding commonly left by the linker in-between code segments (typically OBJ files) in the .text section of PEs. However, there may not always be sufficient space in those 0xCC caves, while it is important to notice that ROPInjector puts this padding space into better use for our purposes as we analyze below.

For the above reasons we choose to append the shellcode to the existing .text section of the executable, and correct all section headers and relocations accordingly. To pass control to it, the default practice is to replace the instructions pointed to by `NT_HEADER.AddressOfEntryPoint` with a jump to the shellcode which is appended those replaced instructions followed by a jump back to the original execution flow. Directly pointing the address of entry point to the shellcode in this case is avoided, since many AVs' heuristics are alarmed by the fact that it points towards the end of .text. An alternative to giving control to the shellcode at program entry, is to hook any calls to `ExitProcess, exit` or other similar functions. This technique in particular, as shown also later by the results, bypasses behavioral profiling by AVs that employ emulation or sandboxing. This can be attributed to the fact that either AVs emulate only a small portion of the executable's entry code due to scanning time constraints, or because of lack of (universal) techniques for triggering a graceful exit (many programs do not handle SIGINT and SIGTERM signals).

An issue that arises when patching signed executables is that their checksum/hash, and thus their certificate, gets invalidated. This is obviously very alarming and would prevent us from testing our methods on popular executables of every-day use. Surprisingly though, hiding the certificate by erasing its pointer in the security data directory of NT_HEADER does not trigger any alarms. Of course, regardless of the certificate, the checksum of the PE file is recalculated and patched accordingly.

5. Experiments and Results

In order to evaluate ROPInjector we used the VirusTotal online antivirus scanning service [4] which at the time of this writing includes 57 AVs. For carrier PEs (i.e., the infected ones), we selected 9 popular 32-bit executable of various sizes that most of them also include certificates (see Table 1).

Regarding the source shellcode, we selected the most popular payload of Metasploit [5]: the Reverse TCP Meterpreter. For each PE and each shellcode we performed 4 patching scenarios as listed in Table 2, resulting in a total of 72 samples. Figure 4 depicts the evasion ratios ($1 - \frac{detection}{total\ \#\ of\ AVs}$) of ROPInjector for each one of the four scenarios for the widely used reverse Meterpreter shellcode. We can observe that the executables generated by the ROPInjector (i.e., "ROP-Exit" scenario) achieve the highest evasion ratio. In particular, in more than half of the test cases the ROPInjector results in 100% AV evasion, while in some PE files (e.g., java.exe), the ROPInjector has evasion ratio greater than 98.5% for the well-known Meterpreter shellcode. This means that in average ROPInjector achieves AV evasion equal to **99.31%,** as depicted in figure 5 (i.e., "ROP-Exit" scenario).

From these results we can deduce that evasion depends almost equally on both code obfuscation/transformation (hence signature evasion) and entry point (hence behavioral profiling evasion). This can be attributed to the fact that some AVs were able to detect ROPInjector despite the fact that there is no signature, due to the ROP polymorphism. It seems that behavioral analysis is equally important to static signatures for some AVs (from the ones that were alarmed) and is mostly performed during entry of executables.

Moreover, a comparison is also made with Shellter v2.2 [1] and PEinject [2] in figure 6. Shellter was used with its default options (i.e. with polymorphic junk code). We can observe that executables generated from our proposed ROPInjector (i.e., "ROP-Exit") have the highest evasion ratio in all conducted experiments compared to Shellter and PEinject. Note also that even the simple "Exit" scenario achieved in some executable files better results than Shellter. Finally, the peinject had the worst evasion ratio.It is also important to notice that besides VirusTotal, we have also tested the effectiveness of ROPInjector against a special piece of software named NCCGroup's "Experimental Windows .text section Patch Detector" [7]. This detector compares the executable sections in memory against the ones on disk to detect modifications/patching. As expected, no executable was detected as patched, since ROPInjector does not alter the .text section in memory (neither does it require to).

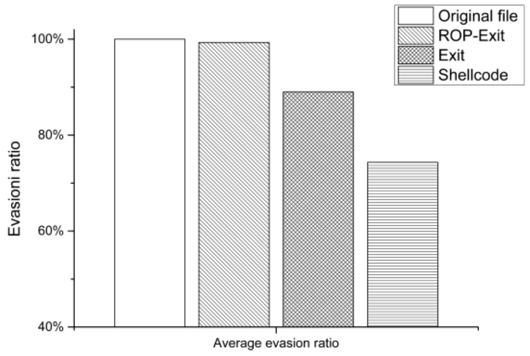

Fig. 5. *Average evasion ratio per combination of methods*

6. Discussion

Most antivirus software relies on string signatures and mild behavioral profiling detection mechanisms. By encoding the shellcode into its return-oriented equivalent and even by performing elementary mutations (unrolling), the former can be bypassed in the vast majority of cases. Dynamic analysis and behavioral profiling can also be avoided by carefully intercepting normal execution flow in points that AVs either cannot emulate or simply cannot derive enough evidence to classify the behavior as malicious. In this work, we presented as a means to the latter the hooking of common calls to process exit resulting in many cases in absolute evasion and in others rates greater than 98%. However, if the anti-dynamic analysis protection of ROPInjector is bypassed than the dynamic analysis of an infected PE can reveal the malicious code hidden in a ROP form. However, we consider this as out of scope since ROPinjector is mainly designed as an obfuscation tool to avoid static detection. Advanced techniques specialized for anti-dynamic analysis such as malwash [16] can be combined and further enhance evasion capabilities of ROPInjector.

The techniques presented can still be mitigated if dealt with individually. For instance, signatures could be created for ROP building instructions and behavioral analysis could be also performed backwards in terms of process life-cycle. However, since slight variations and randomization can again disarm scanners, a more robust countermeasure does not seem straight-forward to design, and/or practical to implement.

Several AVs rely on code statistics (such as entropy, n-grams and more), in order to classify PE files as benign or malicious. Unless such methods are designed to consider separate parts of the file's code, ROPInjector does not affect the statistic metrics of the binary file as a whole. The only metric that is affected and can be used by AVs for detection is the hash of the PE itself. In particular, for well-known PEs (such as Firefox.exe), AVs can cross-check the hash of the PE. If the hash is different from the official version of the executable, then AVs can raise an alarm. However, since new versions with updates are continuously released, maintaining the size of each PE is a challenging task.

Another solution which has a great impact on ROP, since it effectively mitigates software exploitation attacks based on ROP is the well-known Control Flow Integrity (CFI) [12]. The latter is a compile time mechanism for preventing malicious code from redirecting the execution flow of a program. In other words, the goal of CFI is to restrict the set of possible control-flow transfers to those that are strictly required for correct program execution. This prevents code-reuse techniques such as ROP from working because they would cause the program to execute control-flow transfers which are illegal under CFI [13]. By its own definition, we can deduce that CFI will not be triggered, since its purpose is to defend against external threats which will modify the operation of an executable. On the contrary, ROPInjector does not change the execution flow of an executable except for the exit point of the executable in order to trigger ROPInjector and transfer the code execution into the ROP chain, executing the malware. However, since this execution flow modification is not an indirect branch, CFI will not consider this as a policy violation and will not be triggered. We have verified this behavior against executables which are compiled with CFI (i.e., using Control Flow Guard implementation of Windows OS).

Perhaps the most promising direction is towards the strict coupling of the host operating system with the trusted software certificates (or checksums) and a "default distrust all" policy, i.e., whitelisting rather than blacklisting, which may have gain in popularity the last years in Windows domain environments (i.e., application whitelisting [6]), but not in standalone installations of Windows OS in personal environments.

7. Conclusions

In this paper we presented ROPInjector, a software tool which, given any piece of shellcode and any non-packed 32-bit Portable Executable (PE) file, it transforms the shellcode to its ROP equivalent and patches it into (i.e. infects) the PE

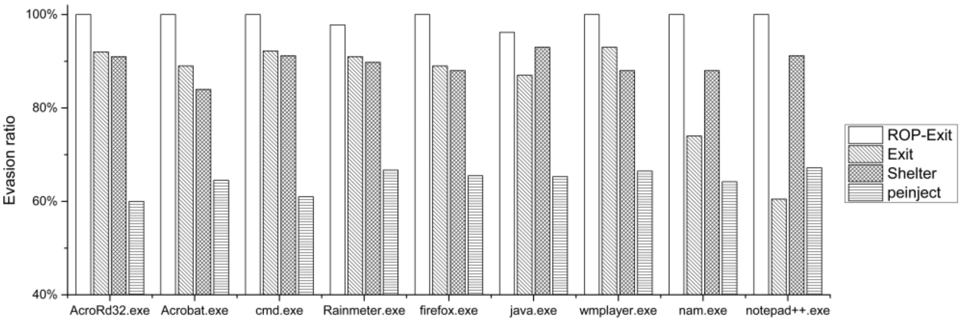

Fig. 6. *Comparison of evasion ratio between "ROP-Exit", "Exit" scenarios with Shelter and PEinject*

file. After trying various combinations of evasion techniques, the results show that ROPInjector can evade nearly and completely all antivirus software employed in the online VirusTotal service. The main outcome of this research is the developed algorithms for: a) analysis and manipulation of assembly code on the x86 instruction set, and b) the automatic chaining of gadgets by ROPInjector to form safe, and functional ROP code that is equivalent to a given shellcode. Currently, we are in the process of porting the implementation of ROPInjector tool, from C to Python programing language. This would benefit the project, since Python is more human friendly and the security community has already implemented libraries for reverse engineering, fact that would greatly help us to reduce the code base and complexity of ROPInjector.

8. Acknowledgments

This work was supported by the European Commission, under the FutureTPM project with Grant Agreement No. 779391.

References

[1] Shellter project, https://www.shellterproject.com, last accessed on April 15, 2018

[2] Injecting Shellcode into a Portable Executable(PE) using Python, http://www.debasish.in/2013/06/injecting-shellcode-into-portable.html, last accessed on April 15, 2018

[3] Shacham, Hovav. "The geometry of innocent flesh on the bone: Return-into-libc without function calls (on the x86)." Proceedings of the 14th ACM conference on Computer and communications security. ACM, 2007.

[4] VirusTotal, https://www.virustotal.com, last accessed on June 15, 2018

[5] Metasploit, http://www.metasploit.com/, last accessed on April 15, 2018

[6] Applocker, https://docs.microsoft.com/en-us/windows/security/threat-protection/windows-defender-application-control/applocker/applocker-overview, last accessed on February 15, 2018

[7] Experimental Windows .text section Patch Detector, https://github.com/nccgroup/WindowsPatchDetector, last accessed on July 15, 2018

[8] Ma, Haoyu, et al. "Software Watermarking using Return-Oriented Programming." Proceedings of the 10th ACM Symposium on Information, Computer and Communications Security. ACM, 2015.

[9] Lu, Kangjie, Siyang Xiong, and Debin Gao. "Ropsteg: program steganography with return oriented programming." Proceedings of the 4th ACM conference on Data and application security and privacy. ACM, 2014.

[10] Dongliang Mu et al. "ROPOB: Obfuscating Binary Code via Return Oriented Programming." International Conference on Security and Privacy in Communication Systems. Springer, Cham, 2017

[11] Jiang Ming et al. "BinSim: Trace-based semantic binary diffing via system call sliced segment equivalence checking." Proceedings of the 26th USENIX Security Symposium. 2017.

[12] Martín Abadi et al. "Control-flow integrity." Proceedings of the 12th ACM conference on Computer and communications security. ACM, 2005

[13] Nathan Burow et al. "Control-flow integrity: Precision, security, and performance." *ACM Computing Surveys (CSUR)* 50.1 (2017): 16.

[14] Shacham, H. et al. "Return-oriented programming: Exploits without code injection." Black Hat USA Briefings, August 2008.

[15] Vishwath Mohan and Kevin W. Hamlen. "Frankenstein: Stitching Malware from Benign Binaries." USENIX Workshop on Offensive Technologies (WOOT 12012): 77-84.

[16] Kyriakos Ispoglou, and Mathias Payer. "malWASH: Washing Malware to Evade Dynamic Analysis." WOOT. 2016.

[17] Giorgos Poulios, Christoforos Ntantogian, and Christos Xenakis. "Ropinjector: Using return oriented programming for polymorphism and antivirus evasion."Blackhat USA (2015).

[18] Giorgos Poulios, “Advanced Antiviurs Evasion Techniques”, Master Thesis, Department of Digital Systems, University of Piraeus, 2015.