How to kill symbolic deobfuscation for free (or: unleashing the potential of path-oriented protections)

Mathilde Ollivier, Sébastien Bardin, Richard Bonichon, Jean-Yves Marion

► To cite this version:

Mathilde Ollivier, Sébastien Bardin, Richard Bonichon, Jean-Yves Marion. How to kill symbolic deobfuscation for free (or: unleashing the potential of path-oriented protections). ACSAC '19: 2019 Annual Computer Security Applications Conference, Dec 2019, San Juan, Puerto Rico, United States. pp.177-189, 10.1145/3359789.3359812 . hal-02564103

HAL Id: hal-02564103

<https://hal.univ-lorraine.fr/hal-02564103>

Submitted on 13 May 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

How to Kill Symbolic Deobfuscation for Free (or: Unleashing the Potential of Path-Oriented Protections)

Mathilde Ollivier
CEA, LIST,
Paris-Saclay, France
mathilde.ollivier2@cea.fr

Richard Bonichon
CEA, LIST,
Paris-Saclay, France
richard.bonichon@cea.fr

Sébastien Bardin
CEA, LIST,
Paris-Saclay, France
sebastien.bardin@cea.fr

Jean-Yves Marion
Université de Lorraine, CNRS, LORIA
Nancy, France
Jean-Yves.Marion@loria.fr

ABSTRACT

Code obfuscation is a major tool for protecting software intellectual property from attacks such as reverse engineering or code tampering. Yet, recently proposed (automated) attacks based on Dynamic Symbolic Execution (DSE) shows very promising results, hence threatening software integrity. Current defenses are not fully satisfactory, being either not efficient against symbolic reasoning, or affecting runtime performance too much, or being too easy to spot. We present and study a new class of anti-DSE protections coined as path-oriented protections targeting the weakest spot of DSE, namely path exploration. We propose a lightweight, efficient, resistant and analytically proved class of obfuscation algorithms designed to hinder DSE-based attacks. Extensive evaluation demonstrates that these approaches critically counter symbolic deobfuscation while yielding only a very slight overhead.

CCS CONCEPTS

• **Security and privacy** → **Software reverse engineering**; *Logic and verification*; Malware and its mitigation; • **Software and its engineering** → *Formal methods*.

KEYWORDS

Reverse Engineering; Code Protection; Obfuscation

ACM Reference Format:

Mathilde Ollivier, Sébastien Bardin, Richard Bonichon, and Jean-Yves Marion. 2019. How to Kill Symbolic Deobfuscation for Free (or: Unleashing the Potential of Path-Oriented Protections). In *2019 Annual Computer Security Applications Conference (ACSAC '19)*, December 9–13, 2019, San Juan, PR, USA. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3359789.3359812>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ACSAC '19, December 9–13, 2019, San Juan, PR, USA

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-7628-0/19/12...\$15.00

<https://doi.org/10.1145/3359789.3359812>

1 INTRODUCTION

Context. Reverse engineering and code tampering are widely used to extract proprietary assets (e.g., algorithms or cryptographic keys) or bypass security checks from software. *Code protection* techniques precisely seek to prevent, or at least make difficult, such *man-at-the-end* attacks, where the attacker has total control of the environment running the software under attack. Obfuscation [21, 22] aims at hiding a program's behavior by transforming its executable code in such a way that the behavior is conserved but the program becomes much harder to understand.

Even though obfuscation techniques are quite resilient against basic automatic reverse engineering (including static attacks, e.g. disassembly, and dynamic attacks, e.g. monitoring), code analysis improves quickly [39]. Attacks based on *Dynamic Symbolic Execution* (DSE, a.k.a. *concolic execution*) [18, 30, 40] use logical formulas to represent input constraints along an execution path, and then automatically solve these constraints to discover new execution paths. DSE appears to be very efficient against existing obfuscations [5, 8, 13, 24, 37, 51], combining the best of dynamic and semantic analysis.

Problem. *The current state of symbolic deobfuscation is actually pretty unclear.* Dedicated protections have been proposed, mainly based on hard-to-solve predicates, like Mixed Boolean Arithmetic formulas (MBA) [52] or cryptographic hash functions [42]. Yet the effect of complexified constraints on automatic solvers is hard to predict [6], cryptographic hash functions may induce significant overhead and are amenable to key extraction attacks (possibly by DSE). On the other hand, DSE has been fruitfully applied on malware and legit codes protected by state-of-the-art tools and methods, including virtualization, self-modification, hashing or MBA [8, 37, 51]. A recent systematic experimental evaluation of symbolic deobfuscation [5] shows that most standard obfuscation techniques do not seriously impact DSE. Only nested virtualization seems to provide a good protection, assuming the defender is ready to pay a high cost in terms of runtime and code size [37].

Goals and Challenges. We want to propose a new class of dedicated anti-DSE obfuscation techniques to render automated attacks based on symbolic execution inefficient. These techniques should be *strong* – making DSE intractable in practice, and *lightweight* – with very low overhead in both code size and runtime performance. While most anti-DSE defenses try to break the symbolic reasoning

part of DSE (constraint solver), we instead target its real weak spot, namely path exploration. Banescu et al. [5] present one such specific obfuscation scheme but with a large space overhead and no experimental evaluation. We aim at proposing a general framework to understand such obfuscations and to define new schemes *both strong and lightweight*.

Contribution. We study *path-oriented* protections, a class of protections seeking to hinder DSE by substantially increasing the number of feasible paths within a program.

- We detail a formal framework describing *path-oriented* protections (Sec. 4). We characterize their desirable properties — namely *tractability*, *strength*, and the key criterion of *single value path* (SVP). The framework is *predictive*, in the sense that our classification is confirmed by experimental evaluation (Sec. 8), allowing both to shed new light on the few existing path-oriented protections and to provide guidelines to design new ones. In particular, no existing protection [5] achieves both tractability and optimal strength (SVP). As a remedy, we propose *the first two obfuscation schemes* achieving both *tractability and optimal strength* (Sec. 5).

- We highlight the importance of the *anchorage policy*, i.e. the way to choose where to insert protection in the code, in terms of protection efficiency and robustness. Especially, we identify a way to achieve *optimal composition* of path-oriented protections (Sec. 6.1), and to completely prevent taint-based and slic-based attacks (two powerful code-level attacks against obfuscation), coined as *resistance by design* (Sec. 6.2).

- We conduct extensive experiments (Sec. 8.3) with two different attack scenarios — *exhaustive path coverage* (Sec. 8.3) and *secret finding*. Results confirm that path-oriented protections are much stronger against DSE attacks than standard protections (including nested virtualization) for only a slight overhead. Moreover, while existing techniques [5] can still be weak in some scenarios (e.g., secret finding), our *new optimal schemes cripple symbolic deobfuscation at essentially no cost in any setting*. Finally, experiments against slice, pattern-matching and taint attacks confirm the quality of our robust-by-design mechanism.

As a practical outcome, we propose a new *hardened deobfuscation benchmark* (Sec. 9), currently out-of-reach of symbolic engines, in order to extend existing obfuscation benchmarks [1, 5, 37].

Discussion. We study a powerful class of protections against symbolic deobfuscation, based on a careful analysis of DSE — we target its weakest point (path exploration) when other dedicated methods usually aim at its strongest point (constraint solving and ever-evolving SMT solvers). We propose a predictive framework allowing to understand these protections, as well as several concrete protections impacting DSE more than three levels of virtualization at essentially no cost. We expect them to be also efficient against other semantic attacks [10, 31] (cf. Sec. 10). From a methodological point of view, this work extends recent attempts at rigorous evaluation of obfuscation methods. We provide both an analytical evaluation, as Bruni et al. [15] for anti-abstract model checking, and a refinement of the experimental setup initiated by Banescu et al. [5].

2 MOTIVATION

2.1 Attacker model

Scenario. We consider man-at-the-end scenarios where the attacker has full access to a potentially protected code under attack. The attacker only has the executable and no access to the source code, is skilled in program analysis but with limited resources (typically: motivated by economic gains [19], short term attack scenarios such as VOD cracking or video games).

As a consequence, this attacker has access to automated state-of-the-art off-the-shelf tools (DSE, etc.), can try to attack protections (tainting [40], slicing [44], patterns) and craft attacks by combining those tools. But our attacker will not invest in crafting dedicated tools going beyond state-of-the-art. Basically, our goal is to delay the attack enough so that the attacker stops because of the cost. *We consider that if the attacker has to craft a dedicated tool beyond state-of-the-art, then the defender has won.*

Scope. We focus on Symbolic Execution and other trace-based semantic attacks as they have proven to be useful automated techniques in recent attacks. We thus aim to remove them from the attacker's toolbox to increase the attack's costs. Typical DSE-based attacks include finding rare behaviors (triggers [13], secrets, etc.) of the whole program or local exhaustive exploration (proofs [8], simplifications [37]). Such attacks can be abstracted by the two following goals: (1) *Secret Finding*; (2) *Exhaustive Path Exploration*.

Caveat. *Part of our experimental evaluations uses source codes, as state-of-the-art source-level DSE tools are much more efficient than binary-level ones. Our experimental conditions actually favors the attacker more, and as a result they show that our approach is all the more effective.*

```
int check_char_0(char chr){
    char ch = chr;
    ch ^= 97;
    return (ch == 31);
}

/* ... 9 other checks ... */

int check_char_10(char chr){ /* ... */ }

int check(char* buf) {
    int retval = 1;
    retval *= check_char_0(buf[0]);
    /* ... check buf[1] to buf[9] ... */
    retval *= check_char_10(buf[10]);
    return retval;
}

int main(int argc, char** argv) {
    char* buf = argv[1];
    if (check(buf)) puts("win");
    else puts("lose");
}
```

Figure 1: Manticore crackme code structure

2.2 Motivating example

Let us illustrate anti-symbolic path-oriented protections on a toy crackme program¹. Fig. 1 displays a skeleton of its source code.

¹<https://github.com/trailofbits/manticore>

main calls check to verify each character of the 11 bytes input. It then outputs "win" for a correct guess, "lose" otherwise. Each sub-function `check_char_ii ∈ [0,10]` hides a secret character value behind bitwise transformations, like xor or shift. *Such a challenge can be easily solved, completely automatically, by symbolic execution tools. KLEE [17] needs 0.03s (on C code) and BINSEC [26] 0.3s (on binary code) to both find a winning input and explore all paths.*

Standard protections. Let us now protect the program with standard obfuscations to measure their impact on symbolic deobfuscation. We will rely on Tigress [23], a widely used tool for systematic evaluation of deobfuscation methods [5, 8, 37], to apply (nested) virtualization, a most effective obfuscation [5]. Yet, Table 1 clearly shows that virtualization does not prevent KLEE from finding the winning output, though it can thwart path exploration – but with a high runtime overhead (40×).

The case for (new) path-oriented protections. To defend against symbolic attackers, we thus need better anti-DSE obfuscation: *path-oriented protections*. Such protections aim at exponentially increasing the number of paths that a DSE-based deobfuscation tool, like KLEE, must explore. Two such protections are SPLIT and FOR, illustrated in Fig. 2 on function `check_char_0` of the example.

FOR	SPLIT
<pre>int func(char chr){ char ch = 0; for (int i=0; i<chr; i++) ch ^= 97; return (ch == 31); }</pre>	<pre>int func(char chr, ch1, ch2) { // new input ch1 and ch2 char garb = 0 // junk char ch = chr; if (ch1 > 60) garb++; else garb--; if (ch2 > 20) garb++; else garb--; ch ^= 97; return (ch == 31); }</pre>

Figure 2: Unoptimized obfuscation of `check_char_0`

For the sake of simplicity, the protections are implemented in a naive form, sensitive to slicing or compiler optimizations. Robustness is discussed afterwards. In a nutshell, SPLIT – an instance of RANGE DIVIDER [5] – adds a number k of conditional statements depending on new fresh inputs, increasing the number of paths to explore by a factor of 2^k . Also, in this implementation we use a junk variable `garb` and two additional inputs `ch1` and `ch2` unrelated to the original code. The novel obfuscation FOR (Sec. 5) adds k loops whose upper bound depends on distinct input bytes and which recompute a value that will be used later, expanding the number of paths to explore by a factor of $2^{8 \cdot k}$ – assuming a 8-bit `char` type. This implementation does not introduce any junk variable nor additional input. In both cases, the obfuscated code relies on the input, forcing DSE to explore *a priori* all paths. Table 1 summarizes the performance of SPLIT and FOR. Both SPLIT and FOR do not induce any overhead, SPLIT is highly efficient (timeout) against coverage but not against secret finding, while FOR is highly efficient for both. FOR ($k = 2$) performs already better than SPLIT ($k = 19$) and further experiments (Sec. 8) shows FOR to be a much more effective path protection than SPLIT.

Question: How to distinguish *a priori* between mildly effective and very strong path-oriented protections?

Note that gcc -Ofast is removes this simple SPLIT, as it is not related to the output (slicing attack). A basic FOR resists this attack, but clang -Ofast is able to remove it by an analysis akin to a *pattern*

Table 1: DSE Attack on the Crackme Example (KLEE)

	Obfuscation type	Slowdown Symbolic Execution		Over- head runtime
		Coverage	Secret	
Standard	Virt	××	××	×1.1
	Virt ×2	×	××	×1.3
	Virt ×3	✓	×	×40
Path-Oriented	SPLIT [5]	$k = 11$	××	×1.0
		$k = 15$	✓	×1.0
		$k = 19$	✓	×1.0
	FOR (new)	$k = 1$	×	×1.0
		$k = 2$	✓	×1.0
		$k = 3$	✓	×1.0

×× $t \leq 1s$ × $30s < t < 5min$ ✓ time out ($\geq 1h30$)

Unobfuscated case: KLEE succeeds in 0.03s

attack. However, a slightly modified FOR (Fig. 3) overcomes such optimizations.

```
int func(char chr) {
    int ch = 0; // int prevents char overflows
    for (int i=0; i<(int)chr; i++) {
        if (i % 2 == 0) ch += 3;
        if (i % 2 != 0) ch--;
    }
    if (i % 2 != 0) ch -= 2; // adjust for odd values
    ch ^= 97;
    return (ch == 31);
}
```

Figure 3: Enhanced FOR – `check_char_0`

Question: How to protect path-oriented protections against code analysis-based attacks (slicing, tainting, patterns)?

The goal of this paper is to define, analyze and explore in a systematic way the potential of path-oriented transformations as anti-DSE protections. We define a *predictive* framework (Sec. 4) and propose several new *concrete* protections (Sec. 5). In particular, our framework allows to precisely explain why FOR is experimentally better than SPLIT. We also discuss how path-oriented protections can be made resistant to several types of attacks (Sec. 6 and 7).

3 BACKGROUND

Obfuscation. Obfuscation [22] aims at hiding a program's behavior or protecting proprietary information such as algorithms or cryptographic keys by transforming the program to protect $\mathcal{P}$ into a program $\mathcal{P}'$ such that $\mathcal{P}'$ and $\mathcal{P}$ are semantically equivalent, $\mathcal{P}'$ is roughly as efficient as $\mathcal{P}$ and $\mathcal{P}'$ is *harder to understand*. While it is still unknown whether applicable theoretical criteria of obfuscation exist [7], practical obfuscation techniques and tools do.

Let us touch briefly on three such important techniques. *Mixed Boolean-Arithmetic* [29, 52] transforms an arithmetic and/or Boolean equation into another using a combination of Boolean and arithmetic operands with the goal to be more complex to understand and more difficult to solve by SMT solvers [9, 46]. *Virtualization* and *Flattening* [47] transform the control flow into an interpreter loop dispatching every instruction. Virtualization even adds a virtual machine interpreter for a custom bytecode program encoding the original program semantic. Consequently, the visible control flow of the protected program is very far from the original control flow.

Virtualization can be *nested*, encoding the virtual machine itself into another virtual machine. *Self-modifying code* and *Packing* insert instructions that dynamically modify the flow of executed instructions. These techniques seriously damage static analyses by hiding the real instructions. However, extracting the hidden code can be done by dynamic approaches [28, 32], including DSE [51].

Dynamic Symbolic Execution (DSE). Symbolic execution [18] simulates the execution of a program along its paths, systematically generating inputs for each new discovered branch condition. This exploration process consider inputs as *symbolic variables* whose value is not fixed. DSE follows a path and each time a conditional statement involving the input is encountered, it adds a constraint to the *symbolic value* related to this input. Solving the constraints automatically (typically with off-the-shelf SMT solvers [46]) then allows to generate *new input values leading to new paths*, progressively covering all paths of the program – up to a user-defined bound. The technique has seen strong renewed interest in the last decade to become a prominent bug finding technique [18, 20, 30].

When the symbolic engine cannot perfectly handle some constructs of the underlying programming language – like system calls or self-modification – the symbolic reasoning is interleaved with a *dynamic analysis* allowing meaningful (and fruitful) approximations – *Dynamic Symbolic Execution* [30]. Typically, (concrete) runtime values are used to complete missing part of path constraints that are then fed to the solver through *concretization* [25]. This feature makes the approach especially robust against complicated constructs found in obfuscated binary codes, typically packing or self-modification, making DSE a strong candidate for automated deobfuscation – *symbolic deobfuscation*: it is as robust as dynamic analysis, with the additional ability to infer *trigger-based conditions*.

4 A FRAMEWORK FOR PATH-ORIENTED PROTECTIONS

This section presents a framework to evaluate *path-oriented* obfuscations, i.e. protections aiming at hindering symbolic deobfuscation by taking advantage of path explosion.

4.1 Basic definitions

Figure 4: Path-Oriented Obfuscation Framework

This paper deals with a specific kind of protections targeting DSE: *path-oriented* protections. Transforming a program $\mathcal{P}$ into $\mathcal{P}'$ using *path-oriented* protections ensures that $\mathcal{P}'$ is functionally equivalent to $\mathcal{P}$ and aims to guarantee $\#\Pi' \gg \#\Pi$, meaning the number of paths in $\mathcal{P}'$ is much greater than the ones in $\mathcal{P}$.

The most basic path-oriented protection consists in one *forking point* inserted in the original code of $\mathcal{P}$.

Definition 1 (Forking Point) A forking point $\mathcal{F}$ is a location in the code that creates at most γ new paths. $\mathcal{F}$ is defined by: an address a , a variable x and a capacity γ . It is written $\mathcal{F}(a, x, \gamma)$.

To illustrate this definition, see the snippet of SPLIT in Figure 2: both if-statements define each a forking point of capacity 2 based on the variable `ch1` and `ch2` respectively.

Now, to obtain a complete path-oriented obfuscation $\mathcal{P}'$ of a program $\mathcal{P}$, we need to insert n forking points throughout the code of $\mathcal{P}$, hence the notion of *obfuscation scheme* (Fig. 4).

Definition 2 (Obfuscation scheme) A (path-oriented protection) obfuscation scheme is a function $f(\mathcal{P}, n)$ that, for every program $\mathcal{P}$, inserts n forking points in $\mathcal{P}$. It comprises a set of forking points and an anchorage policy, i.e. the placement method of the forking points.

4.2 Desirable obfuscation scheme properties

An ideal obfuscation scheme is both strong (high cost for the attacker) and cheap (low cost for the defender). Let us define these properties more precisely.

The **strength** of an obfuscation scheme is intuitively the expected increase of the search space for the attacker. Given an obfuscation scheme f , it is defined as $\Gamma_f(\mathcal{P}, n) = \#\Pi_{f(\mathcal{P}, n)}$, for a program $\mathcal{P}$ and n forking points to insert.

The **cost** is intuitively the *maximal* runtime overhead the defender should worry about. Given an obfuscation scheme f , cost is defined by $\Omega_f(\mathcal{P}, n)$, as the *maximum* trace size of the obfuscated program $f(\mathcal{P}, n)$. Formally, $\Omega_f(\mathcal{P}, n) = \max_i \{|\tau'_i|\}$ where $\{\tau'_i\}$ is the set of execution traces of $f(\mathcal{P}, n)$ and $|\tau'_i|$ is the size of the trace.

We seek strong tractable obfuscations, i.e., yielding enough added paths to get a substantial slowdown, with a low runtime overhead.

Definition 3 (Strong scheme) An obfuscation scheme f is strong if for any program $\mathcal{P}$, we have $\Gamma_f(\mathcal{P}, n) \geq 2^{O(n)} \cdot \#\Pi_{\mathcal{P}}$, where $\Pi_{\mathcal{P}}$ is the set of paths of $\mathcal{P}$. Putting things quickly, it means that the number of paths to explore is multiplied by 2^n .

Definition 4 (Tractable scheme) An obfuscation scheme f is tractable if for any program $\mathcal{P}$, $\Omega_f(\mathcal{P}, n) \leq \max_i \{|\tau_i|\} + O(n)$, where $\max_i \{|\tau_i|\}$ is the size of the longest trace of $\mathcal{P}$. In other words, it is tractable only if the overhead runtime is linear on n .

Combining schemes. Scheme composition preserves tractability (the definition involves an upper bound) but not necessarily strength (the definition involves a lower bound). Hence, we need *optimal composition rules* (Sec. 6.1).

4.3 Building stronger schemes

Strong path-oriented protections, can rely on composition but we saw in Sec. 4.2 that it is not straightforward. But since path-oriented protections first lean on forking execution into many paths, we should also investigate whether some forking points are better than others. The best case is to insert k forking points $(\mathcal{F}(a_i, x_i, \gamma_i))_i$ such that it would ensure that each path created by a forking point $(\mathcal{F}(a_i, x_i, \gamma_i))_i$ corresponds to only one possible value of the variable x_i . This leads us to define this type of forking point as a *Single Value Path* (SVP) protection.

Definition 5 (Single Value Path) A forking point based on variable x is Single Value Path (SVP) if and only if x has only one possible value in each path created by the protection.

A SVP forking point creates a new path for each possible value of variable x (e.g., i.e., 2^{32} new paths are created for an unconstrained C int variable). For example, the FOR obfuscation shown in Figure 2 produces 2^{32} paths, that should be a priori explored since it depends on an input variable. SVP forking points is key to ensure that DSE will need to enumerate all possible input values of the program under analysis (thus *boiling down to brute forcing*) – see Sec. 6.

5 CONCRETE PATH-ORIENTED PROTECTIONS

This section reviews existing path-oriented protection schemes within the framework of Sec. 4, but also details new such schemes achieving both tractability and optimal strength (SVP).

```
int main (int argc, char** argv){
  char* input = argv[1];
  char chr = *input; // inserted by obfuscation
  switch (chr) {      // inserted by obfuscation
    case 1: ... // original code
    break;
    case 2: // obfuscated version of case 1
    break;
    ...
    default: // other obfuscated version of case 1
    break;
  }
  return (*input >= 100);
}
```

Figure 5: RANGE DIVIDER obfuscation

RANGE DIVIDER [5]. RANGE DIVIDER is an anti-symbolic execution obfuscation proposed by Banescu et al.. Branch conditions are inserted in basic blocks to divide the input value range into multiple sets. The code inside each branch of the conditional statement is an obfuscated version of the original code. We distinguish two cases, depending on whether the branch condition uses a `switch` or a `if` statement. *In the remaining part of this paper, SPLIT will denote the RANGE DIVIDER obfuscation with if statement, and RANGE DIVIDER the RANGE DIVIDER obfuscation with switch statement.*

The RANGE DIVIDER (`switch`) scheme introduces an exhaustive `switch` statement over all possible values of a given variable – see example in Fig. 5, thus yielding 2^S extra-paths, with S the bit size of the variable. This scheme enjoys the SVP property as in each branch of the `switch` the target variable can have only one value, and it is also tractable in time provided the `switch` is efficiently compiled into a binary search tree or a jump table, as usual. Yet, while not pointed out by Banescu et al., *this scheme is not tractable in space* (code size) as it leads to *huge* amount of code duplication – the byte case may be fine, but not above.

SPLIT [5]. This transformation (Fig. 6) is similar to RANGE DIVIDER, but the control-flow is split by a condition triggered by a variable. This protection is tractable in both time (only one additional condition check per forking point) and space (only one block duplication per forking point). Yet, the protection is not SVP.

```
int main (int argc, char** argv){
  char* input = argv[1];
  char chr = *input; // inserted by obfuscation
  if (chr < 30) {      // inserted by obfuscation
    ... // original code O
  }
  else ... // obfuscated version of O
}

return (*input >= 100);
}
```

Figure 6: SPLIT obfuscation

FOR (new). The FOR scheme (Fig. 2) replaces assignments $ch := chr$ by loops $ch = 0$; `for` ($i = 0$; $i \leq chr$; $i++$) $ch++$; where chr is an input-dependent variable. Intuitively, such `for` loops can be unrolled n times, for any value n that chr can take at runtime. Hence, a loop controlled by a variable defined over a bit size S generates up to 2^S extra-paths, with additional path length of 2^S . While the achieved protection is excellent, it is *intractable* when $S = 32$ or $S = 64$. **The trick** is to restrict this scheme to byte-size variables, and then chain such forking points on each byte of the variable of interest. Indeed, FOR over a byte-size variable generates up to 2^8 additional paths with an additional path length at most of 2^8 . Chaining k such FOR loops leads up to 2^{8k} extra-paths with an extra-length of only $k \cdot 2^8$, *keeping strength while making runtime overhead reasonable*. (More precisely with a constant time overhead wrt inputs.)

WRITE (new). The WRITE obfuscation adds self-modification operations to the code. It replaces an assignment $a := input$ with a non input-dependent operation $a := k$ (with k an arbitrary constant value) and replaces at runtime this instruction by $a := i$ where i is the runtime value of `input` (self-modification). This is illustrated in Fig. 7, where the offset move at label L1 actually rewrites the constant 0 at L2 to the value contained at the address of the input.

```
L: mov [a], [input]  ⇒  L1: mov L2+off,[input]
                      L2: mov [a], 0
```

Figure 7: WRITE obfuscation

Symbolic execution engines are not likely to relate a and `input`, thus thinking that a is constant across all executions. If the dynamic part of the engine spots that a may have different values, it will iterate over every possible values of `input`, creating new paths each time. The scheme is SVP, and its overhead is negligible (2 additional instructions, independent of the bit size of the targeted variable as long as it can be handled natively by the underlying architecture). WRITE has yet two drawbacks: it can be spotted by a dynamic analysis and needs the section containing the code to be writable.

Table 2: Classification of obfuscation schemes

		Tractable		SVP	Stealth
		Time	Space		
RANGE DIVIDER [5]	switch	✓	✗	✓	✗
SPLIT [5]	if	✓	✓	✗	✓
FOR	word	✗	✓	✓	✓
	byte	✓	✓	✓	✓
WRITE		✓	✓	✓	~

Summary. The properties of every scheme presented so far are summarized in Table 2 – stealth is discussed in Sec. 7.2. *Obfuscation*

schemes from the literature are not fully satisfactory: RANGE DIVIDER is space expensive and easy to spot, SPLIT is not strong enough (not SVP). On the other hand, *the new schemes* FOR (at byte-level) and WRITE *are both strong (they satisfy SVP) and tractable, making them perfect anti-DSE protections*.

As a consequence, we suggest using variations of FOR as the main protection layer, with WRITE deployed only when self-modification and unpacking are already used (so that the scheme remains hidden). RANGE DIVIDER can be used *occasionally* but only on byte variables to mitigate space explosion. SPLIT can add further diversity and code complexity, but it should not be the main defense line. All these protections must be inserted in a *resistant-by-design* manner (Sec. 6.2) together with *diversity of implementation* (Sec. 7.1).

6 ANCHORAGE POLICY

We need to ensure that inserting path-oriented protections into a program gives real protection against DSE and will not be circumvented by attackers.

6.1 Optimal composition

We show how to combine the forking points in order to obtain *strong* obfuscation schemes. The issue with obfuscation scheme combination is that some forking points could hinder the efficiency of other forking points – imagine a `if (x ≥ 100)` split followed by a `if (x ≤ 10)` split: we will end up with 3 *feasible* paths rather than the expected $2 \times 2 = 4$, as one of the path is infeasible ($x > 100 \wedge x \leq 10$). Intuitively, we would like the forking points to be *independent* from each other, in the sense that their efficiency combine perfectly

Definition 6 (Independence) *Let us consider a program $\mathcal{P}$ and σ a path of $\mathcal{P}$. We obfuscate this program alternatively with two forking points $\mathcal{F}_1$ and $\mathcal{F}_2$ such that σ encounters both forking points. This results in three obfuscated programs: $\mathcal{P}_1$, $\mathcal{P}_2$ and $\mathcal{P}_{1,2}$. We note $\# \sigma_1$ (resp. $\# \sigma_2$) the set of feasible paths created from σ when encountering only $\mathcal{F}_1$ in $\mathcal{P}_1$ (resp. $\mathcal{F}_2$ in $\mathcal{P}_2$) and $\# \sigma_{1,2}$ the set of feasible paths created from σ when encountering both $\mathcal{F}_1$ and $\mathcal{F}_2$ in $\mathcal{P}_{1,2}$. $\mathcal{F}_1$ and $\mathcal{F}_2$ are independent over a program $\mathcal{P}$ if for all path σ passing through $\mathcal{F}_1$ and $\mathcal{F}_2$: $\# \sigma_{1,2} = \# \sigma_1 \times \# \sigma_2$*

An easy way to obtain forking point independence is to consider forking points built on independent variables – variables are *independent* if their values are not computed from the same input values. Actually, if independent forking points are well placed in the program, path-oriented protections ensure an exponential increase in the number of paths (cf. Theorem 1, proof in Appendix A).

Theorem 1 (Optimal Composition) *Suppose that $\mathcal{P}'$ is obtained by obfuscating the program $\mathcal{P}$. If every original path of $\mathcal{P}$ goes through at least k independent forking points of $\mathcal{P}'$ inserting at least θ feasible paths, then $\# \Pi_{\mathcal{P}'} \geq \# \Pi_{\mathcal{P}} \cdot \theta^k$*

By choosing enough independent SVP forking points (one for each input variable), we can even ensure that DSE will have to enumerate over all possible input values of the program under analysis, hence performing as bad as mere *brute forcing*.

Implementation. Ensuring that each path will go through at least k forking points can be achieved by carefully selecting the points in the code where the forking points are inserted: a control flow graph analysis provides information about where and how many forking points are needed to cover all paths. The easiest way to

impact all paths at once is to select points in the code that are not influenced by any conditional statement. Dataflow analysis can be used further in order to ensure that the selected variables do not share dependencies with the same input (independent variables).

6.2 Resistance-by-design to taint and slice

Taint analysis [40] and (backward) slicing [44] are two advanced code simplification methods built on the notion of *data flow relations* through a program. These data flow relations can be defined as *Definition-Use* (Def-Use) chains – as used in compilers. Data are *defined* when variables are assigned values or declared, and *used* in expressions. Taint (resp. Slice) uses Def-Use chains to replace input-independent by its constant effect (resp. remove code not impacting the output). If there exists a Def-Use chain linking data x to data y , we write: $x \rightsquigarrow y$.

Definition 7 (Relevant Variable) x is **relevant** if there exists at least two Def-Use chains such that $\text{input} \rightsquigarrow x$ and $x \rightsquigarrow \text{output}$.

A *sound* taint analysis (resp. slice analysis) marks *at least* all variables (x, a) such that $\text{input} \rightsquigarrow (x, a)$ (resp. $(x, a) \rightsquigarrow \text{output}$). Unmarked variables are then safely removed (slicing) or set to their constant value (tainting). Thus, in order to resist by design such attacks, protections must rely on code that will be marked by both slicing and tainting.

Here, we refine the definition of a forking point $\mathcal{F}$: it can be viewed as two parts, a guard $\mathcal{G}$ – the condition – and an action $\mathcal{A}$ – the code in the statement. We denote by $\text{Var}(\mathcal{F})$ the set of variables in $\mathcal{G}$ and $\mathcal{A}$. We say that $\mathcal{F}$ is *built upon relevant variables* if all variables in $\text{Var}(\mathcal{F})$ are relevant.

Theorem 2 (Resistance by design) *Let us consider a program $\mathcal{P}$ and a forking point $\mathcal{F}$. Assuming $\mathcal{F}$ is built upon relevant variables, then $\mathcal{F}$ is slice and taint resistant.*

Note. Our method protects against sound analyses (i.e., overapproximations). The proposed technique does not aim to defend against an underapproximated analysis, which may provide relevant analysis results by chance (with a low probability). However, an underapproximated analysis may yield undue code simplifications.

Implementation. Relevant variables can be identified by modifying standard compiler algorithms computing *possible* Def-Use chains in order to compute *real* Def-Use chains – technically, going from a *may analysis* to a *must analysis*. A more original solution observes at runtime a set of real Def-Use chains and deduces accordingly a set of relevant variables. This method does not require any advanced static analysis, only basic dynamic tracing features.

7 THREATS

In this section we discuss possible threats to path-oriented protections and propose adequate mitigations. Indeed, when weaving the forking points within the code of a program, we need to ensure that they are hard to discover or remove. Three main attacks seem effective against path-oriented protections: (1) taint analysis, (2) backward slicing, (3) and pattern attacks. We showed how path-oriented protections can be made *resistant by-design* to Taint and Slice in Sec. 6.2. We now discuss pattern attacks, as well as stealth issues and the unfriendly case of programs with small input space.

7.1 Pattern attacks

Pattern attacks search for specific patterns in the code of a program to identify, and remove, known obfuscations. This kind of analysis assumes more or less similar constructions across all implementations of an obfuscation scheme. A general defense against pattern attacks is *diversity*. It works well in the case of path-oriented protections: on the one hand the schemes we provide can be implemented in many ways, and on the other hand our framework provides guidelines to design new schemes – altogether, it should be enough to defeat pattern attacks. Regarding diversity of implementations, the standard FOR forking point can be for example replaced by a **while** loop, (mutually) recursive function(s), the loop body can be masked through MBAs, etc. These variants can be combined as in Fig. 8, and we can imagine many other variations.

The other schemes as well can be implemented in many ways, and we could also think of ROP-based encoding [41] or other diversification techniques. Altogether, it should provide a powerful enough mitigation against pattern attacks.

```

① for (int i = 0; i++; i < input) a++;
② for (int i = 0; i++; i < input)
  a = (a ^ 1) + 2 * (a & 1);
③ int i = 0;
  while (i < input) { i++; a++; }
④ int f(int x) {
  return (x <= 0 ? 0 : f(x - 1) + 1); }

a = f(input);
⑤ #define A // arbitrary value

int f(int x) {
  return x <= 0 ? 0 : A + g(x - 1); }

int g(int x) {
  return !x ? 1 - A : 2 - A + f(-x); }

a = f(input);

```

Figure 8: Several encodings of protection FOR

7.2 Stealth

In general, code protections are better when hard to identify, in order to prevent human-level attacks like stubbing parts of the code or designing targeted methods. Let us evaluate the stealthiness of path-oriented protections (summary in Table 2). SPLIT and FOR do not use rare operators or exotic control-flow structures, only some additional conditions and loops scattered through the program. Hence SPLIT and FOR are considered hard to detect on binary code, though FOR especially may be visible at source level. RANGE DIVIDER is easy to spot at source level: **switch** statements with hundreds of branches are indeed distinctive. Compilation makes it harder to find but the height of the produced binary search trees or the size of the generated jump table are easily perceptible. WRITE stands somewhere in between. It cannot be easily discovered statically, but is trivial to detect dynamically. However, since self-modification and unpacking are common in obfuscated codes, WRITE could well be mistaken for one of these more standard (and less damaging) protections.

7.3 Beware: programs with small input space

Resistance by design (Sec. 6.2) relies on *relevant variables*, so we only have limited room for forking points. In practice it should not be problematic as Sec. 8 shows that we already get very strong protection with only 3 input bytes – assuming a SVP scheme. Yet, for programs with very limited input space, we may need to add (*fake*) *crafted inputs* for the input space to become (apparently) larger – see SPLIT example in Fig. 2. In this case, our technique still ensures resistance against tainting attacks, but slicing attacks may now succeed. The defender must then rely on well-known (but imperfect) anti-slicing protections to blur code analysis through hard-to-reason-about constructs, such as pointer aliasing, arithmetic and bit-level identities, etc.

8 EXPERIMENTAL EVALUATION

The experiments below seek to answer four Research Questions²:

- RQ1** What is the impact of path-oriented protections on semantic attackers? Especially, we consider DSE attack and two different attacker goals: Path Exploration (**Goal 1**) and Secret Finding (**Goal 2**).
- RQ2** What is the cost of path-oriented protections for the defender in runtime overhead and code size increase?
- RQ3** What is the effectiveness of our resistance-by-design mechanism against taint and slice attacks?
- RQ4** What is the difference between standard protections, path-oriented protections and SVP protections?

8.1 Experimental setup

Tools. Our attacker mainly comprises the state-of-the-art source-level DSE tool KLEE (version 1.4.0.0 with LLVM 3.4, POSIX runtime and STP solver). KLEE is highly optimized [17] and works from source code, so it is *arguably the worst case off-the-shelf DSE-attacker we can face* [5]. We used all standard search heuristics (DFS, BFS, Non-Uniform Random Search) but report only about DFS, the best performer (see Appendix). We also used three *binary-level* DSE tools (BINSEC [26], TRITON [38], ANGR [43]), with similar results.

Regarding standard defense, we use Tigress [23], a freely available state-of-the-art obfuscator featuring many standard obfuscations and allowing to precisely control which ones are used – making Tigress a tool of choice for the systematic evaluation of deobfuscation methods [5, 8, 37].

Protections. We mainly consider tractable path-oriented protections and select both a new SVP scheme (FOR) and an existing non-SVP scheme (SPLIT), inserted in a robust-by-design way. We vary the number of forking points per path (parameter k).

We also consider standard protections: *Virtualization* (up to 3 levels), *arithmetic encoding* and *flattening* [48]. Previous work [5] has shown that nested virtualization is the sole standard protection useful against DSE. Our results confirm that, so we report only results about virtualization (*other results partly in Appendix*).

8.2 Datasets

We select small and medium programs for experiments as they represent the worst case for program protection. If path-oriented protections

² Download at <https://bit.ly/2wYSEDG>

can slow down DSE analysis substantially on smaller programs, then those protections can only give better results for larger programs.

Dataset #1. This synthetic dataset from Banescu *et al.*³ [5] offers a valuable diversity of functions and has already been used to assess resilience against DSE. It has 48 C programs (between 11 and 24 lines of code) including control-flow statements, integer arithmetic and system calls to `printf`. We exclude 2 programs because reaching full coverage took considerably longer than for the other 46 programs and blurred the overall results. Also, some programs have only a 1-byte input space, making them too easy to brute force (Goal 2). We turn them into equivalent 8-byte input programs with same number of paths – additional input are not used by the program, but latter protections will rely on them. The maximum time to obtain full coverage on the 46 programs with KLEE is 33s, mean time is 2.34s (Appendix B).

Dataset #2. The second dataset comprises 7 larger realistic programs, representative of real-life protection scenarios: 4 hash functions (City, Fast, Spooky, md5), 2 cryptographic encoding functions (AES, DES) and a snippet from the GRUB bootloader. Unobfuscated programs have between 101 and 934 LOCs: KLEE needs at most 33.31s to explore all paths, mean time is 8s (Appendix B).

8.3 Impact on Dynamic Symbolic Execution

Protocol. To assess the impact of protections against DSE, we consider the induced *slowdown* (time) of symbolic execution on an obfuscated program w.r.t. its original version. For more readable results, we only report whether DSE achieves its goal or times out.

For Path Exploration (**Goal 1**), we use programs from Datasets #1 and #2, add the protections and launch KLEE until it reports full coverage or times out – 3h for Dataset #1, or a 5,400× average slowdown, 24h for Dataset #2, or a 10,000× average slowdown.

For Secret finding (**Goal 2**), we modify the programs from both datasets into “secret finding” oriented code (e.g., *win / lose*) and set up KLEE to stop execution as soon as the secret is found. We take the whole Dataset #2, but restrict Dataset #1 to the 15 programs with 16-byte input space. We set smaller timeouts (1h for Dataset #1, 3h and 8h for Dataset #2) as the time to find the secret with KLEE on the original programs is substantially lower (0.3s average).

Table 3: Impact of obfuscations on DSE

Transformation (#TO/#Samples)	Dataset #1		Dataset #2		
	Goal 1 3h TO	Goal 2 1h TO	Goal 1 24h TO	Goal 2 3h TO	Goal 2 8h TO
Virt	0/46	0/15	0/7	0/7	0/7
Virt ×2	1/46	0/15	0/7	0/7	0/7
Virt ×3	5/46	2/15	1/7	0/7	0/7
SPLIT (k = 10)	1/46	0/15	0/7	0/7	0/7
SPLIT (k = 13)	4/46	0/15	1/7	1/7	0/7
SPLIT (k = 17)	18/46	2/15	3/7	2/7	1/7
FOR (k = 1)	2/46	0/15	0/7	0/7	0/7
FOR (k = 3)	30/46	8/15	3/7	2/7	1/7
FOR (k = 5)	46/46	15/15	7/7	7/7	7/7

Results & Observations. Table 3 shows the number of timeouts during symbolic execution for each obfuscation and goal. For example, KLEE is always able to cover all paths on Dataset #1 against

simple Virtualization (0/46 TO) – the protection is useless here, while it fails on $\approx 40\%$ of the programs with SPLIT ($k = 17$), and never succeeds with FOR ($k = 5$).

As expected, higher levels of protections (more virtualization layers or more forking points) result in better protection. Yet, results of Sec. 8.4 will show that while increasing forking points is cheap, increasing levels of virtualization is quickly prohibitive.

Virtualization is rather weak for both goals – only 3 levels of virtualization manage some protection. FOR performs very well for both goals: with $k = 3$ and Dataset #1, FOR induces a timeout for more than half the programs for both goals, which is significantly better than Virt×3. With $k = 5$, all programs timeout. In between, SPLIT is efficient for Goal 1 (even though it requires much higher k than FOR) but not for Goal 2 – see for example Dataset #1 and $k = 17$: 39% timeouts (18/46) for Goal 1, only 13% (2/15) for Goal 2.

Other (unreported) results. All standard protections from Tigris we used turns out to be ineffective against DSE – for example *Flattening* and *EncodeArithmetic* on Dataset#1 slows path exploration by a maximum factor of 10, which is far from critical. Search heuristics obviously do not make any difference in the case of Path Exploration (Goal 1). Still, DFS tends to perform slightly better than BFS and NURS against SPLIT in the case of Secret Finding (Goal 2). No other difference is visible. Experiments with three binary-level DSE engines supported by different SMT solvers (BINSEC [26] with Boolector[14], TRITON [38] and ANGR [43] with Z3[27]) are in line with those reported here.

Regarding WRITE, most state-of-the-art tools (KLEE, BINSEC, TRITON) do not support self-modification. ANGR has a specific option, however after testing we found out it did not consider the WRITE pattern as symbolic thus missing the majority of paths (Goal 1) and failing to recover the secret (Goal 2).

Conclusion. As already stated in the literature, standard protections such as nested virtualization are mostly inefficient against DSE attacks. Path-oriented protections are shown here to offer a stronger protection. Yet, care must be taken. Non-SVP path protections such as SPLIT do improve over nested virtualization (SPLIT with $k = 13$ is roughly equivalent to Virt ×3, with $k = 17$ it is clearly superior), but they provide only a weak-to-mild protection in the cases of Secret Finding (Goal 2) or large time outs. On the other hand, SVP protections (represented here by FOR) are able to discard all DSE attacks on our benchmarks for both Path Exploration and Secret Finding with only $k = 5$, demonstrating a protection power against DSE far above those of standard and non-SVP path protections.

To conclude, path-oriented protections are indeed a tool of choice against DSE attacks (**RQ1**), much stronger than standard ones (**RQ4**). In addition, SVP allows to predict the strength difference of these protections (**RQ4**), against Coverage or Secret Finding.

8.4 Impact on Runtime Performance

Protocol. We evaluate the cost of path-oriented protections by measuring the *runtime overhead* (RO) and the (binary-level) *code size increase* (CI) of an obfuscated program w.r.t. its original version. We consider also WRITE and two variants of FOR – its recursive encoding REC (Sec. 7.1) and the more robust P2 encoding (Sec. 8.5), as well as the untractable word-level For scheme (Sec. 5), coined WORD.

³<https://github.com/tum-i22/obfuscation-benchmarks>

Table 4: Impact of obfuscations on runtime performance

Transformation	Dataset #1		Dataset #2	
	RO	CI	RO	CI
Virt	×1.5	×1.5	×1.7	×1.4
Virt ×2	×15	×2.5	×5.1	×2.1
Virt ×3	×1.6 · 10 ³	×4	×362	×3.0
SPLIT (<i>k</i> = 10)	×1.2	×1.0	×1.0	×1.0
SPLIT (<i>k</i> = 13)	×1.2	×1.0	×1.0	×1.0
SPLIT (<i>k</i> = 50)	×1.5	×1.5	×1.1	×1.0
FOR (<i>k</i> = 1)	×1.0	×1.0	×1.0	×1.0
FOR (<i>k</i> = 3)	×1.1	×1.0	×1.0	×1.0
FOR (<i>k</i> = 5)	×1.3	×1.0	×1.1	×1.0
FOR (<i>k</i> = 50)	×1.5	×1.5	×1.2	×1.1
FOR (<i>k</i> = 5) P2	×1.3	×1.0	×1.1	×1.0
FOR (<i>k</i> = 5) REC	×3.0	×1.0	×2.7	×1.0
FOR (<i>k</i> = 1) WORD	×2.6 · 10 ³	×1.0	×2.1 · 10 ³	×1.0
WRITE (<i>k</i> = 5)	×1.0	×1.0	×1.0	×1.0
WRITE (<i>k</i> = 50)	×1.0	×2.1	×1.0	×1.2

Results & Observations. Results are shown in Table 4 as average values over all programs in the datasets. As expected, nested virtualization introduces a significant and prohibitive runtime overhead (three layers: ×1.6 · 10³ for Dataset #1 and ×362 for Dataset #2), and each new layer comes at a high price (from 1 to 2: between ×3 and ×10; from 2 to 3: between ×70 and ×100). Moreover, the code size is also increased, but in a more manageable way (still, at least ×3 for three layers). On the other hand, SPLIT, FOR and WRITE introduce only very low runtime overhead (at most ×1.3 on Dataset #1 and ×1.1 on Dataset #2). No significant code size increase is reported even for *k* = 50 for larger programs, however we predictably see some code size increase for small programs of dataset #1. Regarding variants of FOR, P2 does not show any overhead w.r.t. FOR, while the recursive encoding REC comes at a higher price. Finally, as predicted by our framework, WORD is intractable.

Conclusion. As expected, tractable path-oriented protections indeed yield only a very slight overhead, both in terms of time or code size (RQ2), and improving the level of protection (*k*) is rather cheap, while nested virtualization comes at a high price (RQ4). Coupled with results of Sec. 8.3, it turns out that path-oriented protections offer a much better anti-DSE protection than nested virtualization at a runtime cost several orders of magnitude lower. Also, the code size increase due to path-oriented protections seems compatible with strict memory requirements (e.g., embedded systems) where it is not the case for nested virtualization.

8.5 Robustness to taint and slice attacks

Protocol. We consider the clang & GCC compilers (many simplifications including slicing), the industrial-strength Frama-C static code analyzer (both its Taint and Slice plugins together with precise interprocedural range analysis) as well as TRITON (featuring tainting) and KLEE. We focus on the 8 programs from dataset #1 with 16-byte input space and all programs from dataset #2. We remove programs with 1-byte input space from dataset #1 because we added fake inputs that would obviously not resist analysis. This issue is further discussed in Sec. 7.3. We use the FOR scheme (*k*=3) weaved

Table 5: Robustness of path-oriented protections

Tool	Robust ?		
	P1 (basic)	P2 (obfuscated)	P3 (weak)
GCC -Ofast	✓	✓	✗
clang -Ofast	✗	✓	✗
Frama-C Slice	✓	✓	✗
Frama-C Taint	✓	✓	✓
TRITON (taint)	✓	✓	✓
KLEE	✓	✓	✓

✓: no protection simplified ✗: ≥ 1 protection simplified

into the code following our robust-by-design method (Sec. 6.2). Actually we consider 3 variants of the scheme: **P1**, **P2** and **P3**. P1 is the simple version of FOR presented in Fig. 8, P2 is a mildly obfuscated version (adds a if statement always evaluating to true in the loop – opaque predicate) and P3 naively relies on fake inputs, as Fig. 2’s SPLIT (a dangerous construction discussed in Sec. 7.3). A protection will be said to be *simplified* when the number of explored paths for full coverage is much lower than expected (DSE tools), no protection code is marked by the analysis tool (Frama-C) or running KLEE on the produced code does not show any difference (compilers).

Results & Observations. Results in Table 5 confirm expectations. No analyzer but clang is able to simplify our robust-by-design protections (P1 and P2), whereas the weaker P3 is broken by slicing (GCC, clang, Frama-C) but not by tainting – exactly as pointed out in Sec. 7.3. Interestingly, clang -Ofast simplifies scheme P1, *not due to slicing* (this is resistant by design), but thanks to some loop simplification more akin to a pattern attack, relying on finding an affine relation between variables and loop counters. The slightly obfuscated version P2 is immune to this particular attack.

Conclusion. Our robust-by-design method experimentally works as expected against taint and slice (RQ3). Yet, care must be taken to avoid pattern-like simplifications. Note that in a real scenario, the attacker must work on binary code, making static analysis much more complicated. Also, virtualization, unpacking or self-modification can be added to path-oriented protections to preclude static analysis.

9 APPLICATION: HARDENED BENCHMARK

We propose new benchmarks containing 4 programs from Banescu’s dataset and Sec. 8.2’s 6 real-world programs (GRUB excluded) to help advance the state of the art of symbolic deobfuscation.

Each program comes with two setups, Path Exploration and Secret Finding, obfuscated with both a path-oriented protection (FOR *k*=5, taint- and slice- resistant) and a virtualization layer against human and static attacks⁴. Table 6 shows the performance of KLEE, TRITON, BINSEC and ANGR (Secret Finding, 24h timeout). Hardened codes remain unsolved within the timeout, for every tool.

⁴Sources available at <https://bit.ly/2GNxNv9>

Table 6: Results on 10 hardened examples (secret finding)

	Unprotected (TO = 10 sec)	Virt ×1 (TO = 5 min)	Hardened – FOR (k=5) (TO = 24h)
KLEE	10/10	10/10	0/10 ✓
BINSEC	10/10	10/10	0/10 ✓
TRITON	10/10	10/10	0/10 ✓
ANGR	10/10	10/10	0/10 ✓

10 DISCUSSION

10.1 On the methodology

We discuss potential biases of our experimental evaluation.

Metrics. We add overhead metrics (runtime, code size) to the commonly used “DSE slowdown” measure [5, 37], giving a better account of the pros and cons of obfuscation methods.

Obfuscation techniques & tools. We consider the strongest standard obfuscation methods known against DSE, as identified in previous systematic studies [5, 37]. We restrict ourselves to their implementation in Tigress, a widely respected and freely available obfuscation tool considered state-of-the-art – studies including Tigress along packers and protected malware [8, 51] do not report serious deficiencies about its protections. Anyway, the evaluation of the path-oriented protections is independent of Tigress.

DSE engines. We use 4 symbolic execution engines (mostly KLEE, also: BINSEC, TRITON, ANGR) working on different program representations (C source, binary), with very similar final results. Moreover, KLEE is a highly respected tool, implementing advanced path pruning methods (*path merging*) and solving strategies. It also benefits from *code-level optimizations* of Clang as it operates on LLVM bytecode. Previous work [5] considers KLEE as the *worst-case off-the-shelf attacker*, in front of TRITON [38] and ANGR [43].

Benchmarks. Our benchmarks include Banescu et al.’s synthetic benchmarks [5], *enriched by 7 larger real-life programs* consisting essentially of hash functions (a typical software asset one has to protect) [37]. We also work both on source and binary code to add another level of variability. As already said, the considered programs are rather small, *on purpose*, to embody *the defender worst case*. Note that this case still represents real life situations, e.g., protecting small critical assets from *targeted* DSE attacks.

10.2 Generality of path-oriented protections

Path-oriented protections should be effective on a larger class of attacks besides DSE – actually, all major semantic program analysis techniques. Indeed, all path-unrolling methods will suffer from path explosion, including Bounded model checking [10], backward bounded DSE [8] and abstract interpretation with aggressive trace partitioning [33]. Model checking based on counter-example guided refinement [31] will suffer both from path explosion and Single Value Path protections – yielding ineffective refinements in the vein of [15]. Finally, standard abstract interpretation [4] will suffer from significant precision loss due to the many introduced *merge points* – anyway purely static techniques cannot currently cope with self-modification or packing.

10.3 Countermeasures and mitigations

Slicing, tainting and *pattern attacks*, are thoroughly discussed in Sec. 6.2 and 7.

Advanced program analysis techniques for loops is a very hot research topic, still largely open in the case of under-approximation methods such as DSE. The best methods for DSE are based on path merging [3], but they lack a generalization step allowing to completely capture loop semantics. Even though KLEE implements such path merging, it still fails against our protections. Widening in abstract interpretation [4] over-approximates loop semantics, but the result is often very crude: using such over-approximations inside DSE is still an open question. Anti-implicit flow techniques [34, 35] may identify dataflow hidden as control-flow (it identified for instance a FOR forking point), yet they do not recover any precise loop semantics and thus cannot reduce path explosion.

Finally, note that: (1) obfuscation schemes can easily be scattered along several functions (see alternative FOR encodings in Sec. 7.1) to bar expensive but targeted intra-procedural attacks – attackers will need (costly) precise inter-procedural methods, (2) real-life attacks are performed on binary code – binary-level static analysis is known to be extremely hard to get precise; and (3) static analysis is completely broken by packing or self-modification.

11 RELATED WORK

We have already discussed obfuscation, symbolic execution and symbolic deobfuscation throughout the paper, including successful applications to deobfuscation [8, 24, 37, 51]. In addition, Schrittwieser et al. [39] give an exhaustive survey about program analysis-based deobfuscation, while Schwartz et al. [40] review DSE and tainting for security.

Limits of symbolic execution. Anand et al. [2] describe, in the setting of automatic testing, the three major weaknesses of DSE: *Path explosion*, *Path divergence* and *Complex constraints*. Cadar [16] shows that compiler optimizations can sensibly alter the performance of a symbolic analyzer like KLEE, confirming the folklore knowledge that strong enough compiler optimizations resemble code obfuscations. That said, the performance penalty is far from offering a strong defense against symbolic deobfuscation.

Constraint-based anti-DSE protections. Most anti-DSE techniques target the constraint solving engine through hard-to-solve predicates. The impact on symbolic deobfuscation through the complexification of constraints has been studied by Banescu et al. [6]. Biondi et al. [11] propose an obfuscation based on *Mixed Boolean-Arithmetic* expressions [52] to complexify *points-to functions*, making it harder for solvers to determine the trigger. Eyrolles et al. [29] present a similar obfuscation together with a MBA expression simplifier based on pattern matching and arithmetic simplifications. Cryptographic hash functions hinder current solvers and can replace MBA [42]. In general, formula hardness is difficult to predict, and solving such formulas is a hot research topic. Though cryptographic functions resist solvers up to now, promising attempts exist [36]. More importantly, private keys must also be protected against symbolic attacks, yielding a potentially easier deobfuscation subgoal – a standard whitebox cryptography issue.

Other anti-DSE protections. Yadegari and Debray [50] describe obfuscations thwarting standard byte-level taint analysis, possibly

resulting in missing legitimate paths for DSE engines using taint analysis (TRITON does, KLEE and BINSEC do not). It can be circumvented in the case of taint-based DSE by bit-level tainting [50]. Symbolic Code combines this idea with trigger-based self modifications. Solutions exist but must be carefully integrated [12, 50]. Wang et al. [49] propose an obfuscation based on mathematical conjectures (e.g., Collatz) to conceal a trigger condition. This transformation increases the number of paths through specifically crafted loops, ensured by the conjecture to always converge to the same result. The main differences are that (1) we seek to protect the entire code, and (2) we do not rely on conjectures. In addition, the method is highly susceptible to pattern attacks since there are few well-suited conjectures. Banescu et al. [5] propose an anti-DSE technique based on encryption and proved to be highly effective, but it requires some form of secret sharing (the key) and thus falls outside the strict scope of MATE attacks that we consider here. Stephens et al. [45] recently proposed an obfuscation based on covert channels (timing, etc.) to hide data flow within invisible states. Current tools do not handle correctly this kind of protections. However, the method ensures only probabilistic correctness and thus cannot be applied in every context.

Systematic evaluation of anti-DSE techniques. Banescu et al. [5] set the ground for the experimental evaluation of symbolic deobfuscation techniques. Our own experimental evaluation extends and refines their method in several ways: new metrics, different DSE settings, larger examples. Bruni et al. [15] propose a mathematically proven obfuscation against Abstract Model Checking attacks.

12 CONCLUSION

Code obfuscation intends to protect proprietary software assets against attacks such as reverse engineering or code tampering. Yet, recently proposed (automated) attacks based on symbolic execution (DSE) and semantic reasoning have shown a great potential against traditional obfuscation methods. We explore a new class of anti-DSE techniques targeting the very weak spot of these approaches, namely path exploration. We propose a predictive framework for understanding such path-oriented protections, and we propose new lightweight, efficient and resistant obfuscations. Experimental evaluation indicates that our method critically damages symbolic deobfuscation while yielding only a very small overhead.

REFERENCES

- [1] Tigress challenge. <http://tigress.cs.arizona.edu/challenges.html>.
- [2] S. Anand, E. K. Burke, T. Y. Chen, J. Clark, M. B. Cohen, W. Grieskamp, M. Harman, M. J. Harrold, and P. McMinn. An orchestrated survey of methodologies for automated software test case generation. *Journal of Systems and Software*, 2013.
- [3] Thanassis Avgerinos, Alexandre Rebert, Sang Kil Cha, and David Brumley. Enhancing symbolic execution with veritesting. *Commun. ACM*, 59(6), 2016.
- [4] Gogul Balakrishnan and Thomas W. Reps. WYSINWYX: what you see is not what you execute. *ACM Trans. Program. Lang. Syst.*, 32, 2010.
- [5] Sebastian Banescu, Christian S. Collberg, Vijay Ganesh, Zack Newsham, and Alexander Pretschner. Code obfuscation against symbolic execution attacks. In *Annual Conference on Computer Security Applications, ACSAC 2016*, 2016.
- [6] Sebastian Banescu, Christian S. Collberg, and Alexander Pretschner. Predicting the resilience of obfuscated code against symbolic execution attacks via machine learning. In *USENIX Security Symposium*, 2017.
- [7] Boaz Barak, Oded Goldreich, Russell Impagliazzo, Steven Rudich, Amit Sahai, Salil P. Vadhan, and Ke Yang. On the (im)possibility of obfuscating programs. In *Advances in Cryptology - CRYPTO*, 2001.
- [8] Sébastien Bardin, Robin David, and Jean-Yves Marion. Backward-bounded DSE: targeting infeasibility questions on obfuscated codes. In *2017 IEEE Symposium on Security and Privacy, SP*, 2017.
- [9] Clark Barrett and Cesare Tinelli. *Satisfiability Modulo Theories*. Springer International Publishing, 2018.
- [10] Armin Biere. Bounded Model Checking. In *Handbook of Satisfiability*. 2009.
- [11] Fabrizio Biondi, Sébastien Josse, Axel Legay, and Thomas Sirvent. Effectiveness of synthesis in concolic deobfuscation. *Computers & Security*, 70, 2017.
- [12] Guillaume Bonfante, José M. Fernandez, Jean-Yves Marion, Benjamin Rouxel, Fabrice Sabatier, and Aurélien Thierry. Codisasm: Medium scale concolic disassembly of self-modifying binaries with overlapping instructions. In *Conference on Computer and Communications Security*, 2015.
- [13] David Brumley, Cody Hartwig, Zhenkai Liang, James Newsome, Dawn Xiaodong Song, and Heng Yin. Automatically identifying trigger-based behavior in malware. In Wenke Lee, Cliff Wang, and David Dagon, editors, *Botnet Detection: Countering the Largest Security Threat*, volume 36 of *Advances in Information Security*, pages 65–88. Springer, 2008.
- [14] Robert Brummayer and Armin Biere. Boolector: An efficient SMT solver for bit-vectors and arrays. In *International Conference on Tools and Algorithms for the Construction and Analysis of Systems, TACAS*, 2009.
- [15] Roberto Bruni, Roberto Giacobazzi, and Roberta Gori. Code obfuscation against abstract model checking attacks. In *Verification, Model Checking, and Abstract Interpretation - 19th International Conference, VMCAI*, 2018.
- [16] Cristian Cadar. Targeted program transformations for symbolic execution. In *Meeting on Foundations of Software Engineering, ESEC/FSE*, 2015.
- [17] Cristian Cadar, Daniel Dunbar, and Dawson R. Engler. KLEE: unassisted and automatic generation of high-coverage tests for complex systems programs. In *8th USENIX Symposium on Operating Systems Design and Implementation, OSDI*, 2008.
- [18] Cristian Cadar and Koushik Sen. Symbolic execution for software testing: three decades later. *Commun. ACM*, 56(2), 2013.
- [19] Mariano Ceccato, Paolo Tonella, Cataldo Basile, Paolo Falcarin, Marco Torchiano, Bart Coppens, and Bjorn De Sutter. Understanding the behaviour of hackers while performing attack tasks in a professional setting and in a public challenge. *Empirical Software Engineering*, 24(1):240–286, Feb 2019.
- [20] Sang Kil Cha, Thanassis Avgerinos, Alexandre Rebert, and David Brumley. Unleashing mayhem on binary code. In *Symposium on Security and Privacy, SP*, 2012.
- [21] Christian Collberg and Jasvir Nagra. *Surreptitious Software: Obfuscation, Watermarking, and Tamperproofing for Software Protection*. Addison-Wesley Professional, 1st edition, 2009.
- [22] Christian Collberg, Clark Thomborson, and Douglas Low. A taxonomy of obfuscating transformations, 1997.
- [23] Christian S. Collberg, Sam Martin, Jonathan Myers, and Jasvir Nagra. Distributed application tamper detection via continuous software updates. In *Annual Computer Security Applications Conference, ACSAC*, 2012.
- [24] Kevin Coogan, Gen Lu, and Saumya K. Debray. Deobfuscation of virtualization-obfuscated software: a semantics-based approach. In *Conference on Computer and Communications Security, CCS*, 2011.
- [25] Robin David, Sébastien Bardin, Josselin Feist, Laurent Mounier, Marie-Laure Potet, Thanh Dinh Ta, and Jean-Yves Marion. Specification of concretization and symbolization policies in symbolic execution. In *International Symposium on Software Testing and Analysis, ISSTA 2016*, 2016.
- [26] Robin David, Sébastien Bardin, Thanh Dinh Ta, Laurent Mounier, Josselin Feist, Marie-Laure Potet, and Jean-Yves Marion. BINSEC/SE: A dynamic symbolic execution toolkit for binary-level analysis. In *IEEE 23rd International Conference on Software Analysis, Evolution, and Reengineering, SANER*, 2016.
- [27] Leonardo Mendonça de Moura and Nikolaj Bjørner. Z3: an efficient SMT solver. In *Tools and Algorithms for the Construction and Analysis of Systems, TACAS*, 2008.
- [28] Saumya K. Debray and Jay Patel. Reverse engineering self-modifying code: Unpacker extraction. In *Working Conference on Reverse Engineering, WCRE*, 2010.
- [29] Ninon Eyrolles, Louis Goubin, and Marion Videau. Defeating mba-based obfuscation. In *Proceedings of the 2016 ACM Workshop on Software PROtection, SPRO@CCS 2016*, 2016.
- [30] Patrice Godefroid, Michael Y. Levin, and David A. Molnar. SAGE: whitebox fuzzing for security testing. *Commun. ACM*, 55(3), 2012.
- [31] Thomas A. Henzinger, Ranjit Jhala, Rupak Majumdar, and Grégoire Sutre. Lazy abstraction. In *The 29th SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*, 2002.
- [32] Min Gyung Kang, Pongsin Poosankam, and Heng Yin. Renovo: a hidden code extractor for packed executables. In *ACM Workshop Recurring Malcode (WORM)*. ACM, 2007.
- [33] Johannes Kinder. Towards static analysis of virtualization-obfuscated binaries. In *19th Working Conference on Reverse Engineering, WCRE*, 2012.
- [34] Dave King, Boniface Hicks, Michael Hicks, and Trent Jaeger. Implicit flows: Can't live with 'em, can't live without 'em. In *Information Systems Security, 4th International Conference, ICISS*, 2008.
- [35] Yin Liu and Ana Milanova. Static information flow analysis with handling of implicit flows and a study on effects of implicit flows vs explicit flows. In *14th European Conference on Software Maintenance and Reengineering, CSMR*, 2010.

- [36] Saeed Nejati, Jia Hui Liang, Catherine H. Gebotys, Krzysztof Czarnecki, and Vijay Ganesh. Adaptive restart and cegar-based solver for inverting cryptographic hash functions. In *VSTTE*, 2017.
- [37] Jonathan Salwan, Sébastien Bardin, and Marie-Laure Potet. Symbolic deobfuscation: from virtualized code back to the original. In *5th Conference on Detection of Intrusions and malware & Vulnerability Assessment (DIMVA)*, 2018.
- [38] Florent Sadel and Jonathan Salwan. Triton : Framework d'exécution concolique. In *SSTIC*, 2015.
- [39] Sebastian Schrittwieser, Stefan Katzenbeisser, Johannes Kinder, Georg Merz-dovnik, and Edgar Weippl. Protecting software through obfuscation: Can it keep pace with progress in code analysis? *ACM Comput. Surv.*, 49(1), 2016.
- [40] Edward J. Schwartz, Thanassis Avgerinos, and David Brumley. All you ever wanted to know about dynamic taint analysis and forward symbolic execution (but might have been afraid to ask). In *Symposium on Security and Privacy, S&P*, 2010.
- [41] Hovav Shacham. The geometry of innocent flesh on the bone: return-into-libc without function calls (on the x86). In *Conference on Computer and Communications Security, CCS*, 2007.
- [42] Monirul I. Sharif, Andrea Lanzani, Jonathon T. Giffin, and Wenke Lee. Impeding malware analysis using conditional code obfuscation. In *Network and Distributed System Security Symposium, NDSS*, 2008.
- [43] Yan Shoshitaishvili, Ruoyu Wang, Christopher Salls, Nick Stephens, Mario Polino, Andrew Dutcher, John Grosen, Siji Feng, Christophe Hauser, Christopher Krügel, and Giovanni Vigna. SOK: (state of) the art of war: Offensive techniques in binary analysis. In *IEEE Symposium on Security and Privacy, SP*, 2016.
- [44] Venkatesh Srinivasan and Thomas W. Reps. An improved algorithm for slicing machine code. In *Proceedings of the 2016 ACM SIGPLAN International Conference on Object-Oriented Programming, Systems, Languages, and Applications, OOPSLA 2016*. ACM, 2016.
- [45] Jon Stephens, Babak Yadegari, Christian S. Collberg, Saumya Debray, and Carlos Scheidegger. Probabilistic obfuscation through covert channels. In *European Symposium on Security and Privacy, EuroS&P*, 2018.
- [46] Julien Vanegue and Sean Heelan. SMT solvers in software security. In *6th USENIX Workshop on Offensive Technologies, WOOT'12*, 2012.
- [47] Chenxi Wang, Jonathan Hill, John Knight, and Jack Davidson. Software tamper resistance: Obstructing static analysis of programs. Technical report, Charlottesville, VA, USA, 2000.
- [48] Chenxi Wang, Jonathan Hill, John C. Knight, and Jack W. Davidson. Protection of software-based survivability mechanisms. In *International Conference on Dependable Systems and Networks (DSN)*, 2001.
- [49] Zhi Wang, Jiang Ming, Chunfu Jia, and Debin Gao. Linear obfuscation to combat symbolic execution. In *European Symposium on Research in Computer Security, ESORICS*, 2011.
- [50] Babak Yadegari and Saumya Debray. Symbolic execution of obfuscated code. In *Conference on Computer and Communications Security (CCS)*, 2015.
- [51] Babak Yadegari, Brian Johannsmeyer, Ben Whitely, and Saumya Debray. A generic approach to automatic deobfuscation of executable code. In *Symposium on Security and Privacy, SP*, 2015.
- [52] Yongxin Zhou, Alec Main, Yuan Xiang Gu, and Harold Johnson. Information hiding in software with mixed boolean-arithmetic transforms. In *Information Security Applications, WISA*, 2007.

A FORMAL PROOFS

Theorem 1 (Optimal Composition) *Suppose that $\mathcal{P}'$ is obtained by obfuscating the program $\mathcal{P}$. If every original path of $\mathcal{P}$ goes through at least k independent forking points of $\mathcal{P}'$ inserting at least θ feasible paths, then $\#\Pi_{\mathcal{P}'} \geq \#\Pi_{\mathcal{P}} \cdot \theta^k$*

PROOF. Let's consider a program $\mathcal{P}$ with $\#\Pi$ original paths σ_i , $i \in \{1 \dots \#\Pi\}$. We obfuscate $\mathcal{P}$ into $\mathcal{P}'$ with an obfuscation scheme adding n independent forking points inserting $\#\sigma_{1..n}$ feasible paths. The forking points are placed such that every original path now contains at least k forking points.

- The total number of paths of $\mathcal{P}'$ is:

$$\#\Pi_{\mathcal{P}'} = \sum_{\sigma_i} \#\sigma_i, \quad \sigma_i \in \Pi_{\mathcal{P}}$$

- According to the definition of independence, one original path σ_i with **at least** k forking points inserting $\#\sigma_{i,\{1..k\}}$ feasible paths creates $\#\sigma_i \geq \prod_{j=1}^k \#\sigma_{i,j}$ new paths
- Then,

$$\#\Pi_{\mathcal{P}'} \geq \sum_{\sigma_i} \left(\prod_{j=1}^k \#\sigma_{i,j} \right), \quad \sigma_i \in \Pi_{\mathcal{P}}$$

We write $\theta = \min_{i,j} (\#\sigma_{i,j})$

$$\#\Pi_{\mathcal{P}'} \geq \sum_{\sigma_i} (\theta^k), \quad \sigma_i \in \Pi_{\mathcal{P}}$$

$$\#\Pi_{\mathcal{P}'} \geq \#\Pi_{\mathcal{P}} \times \theta^k$$

■

Theorem 2 (Resistance by design) *Let us consider a program $\mathcal{P}$ and a forking point $\mathcal{F}$. Assuming $\mathcal{F}$ is built upon relevant variables, then $\mathcal{F}$ is slice and taint resistant.*

PROOF. By definition, a sound taint analysis $\mathcal{A}_{\mathcal{T}}$ will mark any relevant variable (as they depend from input). Hence, if $\mathcal{F}$ is built upon relevant variables, then all variables $v \in \text{Var}(\mathcal{F})$ will be marked by $\mathcal{A}_{\mathcal{T}}$, hence taint analysis $\mathcal{A}_{\mathcal{T}}$ will yield no simplification on $\mathcal{F}$. In the same manner, a sound slice analysis $\mathcal{A}_{\mathcal{S}}$ will mark any relevant variable (as they impact the output), implying that if $\mathcal{F}$ is built upon relevant variables, then analysis $\mathcal{A}_{\mathcal{S}}$ will yield no simplification on $\mathcal{F}$. ■

B STATISTICS ON DATASETS

We present additional statistics on Dataset #1 (Appendix Table 7) and Dataset #2 (Appendix Table 8). For Dataset #1, recall that 1-byte input programs from the original dataset from Banescu et al. [5] are automatically turned into equivalent 8-byte input programs with same number of paths: additional input are not used by the program, but latter protections will rely on them. We must do so as they are otherwise too easy to enumerate.

C ADDITIONAL EXPERIMENTS

Search heuristics. Results in Appendix Table 9 shows that DSE search heuristics does not impact that much overall results (cf. Table 3). Depth-first search appears to be slightly better than the two other ones for SPLIT, and non-uniform random search appears to be

Table 7: Statistics on Dataset #1 (46 programs)

Entry size	#LOC		KLEE exec. (s)	
	average	StdDev.	average	StdDev.
16 bytes	21	1.9	2.6s	6.2
1 byte (*)	17	2.2	1.8s	6.2

loc: line of code

(*) 1-byte input programs are automatically turned into equivalent 8-byte input programs with same number of paths. We report KLEE execution time on the modified versions.

Table 8: Statistics on Dataset #2 (7 programs)

Program	locs	KLEE exec. (s)
City hash	547	7.41
Fast hash	934	7.74
Spooky hash	625	7.12
MD5 hash	157	33.31
AES	571	1.42
DES	424	0.15
GRUB	101	0.06

slightly worse than the two other ones for FOR. Nothing dramatic yet.

Table 9: Impact of search heuristics – Dataset #1 – secret finding – 1h TO

	Timeouts			
	NURS	BFS	DFS	allpath
Virt	0/15	0/15	0/15	0/15
Virt ×2	0/15	0/15	0/15	0/15
Virt ×3	1/15	1/15	1/15	2/15
Flat-Virt	0/15	0/15	0/15	0/15
Flat-MBA	0/15	0/15	0/15	0/15
SPLIT (×10)	0/15	0/15	0/15	0/15
SPLIT (×13)	1/15	1/15	0/15	1/15
FOR (k = 1)	0/15	0/15	0/15	0/15
FOR (k = 2)	1/15	1/15	1/15	4/15
FOR (k = 3)	10/15	8/15	8/15	13/15
FOR (k = 4)	15/15	15/15	15/15	15/15

Runtime overhead. We evaluate how the performance penalty evolved for protection FOR on very high values of k . We take the 15 examples of Dataset #1 with large input space, and we vary the size of the input string from 3 to 100000, increasing the number of forking points accordingly (k between 3 and 100000), one forking point (loop) per byte of the input string. We run 15 random inputs 15 times for each size and measure the average runtime overhead. Fig. 9 shows the evolution of runtime overhead w.r.t. the number of FOR loops.

The runtime overhead stays below 5% ($\times 1.05$) for fewer than $k = 250$. This means in particular that one can significantly boost FOR-based protections without incurring big runtime penalties.

Figure 9: Runtime overhead w.r.t. to the number of FOR obfuscation loops

D CODE SNIPPET FOR WRITE

Fig. 10 shows an example of an assembly-level implementation of the WRITE protection over the expression “var = input;”.

```

__asm__ ( "movl %[src], (.L\%=+1) \n\t"
          "jmp .L\%= \n\t"
          "%=: \n\t"
          ".section .L\%=,\"awx\" \n\t"
          "movl \$_0, %[dst] \n\t"
          "jmp \%=b \n\t"
          ".previous \n\t"
          :[dst] "=&a" (var) : [src] "r" (input));

```

Figure 10: ASM encoding of protection WRITE