# PMP: Cost-effective Forced Execution with Probabilistic Memory Pre-planning 

Wei You ${ }^{1}$, Zhuo Zhang ${ }^{1}$, Yonghwi Kwon ${ }^{2}$, Yousra Aafer ${ }^{1}$, Fei Peng ${ }^{1}$, Yu Shi ${ }^{1}$, Carson Harmon ${ }^{1}$, Xiangyu Zhang ${ }^{1}$<br>${ }^{1}$ Department of Computer Science, Purdue University, Indiana, USA<br>${ }^{2}$ Department of Computer Science, University of Virginia, Virginia, USA<br>Email: \{you58, zhan3299, yaafer, pengf, shi442, harmon35, xyzhang\}@purdue.edu, yongkwon@virginia.edu


#### Abstract

Malware is a prominent security threat and exposing malware behavior is a critical challenge. Recent malware often has payload that is only released when certain conditions are satisfied. It is hence difficult to fully disclose the payload by simply executing the malware. In addition, malware samples may be equipped with cloaking techniques such as VM detectors that stop execution once detecting that the malware is being monitored. Forced execution is a highly effective method to penetrate malware self-protection and expose hidden behavior, by forcefully setting certain branch outcomes. However, an existing state-of-the-art forced execution technique X-Force is very heavyweight, requiring tracing individual instructions, reasoning about pointer alias relations on-the-fly, and repairing invalid pointers by on-demand memory allocation. We develop a light-weight and practical forced execution technique. Without losing analysis precision, it avoids tracking individual instructions and on-demand allocation. Under our scheme, a forced execution is very similar to a native one. It features a novel memory pre-planning phase that pre-allocates a large memory buffer, and then initializes the buffer, and variables in the subject binary, with carefully crafted values in a random fashion before the real execution. The pre-planning is designed in such a way that dereferencing an invalid pointer has a very large chance to fall into the pre-allocated region and hence does not cause any exception, and semantically unrelated invalid pointer dereferences highly likely access disjoint (pre-allocated) memory regions, avoiding state corruptions with probabilistic guarantees. Our experiments show that our technique is $\mathbf{8 4}$ times faster than $\mathbf{X}$-Force, has 6.5X and $10 \%$ fewer false positives and negatives for program dependence detection, respectively, and can expose $98 \%$ more malicious behaviors in $\mathbf{4 0 0}$ recent malware samples.


## I. Introduction

The proliferation of new strains of malware every year poses a prominent security threat. Recently reported attacks demonstrate the emergence of new attacking trends, where malware authors are designing for stealth and leaving lighter footprints. For example, Fileless malware [5] infects a target host through exploiting built-in tools and features, without requiring the installation of malicious programs. Clickless infections [1] avoid end-user interaction through exploiting shared access points and remote execution exploits. Cryptocurrency malware [4] allow attackers to generate huge revenues by illegally running mining algorithms using victim's system resources. According to [3], a massive cryptocurrency mining botnet has generated $\$ 3$ million revenue in 2018. Under this new threatscape, malicious payloads have evolved and look much different than traditional ones. Thus, a critical challenge the security community is facing today is to understand and analyze emerging malware's behavior in an effort to prevent potentially epidemic consequences.

A popular approach to understanding malware behavior is to run it in a sandbox. However, a well-known difficulty is that the needed environment or setup may not be present (e.g., C\&C server is down and critical libraries are missing) such that the malware cannot be executed. In addition, recent malware often makes use of time-bomb and logic-bomb that define very specific temporal and contextual conditions to release payload, and some samples even use cloaking techniques such as packing, and $\mathrm{VM} /$ debugger detectors that prevent execution when the malware is being monitored.

Researchers in [32] proposed a technique called forcedexecution (X-Force) that penetrates these malware selfprotection mechanisms and various trigger conditions. It works by force-setting branch outcomes of some conditional instructions. (e.g., those checking trigger conditions). As forcing execution paths could lead to corrupted states and hence exceptions, X-Force features a crash-free execution model that allocates a new memory block on demand upon any invalid pointer dereference. However, X-Force is a very heavy-weight technique that is difficult to deploy in practice. Specifically, in order to respect program semantics, when XForce fixes an invalid pointer variable (by assigning a newly allocated memory block to the variable), it has to update all the correlated pointer variables (e.g., those have constant offsets with the original invalid pointer). To do so, it has to track all memory operations (to detect invalid accesses) and all move/addition/subtraction operations (to keep track of pointer variable correlations/aliases). Such tracking not only entails substantial overhead, but also is difficult to implement correctly due to the complexity of instruction set and the numerous corner situations that need to be considered (e.g., in computing pointer relations). As a result, the original X-Force does not support tracing into library functions.

In this paper, we propose a practical forced execution technique. It does not require tracking individual memory or arithmetic instructions. Neither does it require on demand memory allocation. As such, the forced execution is very close to a native execution, naturally handling libraries and dynamically generated code. Specifically, it achieves crashfree execution (with probabilistic guarantees) through a novel memory pre-planning phase, in which it pre-allocates a region of memory starting from address 0 , and fills the region with carefully crafted random values. These values are designed in such a way that (1) if they are interpreted as addresses and further dereferenced, the addresses fall into the pre-allocated region and do not cause exception; (2) they have diverse
random values such that semantically unrelated pointer variables unlikely dereference the same random address and avoid causing bogus program dependencies and corrupted states. An execution engine is developed to systematically explores different paths by force-setting different sets of branch outcomes. For each path, multiple processes are spawned to execute the path with different randomized memory pre-planning schemes, further reducing the probability of coincidental failures. The results of these processes are aggregated to derive the results for the particular path. The engine then moves forward to the next path.

Our contributions are summarized as follows.

- We develop a practical forced-execution engine that does not entail any heavy-weight instrumentation.
- We propose a novel memory pre-planning scheme that provides probabilistic guarantees to avoid crashes and bogus program dependencies. The execution under our scheme is very similar to a native execution. Once the memory is pre-planned and initialized at the beginning, the execution just proceeds as normal, without requiring any tracking or on the fly analysis (e.g., pointer correlation analysis).
- We have implemented a prototype called PMP and evaluated it on SPEC2000 programs (which include gcc), and 400 recent real-world malware samples. Our results show that PMP is a highly effective and efficient forced execution technique. Compared to X -Force, PMP is 84 time faster, and the false positive (FP) and false negative (FN) rates are 6.5 X and $10 \%$ lower, respectively, regarding dependence analysis; and detect $98 \%$ more malicious behaviors in malware analysis. It also substantially supersedes recent commercial and academic malware analysis engines Cuckoo [2], Habo [10] and Padawan [8].


## II. Motivation

In this section, we use an example to motivate the problem, explain the limitations of existing techniques, and illustrate our idea. The code snippet in Figure 1 simulates the command and control (C\&C) behavior of a variant of Mirai [7], a notorious IoT malware that launches distributed denial of service attacks when receiving commands from the remote C\&C server. In particular, it reads the maximum number of destination hosts (to attack) from a configuration file (line 9), and allocates a Cmd object with sufficient memory to store destination information in the Dest objects (lines 10-12). When the C\&C server is connectable (line 15), the malware scans the local network for the destination hosts (line 16), receives the requested command (line 17), and performs the corresponding actions on the destination hosts (lines 18-22).
To expose such malicious behavior, analysts could run the sample in a sandbox and monitor its system call sequences and network flows [8]. Unfortunately, a naive execution-based analysis is incomplete and hence cannot reveal all the malicious payloads, especially those that are condition-guarded and environment-specific. In our example, if the configuration file
does not exist or the $\mathrm{C} \& \mathrm{C}$ server is not connectable, the malicious behavior will not be exposed at all. One may consider to construct an input file and simulate the network data. However, such a task is time-consuming and not practical for zeroday malware whose input format and network communication protocol are unknown. In addition, recent malware samples are increasingly equipped with anti-analysis mechanism, which prevents these samples from execution even if they are given valid inputs (please refer to Section IV for real-world cases). This poses great difficulties for dynamic analysis.

Forced execution [32] provides a practical solution to systematically explore different execution paths (and, hence reveal different program behaviors) without any input or environment setup. It works by force-setting branch outcomes of a small set of predicates and jump tables. One critical problem faced by forced execution is invalid memory accesses due to the absence of necessary memory allocations and initializations, which are present in normal execution. Without appropriate handling of invalid memory accesses, the program is most likely to crash before reaching any malicious payload. In our example, the malicious behaviors were supposed to be exposed, if the predicate in line 15 is forced to take the true branch, and the jump table in line 18 is forced to iterate different entries. However, the forced execution fails in line 30, because cmd is not properly allocated and its dests field is not initialized.

X-Force. In X-Force [32], researchers show that simply ignoring exceptions does not work as that leads to cascading failures (i.e., more and more crashes), they propose to recover from invalid memory accesses by performing on-demand memory allocation. In particular, X-Force monitors all memory operations (i.e., allocate, free, read and write) to maintain a list of valid memory addresses. If an accessed memory address is not in the valid list, a new memory block will be allocated on demand for the access. To respect program semantics, when a pointer variable holding an invalid address $x$ is set to the address of the allocated memory, all the other pointer variables that hold a value denoting the same invalid address or its offset (e.g., $x+c$ with $c$ some constant) need to be updated. X-Force achieves this through linear set tracing, which identifies linearly correlated pointer variables that are induced by address offsetting. When a pointer variable is updated, all the correlated pointers in its linear set need to be updated accordingly based on their offsets.

Assume in an execution instance, line 8 takes the false branch and line 15 is forced to take the true branch. In this execution, cmd is a NULL pointer, hence the dests pointer in line 27 points to $0 x 8$ (the offset of dests field is 8 ). The rounded rectangle in Figure 1 illustrates what X-Force does for the memory access of dests [0]->ip in line 30. Linear sets are maintained for each register and each memory address. In particular, $S R(r)$ and $S M(a)$ are used to denote the linear set of register $r$ and address $a$, respectively. After executing instruction $\alpha$, the linear set of register rbx is updated to be the same as that of \& dests, i.e., $S R(\mathrm{rbx}) \leftarrow S M(\& \mathrm{dests})$ such that $S R(\mathrm{rbx})=S M$ (\&dests)=\{0x7ffdfffffed0\}, which

```
typedef struct{char ip[16]; long port;} Dest;
typedef struct{long act; Dest* dests[0];} Cmd;
int main(int argc, char *argv[]) {
    Cmd *cmd = NULL;
    int max = 0;
    if (config_file_exists()) {
        max = read_from_config_file();
        cmd = malloc(sizeof(Cmd) + max*sizeof(Dest*));
        for (int i = 0; i < max; i++)
            cmd->dests[i] = malloc(sizeof(Dest));
        }
    if (cnc_server_connectable()) {
        scan_intranet_hosts(cmd, max);
        cmd->act = get_action_from_cc_server();
        switch (cmd->act) {
            case 1: do_action_1(cmd->dest, max); break;
            case 2: do_action_2(cmd->dest, max); break;
        }
    }
    ...
}
```








33 \}
$\alpha$. mov rbx, [rbp - 0x10] $/ / \mathrm{rbx}=[\mathrm{rbp}-0 \mathrm{x} 10]=[0 \times 7 \mathrm{ffdfffffed} 0]=0 \times 8$
/* Validate Memory Address: get_accessible(0x7ffdfffffed0) = true */
$/ *$ Update Linear Set: $S R(\mathrm{rbx}) \leftarrow S M(\&$ dests $)=\{0 \mathrm{x} 7 \mathrm{ffdfffffed} 0\} *$
$\beta$. mov ecx, $[\mathrm{rbp}-0 \mathrm{x} 14] / / \mathrm{ecx}=[\mathrm{rbp}-0 \mathrm{x} 14]=[0 \mathrm{x} 7 \mathrm{ffdfffffecc}]=0 \times 0$
/* Validate Memory Address: get_accessible(0x7ffdfffffed4) $=$ true */
/* Update Linear Set: $S R(\mathrm{rcx}) \leftarrow S M(\& \mathrm{i})=\{0 \mathrm{x} 7 \mathrm{ffdfffffecc}\} * /$
$\gamma$. lea rdx, $[\mathrm{rbx}+8 * \mathrm{rcx}] / / \mathrm{rdx}=\mathrm{rbx}+8 * \mathrm{rcx}=0 \mathrm{x} 8$
/* Update Linear Set: $S R(\mathrm{rdx}) \leftarrow S R(\mathrm{rbx})=\{0 \mathrm{x} 7 \mathrm{ffdfffffed} 0\} * /$
$\delta$. mov rax, [rdx] // rax $=[\operatorname{rdx}]=[0 \mathrm{x} 8]$
/* Validate Memory Address: get_accessible(0x8) = false (invalid read on 0x8) */
/* Allocate Memory Block: malloc (BLOCK_SIZE) $=0 \times 2531000$ */
$/ *$ Update Reference: $r d x=*(0 x 7 f f d f f f f f e d 0)=0 \times 2531000+0 \times 8=0 \times 2531008 * /$
$\epsilon$. mov rax, $[\mathrm{rax}] / / \operatorname{rax}=[\mathrm{rax}]=[0 \mathrm{x} 0]$
/*Validate Memory Address: get_accessible(0x0) = false (invalid read on 0x0) */
/* Allocate Memory Block: malloc(BLOCK_SIZE) $=0 \times 2532000$ */
$/ *$ Update Reference: $\mathrm{rdx}=*(0 \mathrm{x} 7 \mathrm{ffdfffffed} 0)=0 \times 2532000+0 \times 8=0 \times 2532008 * /$

Fig. 1: Motivation example. The assembly code here is functionally equivalent with the original one for easy understanding.
is the address of dests. Intuitively, the pointer value in $r b x$ is linearly correlated to that in dests. Hence, fixing either one entails updating the other. The linear correlation is further propagated to register $r d x$ after executing instruction $\gamma$, since its value is derived from rbx by address offsetting (i.e., \&dests $[0]=\&$ dests +0 ). When executing instruction $\delta$, X-Force detects an invalid access through the pointer denoted by rdx (i.e., \&dests[0]), holding an invalid address 0x8. Hence, it allocates a memory block with address 0x2531000 and initializes it with zero values. Register $r d x$ is then updated to 0x2531008. The value of \&dest should also be updated, since it linearly correlates with $r d x$. Similar memory recovery operations are needed for instruction $\epsilon$ that accesses dests[0]->ip through an invalid memory address 0x0.

As we can see that each memory operation should be intercepted by X-Force for memory address validation and linear set tracing. Upon the recovery of an (invalid) pointer variable, all the linearly correlated variables need to be updated accordingly. This causes substantial performance degradation. It was reported that X-Force has 473 times runtime overhead over the native execution [32]. Furthermore, since many library functions such as string functions in glibc can lead to linear set explosion (due to substantial heap array operations), XForce chose not to trace into library functions to update linear sets. As a result, its memory recovery is incomplete (see Section IV for a real-world example).

Our technique. We propose a novel randomized memory preplanning technique (called PMP) to handle invalid memory accesses with probabilistic guarantees. Instead of allocating new memory blocks on demand, PMP pre-allocates a large memory block with a fixed size (e.g., 16 KB ) when the program is loaded. The pre-allocated memory area (PAMA) is filled with carefully crafted random values such that if these values are interpreted as memory addresses, the corresponding
accesses still fall into PAMA. We call this self-contained memory behavior (SCMB). In addition, these random values are designed in a way that they are self-disambiguated. That is, it is highly unlikely that two semantically unrelated memory operations access the same random address, causing bogus dependencies. We call this self-disambiguated memory behavior (SDMB). For example, the simplest way to achieve SCMB is to pre-allocate a chunk of memory starting at $0 \times 00$ and fill it with $0 x 00$. As such, dereferences of null pointers (e.g., $* p$ with $p=0$ ) or pointers with some offset from null (e.g., $*(p+8)$ ), yield value $0 \times 00$ due to the initialization. If the yielded value $0 x 00$ is further interpreted as a pointer, its dereference continues to yield $0 x 00$, without causing any memory exception. However, such a scheme leads to substantial bogus program dependencies as semantically unrelated memory operations through uninitialized/invalid pointer variables all end up accessing address $0 x 00$. For example, assume $p$ and $q$ are not properly initialized and both have a null value due to forced execution and there are two pointer dereference statements " $1 . * p=\ldots ; \quad 2 . \ldots=* q$ ". A bogus dependence will be introduced between 1 and 2 . Such bogus dependencies further lead to highly corrupted program states. SDMB is to ensure that unrelated pointer variables have a high likelihood to contain disjoint addresses such that it is like they were all properly allocated and initialized. Intuitively, PMP diversifies the values filled in the pre-allocated large memory region such that dereferences at different offsets yield different values. Consequently, follow-up dereferences (of these values) can continue to disambiguate themselves.

In addition to the aforementioned pre-planning, during execution, PMP also initializes global, local variables, and heap regions allocated by the original program logic with random values pointing to PAMA. Note that otherwise they are initialized to 0 by default. As such, when these variables are interpreted as pointers and dereferenced without being


Fig. 2: Pre-allocated memory area. The data is presented in the little-endian format for the x86_64 architecture. The bytes in gray are free to be filled with 8 -multiple random values.
properly initialized along some forced path, the accesses still fall in PAMA and also have low likelihood to collide (on the same address). Through SCMB, PMP enables crash-free memory operations, which are critical for forced execution. Since it does not require tracing memory operations or performing on-demand allocation, it is 84 times faster than X Force (Section IV). Through SDMB, PMP respects program semantics such that it can faithfully expose (hidden) program behaviors with probabilistic guarantees. As shown in our evaluation (Section IV), PMP has fewer false positives (FP) and false negatives (FN) than X-Force as well.

Figure 2 illustrates a $64-\mathrm{KB}$ pre-allocated memory area mapped in the address space from $0 x 0$ to $0 x f f f f$. Note that although this memory region may overlap with some reserved address ranges, we leverage QEMU's address mapping to avoid such overlap (see Section III-E). It is filled with crafted random values that ensure both SCMB and SDMB. For our motivation example, instruction $\delta$ reads the memory unit at address $0 x 8$ (i.e., \&dests[0]) and gets the value $0 x 3850$. Subsequently, the instruction $\epsilon$ uses $0 x 3850$ as the address to access dests[0]->ip. These two accessed addresses ( $0 \times 8,0 \times 3850$ ) are contained in the PAMA, hence no memory exception occurs. The data dependence between these two addresses are also faithfully exposed, without undesirable address collision. Observe that there is no memory validation and linear set tracing required.

We want to point out while SCMB and SDMB can be effectively ensured in forced execution, they may not be as effective in regular execution. Otherwise, dynamic memory allocation could be completely avoided. The reason is that forced execution aims to achieve good coverage to expose program behaviors such that it bounds loop iterations [32]. As a result, linear scannings of large memory regions are mostly avoided, allowing to establish SCMB and SDMB effectively and efficiently. Intuitively, one can consider that our design is equivalent to pre-allocating many small regions that are randomly distributed. This is particularly suitable for heap accesses in forced-execution as they tend to happen in smaller memory regions. Even if overflows might happen, the likelihood of critical data being over-written is low due to the random distribution.

## III. DESIGN

## A. Overview

Figure 3 presents the architecture of PMP, which consists of three components: the path explorer, the dispatcher and the


Fig. 3: Architecture of PMP.
executors. Given a target binary, the path explorer systematically generates a sequence of branch outcomes to enforce, including the PCs of the conditional instructions and their true/false values. We call it a path scheme. Note that like X-Force, PMP does not enforce the branch outcome of all predicates, but rather just a very small number of them (e.g., less than 20). The other predicates will be evaluated as usual. PMP operates in rounds, each round executing a path scheme. For each path scheme, PMP further generates multiple versions of variable initializations, each having different initial values but satisfying both SCMB and SDMB. We call them memory schemes. The reason of having multiple memory schemes is to reduce the likelihood of coincidental address collisions. A process is forked for each path and memory scheme and distributed to an executor for execution. At the end of a round, the dispatcher aggregates the results from the executors (e.g., coverage). Another path scheme is then computed by the path explorer to get into the next round, based on the results from previous rounds.

Path Explorer. In essence, path exploration is a search process that aims to cover different parts of the subject binary. In each round, a new path scheme is determined by switching additional/different predicates, or enforcing additional/different jump table entries, to improve code coverage. Since the search space of all possible paths is prohibitively large for real-world binaries, PMP follows the same path exploration strategies in X-Force [32], including the linear search, the quadratic search and the exponential search. In particular in each round, the linear search selects a new predicate or jump table entry to enforce, which is usually the last one that does not have all its branches covered in previous rounds. The exponential strategy aims to explore all combinations of branch outcomes and is hence the most expensive. It is only used to explore some critical code regions. Quadratic search falls in between the two. Since these are not our contributions, interested readers are referred to the X-Force project [32].
Dispatcher. The dispatcher aggregates execution results (e.g., code coverage and program dependencies) of multiple executors in a conservative fashion. Specifically, it considers a result valid if and only if it is agreed by $n$ executors, with $n$ configurable. In our experience, $n=2$ is good enough in practice. Such aggregation further improves our


Fig. 4: Workflow of Memory-preplanning.
probabilistic guarantees. Intuitively, assume PMP ensures that a reported result has lower than $p \in[0,1]$ probability to be incorrect during a single execution (on an executor), due to the inevitable accidental violations of SCMB or SDMB. The aggregation further reduces the probability to $p^{n}$ if the memory schemes on the various executors are truly randomized (and hence independent).

Executors. All executors are forked from the same main process with the same initialized PAMA. Each executor then enforces a given path and memory scheme assigned to it. Such a design avoids the redundant initialization of PAMA. Note that all memory accesses must start from some variable, whose value is fully randomized across executors.

The rest of this section will explain in details the memory pre-planning step and the probability analysis for SCMB and SDMB guarantees. Execution result aggregation is omitted due to its simplicity.

## B. Memory Pre-planning

Overview. Figure 4 presents the workflow of memory preplanning. When a program is loaded, a pre-allocated memory area (PAMA) is prepared by invoking the mmap system call to map a crafted file to the program address space. The file content is randomly generated beforehand. During execution, program variables (including global, local variables and heap regions) are initialized by PMP with random eight-multiple values pointing to PAMA. Specifically, PMP intercepts: 1) the program entry point for initializing global variables; 2) call instructions for initializing local variables; and 3) memory allocations for initializing heap regions. Note that PAMA preparation happens a priori and incurs negligible runtime overhead, while variable initialization occurs on-the-fly during execution. Both are generic and do not require case-by-case crafting. We further discuss these steps in the following.
PAMA Preparation. PAMA is mapped at the lower part of the address space starting from $0 x 0$, in order to accommodate null pointers or pointers with invalid small values. The wordaligned addresses within PAMA (i.e., those having 0 at the lowest three bits) are filled with carefully crafted random values, such that if these values are interpreted as addresses, they fall within PAMA. As such, the range of random values that we can fill is dependent on the size of PAMA. For a 64-KB PAMA (i.e., in the address range of [0, 0xffff]), the first two least-significant bytes of a filling value are free to be set with a random eight-multiple value. Other bytes are fixed to zero. Note that such a value is essentially a valid
word-aligned address in PAMA. For a $64-\mathrm{MB}$ PAMA, the first three least-significant bytes of a filling value can be set randomly, providing better SDMB. The maximum PAMA can be as large as 128 TB, as a larger PAMA would overlap with the kernel space. While a feasible design is to change the entire virtual space layout (by changing kernel), it would hinder the applicability of our technique. In practice, we find that 4-MB of PAMA provides a good balance of SCMB and SDMB.

Global Variable Initialization. In an ELF binary, the uninitialized or zero-initialized global variables are stored in the . bss segment. During loading, PMP reads the offset and size information of the .bss segment from the ELF header. PMP then initializes the segment like a heap region.

Heap Initialization. Pre-planning heap regions that are dynamically allocated by instructions in the subject binary is relatively easier. PMP intercepts all memory allocations and set the allocated regions to contain random word-aligned PAMA addresses. Note that PMP writes these values to each word-aligned address in the heap region. If a regular compiler is used to generate the subject binary, the compiler would enforce pointer-related memory accesses to be word-aligned through padding. However, malware may intentionally introduce pointer accesses that are not word-aligned. Section III-E will discuss how PMP handles such cases. In the following discussion, we always assume word alignment.

Local Variable Initialization. Initializing local variables is more complex. After initializing PAMA and before spawning the executors, PMP initializes the entire stack region like a heap region. Note that stack frames are pushed and popped frequently and the same stack address space may be used by many function calls. As such, the stack space may need to be re-initialized. A plausible solution is to identify stack frame allocations (e.g., updates of $r s p$ register) and conduct initialization after each allocation. However, due to the flexibility of stack allocations, it is difficult to precisely identify them. Inspired by stack canaries used to detect stack overflows, PMP uses the following design to initialize stack regions. It intercepts each function invocation. Then starting from the current address denoted by $r s p$, it randomly checks eight ${ }^{1}$ unevenly distributed addresses lower than the rsp address (i.e., the potential stack space to be allocated), in the order from high to low, to see if they are PAMA addresses (meaning that they were not overwritten by previous function invocations). We also call these addresses canaries without causing confusion in our context and use $C_{i}$ to denote the $i$ th canary. PMP identifies the lowest (last) canary that is not PAMA address, say $C_{t}$, and then re-initializes $\left[C_{t+1}, r s p\right.$ ] (note that stack grows from high address to low address). If all eight canaries are overwritten, PMP continues to check the next eight. Observe that since stack writes may not be continuous, the detection scheme has only probabilistic guarantees. In practice, our scheme is highly

[^0]```
typedef struct{double *f1; long *f2;} T;
typedef struct{char f3; long *f4; long *f5;} G;
G *g;
void case3() {
        long *e = NULL, *f = NULL;
        if (condl()) init(e, f);
        if (cond2()) {
            *e = 0x6038; // [0\times0000] = 0\times6038
        long tmp = *f; // tmp = [0x0000]: bogus dep!
        }
}
void case4() {
    if (condl()) init(g);
        if (cond2()) {
        *(g->f4) = 0x0830;
        long tmp = *(g->f5); // &(g->f5) = 0x10000
    }
}
```

```
void case1() {
    long **a = malloc(...);
    T *b;
    if (condl()) init(b);
    if (cond2()) {
        long *alias = b->f2;
        *(b->f2) = **a; // [0x0008] = [0x0010]
        * (b->f1) = 0.1; // [0xffd0] = 0.1
        long tmp = *alias;
    }
}
void case2() {
    long *c; double **d;
    if (cond1()) init(c, d);
    if (cond2()) {
        *c = 0xdeadbeef; // [0xffd8] = 0xdeadbeef
        double tmp = **d; // [Oxdeadbeef]: error!
    }
}
```

(a) code snippet.

(b) memory scheme.

Fig. 5: Memory pre-planning.
effective and we haven't encountered any problems caused by incorrect stack initialization.

Example. We use the code snippet shown in Figure 5a as an example to explain the memory pre-planning process. In the code, a global variable $g$ is defined at line 3, two local variables $a$, b are defined in function case1 (). Assume in an execution instance, line 24 takes the false branch and $b$ is not allocated and initialized; and line 25 is forced to take the true branch. Although a is initialized by the original program code with an allocated heap region, the data in the heap region is not initialized. Without memory pre-planning, the program would have exception at any of the memory operations in lines 26-29.

In this example, the global variable $g$ is set to a random PAMA address at the beginning. Upon calling case1(), PMP checks the canaries at $C_{1}, C_{2}$, and so on (see the stack frame in the top-left corner of Figure 5b), and then identifies, say, the region from $\left[C_{3}, r s p\right]$ needs re-initialization, which includes local variables a and b. Inside the function body, $a$ is set to a dynamically allocated heap region at line 22 , but other variables such as $g$ and $b$ keep their initial PAMA address value (as line 24 is not executed). Specifically, $g$ and b point to 0xfff0 and 0x20 (in PAMA), respectively. Consider the read operation at line 28 that triggers pointer dereferences on
b and then $b->f 1$. The former dereferences address $0 \times 20$ and yields value 0xffd0, which is further interpreted as an address in the follow-up dereference of $b->f 1$, yielding another valid PAMA address. Observe that any following dereferences will be within PAMA and do not cause any exceptions, illustrating the SCMB property. The value of $b->£ 1$ (i.e., $0 x f f d 0$ ) dereferenced at line 28 is different from that of $b->£ 2$ (i.e. 0x08) dereferenced at line 27, and hence disambiguate themselves, illustrating SDMB.

## C. Other PAMA Memory Behavior and Interference with Regular Memory Operations.

Memory pre-planning is particularly designed to handle exceptional memory operations (caused by forced execution). As such, all the values filled in PAMA are essentially in preparation for these values being interpreted as addresses and further dereferenced. It is completely possible that the subject binary does not interpret values from PAMA as addresses. For example, it may interpret a PAMA region as a string and access individual bytes in the region. In such cases, the accessed values are just random values. This is equivalent to how X-Force handles uninitialized/undefined buffers.
A PAMA location can be written to and later read from by instructions in the subject binary, dictated by the program semantics. Program dependencies induced by PAMA are no
different from those induced through regular memory regions. For example, the code at line 26 in Figure 5a establishes an alias between variable alias and b->£2. At line 27, a memory write is conducted on $\mathrm{b}->£ 2$. At line 29 , a memoryread is conducted on alias. PMP can correctly establish the dependence between line 27 and line 29 , since they both point to the same memory address 0x8.

It may happen that a PAMA location is written to by the subject binary and then read through a semantically unrelated invalid pointer dereference later. As the written value may not be a legitimate PAMA address, the later read causes exception. For example, line 37 at function case2() of Figure 5a writes a value 0xdeadbeef that is not a word-aligned address within PAMA to the address indicated by pointer c. Assume c happens to have the same value $0 x f f d 8$ as an unrelated pointer $d$. The write to $* c$ also changes the value in $\star d$ to 0xdeadbeef. As such at line 38, an exception is triggered for the read of $* * \mathrm{~d}$. In the next subsection, our probability analysis shows that such cases rarely happen as the likelihood for two semantically unrelated pointers are initialized to the same random value is very low. Furthermore, PMP employs different memory schemes in multiple executors, further reducing such possibility.

In the worst situation, the subject binary uses its own instructions to set semantically unrelated pointers to null. In normal execution, these pointers would point to different properly allocated memory regions. However in forced execution, they may not be allocated, and all point to address 0 . In such cases, PMP cannot disambiguate the accesses of these variables, and lead to bogus dependencies. For example, the local variables $e$ and $f$ in function case 3 () of Figure 5a are explicitly set to null by the original program code. In forced execution where line 7 is not executed, they point to the same address $0 x 0$, resulting in bogus dependence (e.g., between lines 9 and 10). Our experimental results in Section IV show that such cases rarely happen.

## D. Probability Analysis

In this section, we study the probabilistic guarantee of PMP for the SCMB and SDMB properties. Violations of SCMB lead to exceptions whereas violations of SDMB lead to bogus dependences and corrupted variable values. To facilitate discussion, we introduce the following definitions. Let PA be the set of all possible addresses within PAMA, and WA be its word-aligned subset. Assume the size of PAMA is $S$. Then, on a 64-bit architecture, we have equation (1).

$$
\begin{equation*}
S=|\mathrm{PA}|=|\mathrm{WA}| \times 8 \tag{1}
\end{equation*}
$$

In addition, let FV be a random subset of WA, called the filling value set, whose elements are used as the values to be filled in PAMA. Without loss of generality, we assume 0 belongs to FV . We define the ratio between the size of FV and the size of WA as diversity, denoted as $d$. Then, we have equation (2).

$$
\begin{equation*}
|\mathrm{FV}|=|\mathrm{WA}| \times d=\frac{d \cdot S}{8} \tag{2}
\end{equation*}
$$

The initialization of PAMA can be formulated as a mapping $f: W A \mapsto F V$, which assigns each word (with 8 bytes alignment) in PAMA (i.e., denoted by addresses in WA) with a random value selected from FV. Intuitively, a more diverse FV leads to a more random memory scheme. The initialization that fills the whole PAMA with value 0 can be considered an extremal case where FV contains only a single element 0 . Note that in this case, SCMB is fully respected, while SDMB is substantially violated as all invalid memory operations collide on address 0 .

Probabilistic Guarantee of SCMB. When a pointer variable is initialized (by PMP) with a value indicating an address close to the end of PAMA, dereference of its offset may result in an access out of the bound of PAMA. As an example, consider the dereference of $g->£ 5$ at line 18 of function case 4 () in Figure 5a. Recall that $g$ is set to be 0xfff0 by PMP. The address of $g->£ 5$ is hence $0 \times 10000$, out of the bound of PAMA with 16 KB size.

Theorem 1. Let $x$ be a filling value selected from $\mathrm{FV}, \alpha$ be an offset. The probability $P_{e r r 1}$ of $x+\alpha$ being out of the bound of PAMA is calculated by equation (3).

$$
\begin{equation*}
P_{e r r 1}=P((x+\alpha) \notin \mathrm{PA} \mid x \in \mathrm{FV})=\frac{\alpha}{S-8} \cdot\left(1-\frac{8}{d \cdot S}\right) \tag{3}
\end{equation*}
$$

Proof. For PMP to access an out-of-bound address $x+$ $\alpha, x$ must belong to an address set IA $=W A \cap$ $\{S-\alpha, S-\alpha+1, \ldots, S-1\}$. To simplify discussion, let $\alpha^{\prime}=$ $|I A|=\alpha / 8, S^{\prime}=|W A|$ and $N=|F V|$. Let the size of IA $\cap F V$ be $i$. We can infer conditional probability $P(x \in I A \mid x \in \mathrm{FV})=$ $i / N$, denoted as $P_{i 1}$. Additionally, because there are $\binom{S^{\prime}-1}{N-1}$ possible FVs that could be uniformly chosen from (recall $0 \in \mathrm{FV}$ always holds) and $\binom{\alpha^{\prime}}{i} \cdot\binom{S^{\prime}-\alpha^{\prime}-1}{N-i-1}$ FVs have $i$ common elements with IA, $P(|F \vee \cap I A|=i)=\binom{\alpha^{\prime}}{i} \cdot\binom{S^{\prime}-\alpha^{\prime}-1}{N-i-1} /\binom{S^{\prime}-1}{N-1}$, denoted as $P_{i 2}$. Enumerating size $i \in\left\{1, \ldots, \alpha^{\prime}\right\}, P_{\text {err } 1}=$ $\sum_{i=1}^{\alpha^{\prime}} P_{i 1} \cdot P_{i 2}=\left(\alpha^{\prime} / N\right) \cdot\left(\binom{S^{\prime}-2}{N-2} /\binom{S^{\prime}-1}{N-1}\right)=\frac{\alpha}{S-8} \cdot\left(1-\frac{8}{d \cdot S}\right)$

Intuitively, the larger the pre-allocated memory area (i.e., $S)$ and the lower the diversity (i.e., $d$ ), the lower the $P_{e r r}$. In particular, the $P_{\text {err } 1}$ of a naive initialization that fills PAMA with value 0 is 0 . In a typical setting of $S=0 \times 400000, \alpha=8$ and $d=1, P_{\text {err } 1}=1.9073 \mathrm{e}-06$, illustrating a very low chance of exception. A plausible way to completely avoid SCMB violation is to avoid using address values close to the end of PAMA. However this requires knowing the largest possible offset, which is difficult in practice.

Probabilistic Guarantee of SDMB. SDMB will be compromised when two unrelated pointers are initialized to the same value by chance. Taking local variables c and $d$ for case2 () in Figure 5a as an example, both of them are initialized to 0xffd8, causing invalid pointer dereference at line 38.

Theorem 2. Let $x$ and $y$ be two filling values independently selected from FV. The probability $P_{\text {err } 2}$ of coincidental address collision, when $x$ and $y$ have the same value, is calculated by equation (4).

$$
\begin{equation*}
P_{e r r 2}=P(x=y \mid x \in \mathrm{FV}, y \in \mathrm{FV})=\frac{8}{d \cdot S} \tag{4}
\end{equation*}
$$

Proof. Recall $x$ and $y$ are independently selected from FV. Thus, fixing $x=v_{0}$ as a constant, we can infer $P_{\text {err } 2}=P(y=$ $\left.v_{0} \mid y \in F V\right)=1 /|F V|=8 /(d \cdot S)$.

With a typical setting $d=1$ and $S=0 \times 400000, P_{\operatorname{err} 2}=$ $1.9073 \mathrm{e}-06$, a very low probability.

$$
\begin{align*}
P_{e r r} 3 & =P(l(x, \beta) \cap l(y, \gamma) \neq \emptyset \mid x \in \mathrm{FV}, y \in \mathrm{FV}) \\
& \leq \frac{64}{d^{2} \cdot S^{2}}+\left(1-\frac{8}{d \cdot S}\right)^{2} \cdot \frac{\beta+\gamma-8}{S-8} \tag{5}
\end{align*}
$$

Proof is elided due to space limitations. With a setting of $\beta=0 \times 1000, \gamma=0 \times 1000$, and the rest as the same before, $P_{\text {err } 3}=0.00195$, still reasonably low. Note that one can always improve the guarantee by having more executors with different pre-plans.

## E. Implementation

PMP is implemented based on the QEMU user-mode emulator [9]. Specifically, PMP instruments conditional jumps and indirect jumps to enforce path scheme. A path scheme is a sequence of branch outcomes that need to be enforced. As an instance, "401a4c:T, 4094fc:F, 40a322\#40a566" is a path scheme that contains three branch outcomes to be enforced in order. Particularly, the predicates at 0x401a4c and 0x4094fc should take the true branch and false branch respectively, the jump table at $0 \times 40 a 322$ should take the entry at $0 \times 40 a 566$. Currently, PMP supports ELF binary on the x86_64 platform. It can be easily extended to support other architectures due to the cross-platform feature of QEMU. We leave it as our future work. In the rest of the subsection, we discuss a number of practical challenges faced by PMP.
Handling File and Network I/O, Infinite Loop and Recursion. Forced execution may result in exceptional program behaviors, such as invalid file/network access, infinite loop and infinite recursion. To make PMP applicable to real-world executables, these issues need to be handled. PMP follows similar solutions to X -Force regarding these problems. The difference lies in that we implement them on QEMU while X-Force was on PIN. We briefly discuss these solutions for the completeness of discussion.

To handle invalid file access, PMP wraps file open functions (e.g., open and fopen). If the file to be opened does not exist, a file padded with random values will be used. To handle infinite loop, PMP adopts the profiling-based approach proposed in [31] to dynamically identify loop structures. For each identified loop structure, PMP resets the loop bound to a pre-define constant. This is more sophisticated than XForce, which uses a fixed global loop bound. To handle infinite recursion, PMP intercepts call and return instructions to maintain a call stack. At each function invocation, PMP checks whether the appearances of the target function in the call stack exceed a pre-defined threshold. If so, PMP skips the function invocation. Note that while maintaining a faithful
shadow call stack is very challenging due to the various strange calling conventions, PMP does not require a precise shadow stack.

Allocation of Large PAMA. PAMA is located at the lower part of the address space starting from $0 x 0$. The default load address for non-position-independent executables is usually $0 x 400000$. If the size of PAMA is larger than 4 MB , there will be overlap between PAMA and the text/data segment of the subject executable, which is problematic.

To support large-size PAMA, we enable the address mapping mechanism provided by QEMU, which translates a guest address (denoted as GA) used by the subject executable to a host address (denoted as HA) used by QEMU. In the usermode emulation, QEMU and the subject executable share the same address space. The address mapping $g 2 h$ is flattened to essentially an offsetting operation, such that $h a=g 2 h(g a)=$ $g a+b a s e$, where $g a \in \mathrm{GA}, h a \in \mathrm{HA}$, and base is a pre-defined base address. We set the base address to the size of PAMA to avoid any overlap. Consequently, we need to adjust the filling values accordingly such that they are mapped to the addresses within PAMA (started from 0x0 in the host space). Formally, let $F V^{\prime}$ be the set of the adjusted filling values. Then we have $\mathrm{FV}^{\prime}=\{x-$ base $\mid x \in \mathrm{FV}\}$.
Misaligned Memory Access. The memory pre-planning of PMP assumes that any pointer field of a structure is wordaligned. It is a reasonable assumption for most real-world applications, since making pointer fields word-aligned (by padding if needed) is the default behavior of compilers. For example, mainstream compilers will place a 7-byte padding between the $£ 3$ field and the $£ 4$ field of the structure $G$ in Figure 5a by default, such that the offset of $£ 4$ is word-aligned.

Although we didn't find any real-world cases in our evaluation, it is possible to disable word-alignment via a special compilation option. The misalignment of a pointer field (within PAMA) may result in invalid memory access. For example, assume the global variable $g$ in Figure 5a points to 0xfff0 set by PMP. If its pointer field $£ 4$ is not wordaligned, its value will be loaded from 0xfff1, which would be 0xe800000000000050. If this value is used as an address, the access falls out of PAMA (even out of the user address space) and causes exception.

We develop the following mechanism in the dispatcher to handle misaligned memory accesses in a demand driven fashion. If a path scheme results in invalid memory access in all the executors (most likely induced by misaligned accesses), the dispatcher checks the QEMU exception $\log$ to acquire the instruction $i$ that accesses misaligned address. Then PMP additionally intercepts the code generation of instruction $i$ to mask the most-significant bytes of the accessed memory address to make it fall within PAMA. Note that while our design anticipates misaligned pointer field accesses are rare, which is true according to our experience (see Section IV), it is possible future malware may purposely introduce lots of such misalignments. In this case, PMP would have to instrument all memory operations to sanitize the addresses.

## IV. Evaluation

## A. Experiment Setup

We evaluate PMP with the SPEC2000 benchmark set as well as a set of malware samples provided by VirusTotal [12] and Padawan [8]. The experiment on SPEC2000 is conducted on a desktop computer equipped with an 8 -core CPU (Intel $®$ ® Core ${ }^{\mathrm{TM}}$ i7-8700 @ 3.20 GHz ) and 16G main memory. The experiment on the malware samples is conducted on a virtual machine (to sandbox their malicious behaviors) hosted on the same desktop. On both experiments, the configuration of PMP is as follows: 4-MB pre-allocated memory area (i.e., $S=0 \times 400000$ ), diversity $d=1$, and 2 executors (i.e., $n=2$ ).

## B. SPEC2000

SPEC2000 is a well-known benchmark set contains 12 real world programs, some of them are large (e.g., 176.gcc). The list of programs and the characteristics of their executables can be found in Appendix A. We choose SPEC2000 for the purpose of comparison as it was used in X-Force. Table I presents the comparative results on different aspects, including forced execution outcomes, code coverage and memory dependence.
Forced Execution. In this experiment, both PMP and X-Force use the same linear path exploration strategy. Specifically, it first executes the binary once without forcing any branch outcome. Then it traverses the executed predicates in the reverse temporal order (the last predicate first) and finds the predicate that has an uncovered branch. A new path scheme is then generated to force-set the uncovered branch. The procedure repeats until there are no more schemes that can lead to new coverage. Column 2 in Table I reports the total execution time when PMP finishes the exploration. Columns 3 and 4 present the number of executions that pass and fail (i.e., encounters an exception), respectively. The number in parentheses denote the number of executions finished per second. Columns 1113 show the corresponding results for X-Force. From these results, we have the following observations. (1) PMP can perform 12.6 forced executions per second on average, which is 84 times faster than X-Force ( 0.15 execution per second). Since PMP uses 2 executors for each path scheme, one may argue that X-Force can be parallelized to use two cores (for fair comparison). We want to point out that first it is unclear how to parallelize the linear search algorithm; and the second executor in PMP is just to provide better probabilistic guarantees. In most cases, such improvement may not have practical impact (see our next experiment). Hence in deployment, additional executors may be turned off. (2) The execution failure rate of PMP is $3.5 \%$, which is reasonably low and comparative with X-Force. Note that the rate is higher than what we identified in the SCMB probability analysis (Section III-D). The reason is that the majority of failures reported by both PMP and X-Force are not caused by memory exceptions, but rather inevitable as the path explorer forces the execution to enter branches that must lead to failures (e.g., forcing the true branch of a stack smash check inserted by the compiler).

Code Coverage. Columns 5~7 and 14~16 show the code coverage of PMP and X-Force, respectively. Observe that on average PMP covers $83.8 \%$ instructions, $79.1 \%$ basic blocks and $91.8 \%$ functions, which is comparable to X-Force. For most of the benchmark programs, PMP achieves more than $80 \%$ code coverage. Specifically, for $m c f$ and gzip, PMP achieves $100 \%$ code coverage.

The worst cases are eon and $g c c$. Further manual inspection shows that this is due to some inherent shortcoming of the linear search strategy. To illustrate, consider the code snippet in Figure 6, which is extracted from $g c c$ that validates function arguments before proceeding. When the check_arg () function is invoked for the first time at line 2 , the true branch of predicate at line is taken by default. The linear path exploration will force the next execution to take the false branch, since it has not been covered before. At the second-time invocation of check_arg() at line 3, the false branch of the predicate at line 8 will not be forced to execute again (hence take the true branch by default), since it has been covered before. That means, the code after line 3 will not get executed due to the validation failure at line 3 .
The essence of the problem is that linear search only focuses on predicates, without considering their context. For example, function check_arg () may be invoked from multiple places, and each calling context should be considered differently. That is, a branch being covered in a context should not prevent it from being explored again in a different context. In our future work, we will explore a context-sensitive path exploration method that can provide probabilistic guarantees. Specifically, we will explore a sampling algorithm that can sample a predicate, together with its unique context, in a specific distribution (e.g., uniform distribution).
Memory Dependence. We also conducted an experiment, in which we detect the program dependencies exercised by forced execution. A dependence is exercised when an instruction writes to some address, which is later read by another instruction. This is to evaluate the SDMB property of PMP. Note that it is intractable to acquire the ground truth of program dependencies, even with source code (due to reasons such as aliasing). Therefore, we use two methods to evaluate the quality of detected dependencies. First, we run the SPEC programs on the inputs provided by the SPEC suite (some of them are large and comprehensive) and collect the dependencies observed. These must be true positive program dependencies. As such, forced execution is supposed to expose most of them. Any missing one is an FN. Second, we built a static type checker to check if the source and destination of a (detected) dependence must have the same type. We developed an LLVM pass to propagate symbolic information to individual instructions, registers, and memory locations such that we know the type of each binary operation and its operands. Note that we need the symbolic information just for this experiment. PMP operates on stripped binaries. Ideally, force execution should report as few mistyped dependencies as possible. Each mistyped dependence must be an FP. Columns $8 \sim 10$ and

TABLE I: SPEC2000 Results

| Benchmark | PMP |  |  |  |  |  |  |  |  | X-Force |  |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | execution status |  |  | code coverage |  |  | memory dependence |  |  | execution status |  |  | code coverage |  |  | memory dependence |  |  |
|  | time (s) | \# run | \# fail | \# insn | \# block | \# func | \# found | \# correct | \# mistyped | time (s) | \# run | \# fail | \# insn | \# block | \# func | \# found | \# correct | \# mistyped |
| 164.gzip | 24.6 | $\begin{gathered} 382 \\ (15.6 / \mathrm{s}) \end{gathered}$ | $\begin{gathered} 11 \\ (3 \%) \\ \hline \end{gathered}$ | $\begin{array}{\|c\|} \hline 7,650 \\ (100 \%) \\ \hline \end{array}$ | $\begin{gathered} 699 \\ (99 \%) \\ \hline \end{gathered}$ | $\begin{gathered} 61 \\ (100 \%) \\ \hline \end{gathered}$ | 3,529 | $\begin{aligned} & 2,824 \\ & (80 \%) \\ & \hline \end{aligned}$ | $\begin{gathered} 0 \\ (0 \%) \\ \hline \end{gathered}$ | 2,112 | $\begin{array}{\|c\|} \hline 369 \\ (0.17 / \mathrm{s}) \\ \hline \end{array}$ | $\begin{gathered} 10 \\ (3 \%) \\ \hline \end{gathered}$ | $\begin{aligned} & 7,420 \\ & (97 \%) \\ & \hline \end{aligned}$ | $\begin{gathered} 669 \\ (95 \%) \\ \hline \end{gathered}$ | $\begin{gathered} 61 \\ (100 \%) \\ \hline \end{gathered}$ | 3,662 | $\begin{aligned} & 2,343 \\ & (64 \%) \\ & \hline \end{aligned}$ | $\begin{gathered} 28 \\ (1 \%) \\ \hline \end{gathered}$ |
| 175.vpr | 76.8 | $\begin{gathered} 1,006 \\ (13.1 / \mathrm{s}) \end{gathered}$ | $\begin{gathered} 82 \\ (8 \%) \\ \hline \end{gathered}$ | $\begin{array}{\|l\|} \hline 26,783 \\ (83 \%) \\ \hline \end{array}$ | $\begin{aligned} & 2,007 \\ & (71 \%) \\ & \hline \end{aligned}$ | $\begin{gathered} 226 \\ (89 \%) \\ \hline \end{gathered}$ | 13,418 | $\begin{aligned} & 8,983 \\ & (67 \%) \\ & \hline \end{aligned}$ | $\begin{aligned} & 333 \\ & (2 \%) \\ & \hline \end{aligned}$ | 9,436 | $\begin{array}{\|c\|} \hline 1,000 \\ (0.10 / \mathrm{s}) \\ \hline \end{array}$ | $\begin{gathered} 79 \\ (8 \%) \\ \hline \end{gathered}$ | $\begin{aligned} & 26,677 \\ & (83 \%) \\ & \hline \end{aligned}$ | $\begin{aligned} & 2,004 \\ & (70 \%) \\ & \hline \end{aligned}$ | $\begin{gathered} 226 \\ (89 \%) \\ \hline \end{gathered}$ | 13,332 | $\begin{aligned} & 7,199 \\ & (57 \%) \\ & \hline \end{aligned}$ | $\begin{aligned} & 2,428 \\ & (18 \%) \\ & \hline \end{aligned}$ |
| 176.gcc | 3490.2 | $\begin{aligned} & \hline 26,524 \\ & (7.6 / \mathrm{s}) \\ & \hline \end{aligned}$ | $\begin{gathered} 822 \\ (3 \%) \\ \hline \end{gathered}$ | $\begin{array}{\|c} \hline 186,310 \\ (49 \%) \\ \hline \end{array}$ | $\begin{aligned} & \hline 16,104 \\ & (44 \%) \\ & \hline \end{aligned}$ | $\begin{array}{r} 1,239 \\ (65 \%) \\ \hline \end{array}$ | 573,375 | $\begin{gathered} 384,161 \\ (67 \%) \\ \hline \end{gathered}$ | $\begin{gathered} 11,467 \\ (2 \%) \\ \hline \end{gathered}$ | 347,014 | $\begin{array}{\|l\|} \hline 26,647 \\ (0.08 / \mathrm{s}) \\ \hline \end{array}$ | $\begin{gathered} \hline 799 \\ (3 \%) \\ \hline \end{gathered}$ | $\begin{array}{\|c\|} \hline 183,280 \\ (48 \%) \end{array}$ | $\begin{aligned} & 16,098 \\ & (43 \%) \end{aligned}$ | $\begin{array}{\|l\|} \hline 1,221 \\ (64 \%) \\ \hline \end{array}$ | 573,926 | $\begin{array}{\|c} \hline 332,303 \\ (58 \%) \end{array}$ | $\begin{gathered} 63,131 \\ (11 \%) \\ \hline \end{gathered}$ |
| 181.mcf | 8.6 | $\begin{gathered} 144 \\ (16.7 / \mathrm{s}) \end{gathered}$ | $\begin{gathered} 2 \\ (1 \%) \\ \hline \end{gathered}$ | $\begin{array}{\|c\|} \hline 2,977 \\ (100 \%) \\ \hline \end{array}$ | $\begin{array}{\|c\|} \hline 213 \\ (100 \%) \\ \hline \end{array}$ | $\begin{gathered} 24 \\ (100 \%) \\ \hline \end{gathered}$ | 1,718 | $\begin{aligned} & 1,248 \\ & (73 \%) \\ & \hline \end{aligned}$ | $\begin{gathered} 0 \\ (0 \%) \\ \hline \end{gathered}$ | 374 | $\begin{array}{\|c\|} \hline 164 \\ (0.43 / \mathrm{s}) \\ \hline \end{array}$ | $\begin{gathered} 2 \\ (1 \%) \end{gathered}$ | $\begin{aligned} & 2,947 \\ & (99 \%) \\ & \hline \end{aligned}$ | $\begin{gathered} 213 \\ (100 \%) \\ \hline \end{gathered}$ | $\begin{gathered} 24 \\ (100 \%) \\ \hline \end{gathered}$ | 1,487 | $\begin{aligned} & \hline 1,011 \\ & (68 \%) \\ & \hline \end{aligned}$ | $\begin{array}{r} 130 \\ (9 \%) \\ \hline \end{array}$ |
| 186.crafty | 860.3 | $\begin{gathered} \hline 2,753 \\ (3.2 / \mathrm{s}) \\ \hline \end{gathered}$ | $\begin{gathered} 15 \\ (0.5 \%) \end{gathered}$ | $\begin{array}{\|l\|} \hline 40,404 \\ (96 \%) \end{array}$ | $\begin{aligned} & 4,237 \\ & (96 \%) \end{aligned}$ | $\begin{array}{\|c\|} \hline 104 \\ (100 \%) \end{array}$ | 22,437 | $\begin{aligned} & 14,300 \\ & (64 \%) \\ & \hline \end{aligned}$ | $\begin{gathered} 20 \\ (0.08 \%) \end{gathered}$ | 99,764 | $\begin{array}{\|c\|} \hline 2,830 \\ (0.03 / \mathrm{s}) \end{array}$ | $\begin{gathered} 13 \\ (0.4 \%) \end{gathered}$ | $\begin{aligned} & 41,685 \\ & (99 \%) \\ & \hline \end{aligned}$ | $\begin{aligned} & 4,381 \\ & (99 \%) \\ & \hline \end{aligned}$ | $\begin{array}{c\|} \hline 104 \\ (100 \%) \\ \hline \end{array}$ | 22,816 | $\begin{aligned} & 12,092 \\ & (53 \%) \\ & \hline \end{aligned}$ | $\begin{aligned} & 2,749 \\ & (12 \%) \end{aligned}$ |
| 197.parser | 98.2 | $\begin{gathered} 1,590 \\ (16.2 / \mathrm{s}) \end{gathered}$ | $\begin{gathered} 68 \\ (4 \%) \end{gathered}$ | $\begin{aligned} & 22,093 \\ & (90 \%) \\ & \hline \end{aligned}$ | $\begin{aligned} & 2,688 \\ & (92 \%) \\ & \hline \end{aligned}$ | $\begin{array}{\|c\|} \hline 279 \\ (94 \%) \\ \hline \end{array}$ | 9,958 | $\begin{aligned} & \hline 6,664 \\ & (67 \%) \\ & \hline \end{aligned}$ | $\begin{array}{r} 887 \\ (9 \%) \\ \hline \end{array}$ | 6,340 | $\begin{array}{\|c\|} \hline 1,685 \\ (0.27 / \mathrm{s}) \\ \hline \end{array}$ | $\begin{gathered} 69 \\ (4 \%) \\ \hline \end{gathered}$ | $\begin{aligned} & 23,331 \\ & (95 \%) \end{aligned}$ | $\begin{aligned} & 2,799 \\ & (96 \%) \\ & \hline \end{aligned}$ | $\begin{array}{\|c\|} \hline 288 \\ (97 \%) \\ \hline \end{array}$ | 11,740 | $\begin{aligned} & 5,870 \\ & (50 \%) \\ & \hline \end{aligned}$ | $\begin{aligned} & 3,682 \\ & (31 \%) \end{aligned}$ |
| 252.eon | 37.2 | $\begin{gathered} 707 \\ (19.0 / \mathrm{s}) \end{gathered}$ | $\begin{gathered} 27 \\ (4 \%) \\ \hline \end{gathered}$ | $\begin{array}{\|l\|} \hline 28,600 \\ (71 \%) \\ \hline \end{array}$ | $\begin{aligned} & 5,560 \\ & (70 \%) \\ & \hline \end{aligned}$ | $\begin{gathered} 502 \\ (82 \%) \\ \hline \end{gathered}$ | 9,521 | $\begin{aligned} & 4,457 \\ & (47 \%) \\ & \hline \end{aligned}$ | $\begin{aligned} & 142 \\ & (1 \%) \\ & \hline \end{aligned}$ | 4,020 | $\begin{array}{\|c\|} \hline 659 \\ (0.16 / \mathrm{s}) \\ \hline \end{array}$ | $\begin{gathered} 26 \\ (4 \%) \\ \hline \end{gathered}$ | $\begin{array}{\|l} \hline 27,622 \\ (69 \%) \\ \hline \end{array}$ | $\begin{aligned} & 5,413 \\ & (68 \%) \\ & \hline \end{aligned}$ | $\begin{array}{\|c\|} \hline 501 \\ (81 \%) \\ \hline \end{array}$ | 9,121 | $\begin{aligned} & 3,557 \\ & (39 \%) \\ & \hline \end{aligned}$ | $\begin{aligned} & 5,669 \\ & (62 \%) \\ & \hline \end{aligned}$ |
| 253.perlbmk | 1,189 | $\begin{aligned} & \hline 10,318 \\ & (8.7 / \mathrm{s}) \\ & \hline \end{aligned}$ | $\begin{gathered} 508 \\ (5 \%) \\ \hline \end{gathered}$ | $\begin{array}{\|c} \hline 118,135 \\ (88 \%) \\ \hline \end{array}$ | $\begin{array}{\|l\|} \hline 11,600 \\ (90 \%) \\ \hline \end{array}$ | $\begin{gathered} 692 \\ (97 \%) \\ \hline \end{gathered}$ | 66,726 | $\begin{gathered} 28,394 \\ (43 \%) \end{gathered}$ | $\begin{aligned} & 4,001 \\ & (6 \%) \\ & \hline \end{aligned}$ | 176,096 | $\begin{array}{\|c\|} \hline 10,400 \\ (0.06 / \mathrm{s}) \\ \hline \end{array}$ | $\begin{gathered} 502 \\ (4 \%) \\ \hline \end{gathered}$ | $\begin{gathered} 119,467 \\ (89 \%) \\ \hline \end{gathered}$ | $\begin{array}{\|l\|} \hline 11,676 \\ (90 \%) \\ \hline \end{array}$ | $\begin{array}{\|c\|} \hline 696 \\ (97 \%) \\ \hline \end{array}$ | 70,611 | $\begin{aligned} & 24,713 \\ & (35 \%) \\ & \hline \end{aligned}$ | $\begin{aligned} & 18,866 \\ & (27 \%) \\ & \hline \end{aligned}$ |
| 254.gap | 1,054 | $\begin{gathered} 7,754 \\ (7.3 / \mathrm{s}) \end{gathered}$ | $\begin{gathered} 310 \\ (4 \%) \\ \hline \end{gathered}$ | $\begin{gathered} 49,869 \\ (54 \%) \\ \hline \end{gathered}$ | $\begin{aligned} & 4,519 \\ & (50 \%) \\ & \hline \end{aligned}$ | $\begin{gathered} 401 \\ (88 \%) \end{gathered}$ | 38,243 | $\begin{aligned} & 20651 \\ & (54 \%) \\ & \hline \end{aligned}$ | $\begin{array}{r} 3,059 \\ (8 \%) \\ \hline \end{array}$ | 103,458 | $\begin{array}{\|c\|} \hline 7,461 \\ (0.07 / \mathrm{s}) \\ \hline \end{array}$ | $\begin{array}{r} 298 \\ (4 \%) \\ \hline \end{array}$ | $\begin{gathered} 49,920 \\ (54 \%) \\ \hline \end{gathered}$ | $\begin{aligned} & 4,521 \\ & (50 \%) \\ & \hline \end{aligned}$ | $\begin{array}{\|c} \hline 401 \\ (88 \%) \\ \hline \end{array}$ | 38,784 | $\begin{aligned} & 18228 \\ & (47 \%) \\ & \hline \end{aligned}$ | $\begin{array}{r} 6,593 \\ (17 \%) \\ \hline \end{array}$ |
| 255.vortex | 487.0 | $\begin{gathered} \hline 7,232 \\ (14.9 / \mathrm{s}) \\ \hline \end{gathered}$ | $\begin{gathered} 157 \\ (2 \%) \\ \hline \end{gathered}$ | $\begin{array}{\|c} \hline 100,718 \\ (92 \%) \end{array}$ | $\begin{aligned} & 15,513 \\ & (91 \%) \\ & \hline \end{aligned}$ | $\begin{gathered} 577 \\ (92 \%) \\ \hline \end{gathered}$ | 55,205 | $\begin{aligned} & 19,939 \\ & (36 \%) \\ & \hline \end{aligned}$ | $\begin{aligned} & 630 \\ & (1 \%) \\ & \hline \end{aligned}$ | 58,646 | $\begin{array}{\|c\|} \hline 7,223 \\ (0.12 / \mathrm{s}) \\ \hline \end{array}$ | $\begin{gathered} 132 \\ (2 \%) \\ \hline \end{gathered}$ | $\begin{gathered} \hline 100,652 \\ (92 \%) \\ \hline \end{gathered}$ | $\begin{aligned} & \hline 15,489 \\ & (91 \%) \\ & \hline \end{aligned}$ | $\begin{array}{\|c\|} \hline 577 \\ (92 \%) \\ \hline \end{array}$ | 54,977 | $\begin{aligned} & 15,393 \\ & (28 \%) \\ & \hline \end{aligned}$ | $\begin{aligned} & 14,072 \\ & (26 \%) \\ & \hline \end{aligned}$ |
| 256.bzip2 | 16.0 | $\begin{gathered} 249 \\ (15.6 / \mathrm{s}) \end{gathered}$ | $\begin{gathered} 13 \\ (5 \%) \end{gathered}$ | $\begin{aligned} & 6,338 \\ & (92 \%) \end{aligned}$ | $\begin{array}{\|c\|} \hline 545 \\ (94 \%) \end{array}$ | $\begin{array}{\|c\|} \hline 60 \\ (95 \%) \end{array}$ | 2,755 | $\begin{aligned} & 2,375 \\ & (86 \%) \end{aligned}$ | $\begin{gathered} 0 \\ (0 \%) \end{gathered}$ | 842 | $\begin{array}{\|c\|} \hline 258 \\ (0.19 / \mathrm{s}) \end{array}$ | $\begin{gathered} 11 \\ (4 \%) \end{gathered}$ | $\begin{aligned} & 5,179 \\ & (76 \%) \end{aligned}$ | $\begin{gathered} 471 \\ (82 \%) \end{gathered}$ | $\begin{gathered} 53 \\ (84 \%) \end{gathered}$ | 2,434 | $\begin{aligned} & 1,849 \\ & (76 \%) \\ & \hline \end{aligned}$ | $\begin{aligned} & 215 \\ & (9 \%) \end{aligned}$ |
| 300.twolf | 221.4 | $\begin{gathered} 2,972 \\ (13.4 / \mathrm{s}) \end{gathered}$ | $\begin{gathered} 97 \\ (3 \%) \\ \hline \end{gathered}$ | $\begin{aligned} & 52,351 \\ & (91 \%) \\ & \hline \end{aligned}$ | $\begin{aligned} & 3,682 \\ & (86 \%) \\ & \hline \end{aligned}$ | $\begin{array}{\|c\|} \hline 165 \\ (99 \%) \end{array}$ | 24,032 | $\begin{aligned} & 10,333 \\ & (43 \%) \\ & \hline \end{aligned}$ | $\begin{gathered} 528 \\ (2 \%) \\ \hline \end{gathered}$ | 21,308 | $\begin{array}{\|c\|} \hline 2,997 \\ (0.14 / \mathrm{s}) \\ \hline \end{array}$ | $\begin{gathered} 90 \\ (3 \%) \\ \hline \end{gathered}$ | $\begin{aligned} & 52,831 \\ & (92 \%) \end{aligned}$ | $\begin{array}{r} 3,749 \\ (88 \%) \\ \hline \end{array}$ | $\begin{gathered} 165 \\ (99 \%) \end{gathered}$ | 25,664 | $\begin{aligned} & 8,212 \\ & (32 \%) \\ & \hline \end{aligned}$ | $\begin{aligned} & 3,132 \\ & (12 \%) \\ & \hline \end{aligned}$ |
| Average | - | 12.6/s | 3.5\% | 83.8\% | 79.1\% | 91.8\% | - | 60.6\% | 2.6\% | - | 0.15/s | 3.4\% | 82.7\% | 81.0\% | 90.9\% | - | 50.6\% | 19.6\% |

```
int some_func (char *arg1, char *arg2) {
    check_arg(arg1);
    check_arg(arg2);
    do_something(); // do nothing
    ...
} }
void check_arg(char *arg) {
    if (strlen(arg) == 0) exit(-1);
        ...
}
```

Fig. 6: Explaining problem of linear search using gcc.

17~19 show the memory dependence results for PMP and X-Force, respectively.

Observe that X-Force has 6.5 times more mis-typed memory dependences compared to PMP ( $19.6 \%$ versus $2.6 \%$ ), that is, 6.5 X more FPs. In addition, the must-be-true memory dependences reported by X-Force are $10 \%$ fewer than those by PMP. That is, X-Force has $10 \%$ more FNs. The main reason is that X -Force does not trace into library execution such that pointer relations are incomplete. We will use a case study to explain this in the next paragraph. Mis-typed dependences (FPs) in PMP are mostly caused by violations of SDMB. The results are consistent with our analysis in Section III-D. Note that our probabilistic guarantee for SDMB was computed for a pair of accesses, whereas the reported value is the expected value over a large number of pairs.

Case Study. We use 181.mcf as a case study to demonstrate the advantages of PMP over X-Force, as well as over a naive memory pre-planning that fills the pre-allocated region and variables with 0 . To reduce the interference caused by the path exploration algorithm, we use the execution traces of the runs on the provided test cases as the path schemes. That is, we enforce the branch outcomes in a way that strictly follows the traces. The test cases fall into three categories: training, test, and reference, with difference sizes (reference tests are

```
long suspend_impl(..){..
    if (is_valid(arc)) {..
        memcpy(new_arc, arc, 0x40);..
        *(arc->tail) = node1;..
        node2 = *(new_arc->tail);..
    }
}
```

Fig. 7: Explaining FPs and FNs by X-Force using $m c f$.
the largest). We use the memory dependences reported while executing the test cases normally as the ground truth to identify the false positives and false negatives for PMP and X-Force. Since both the forced and unforced executions of a test input follow the same path, the comparison particularly measures the effectiveness of the memory schemes. To be more fair, we only run PMP on a single executor.

The results are shown in Table II. The 2nd and 3rd columns compare the execution speed. Observe that PMP is much faster, consistent with our earlier observation. For the memory dependences, PMP has no FPs or FNs while the naive planning method has some; and X-Force has the largest number of FPs and FNs. The former is because SDMB is violated. The latter is due to the incompleteness of pointer relation tracking (i.e., missing the library part). Note that the numbers of FPs and FNs are smaller compared to the previous experiment as these are results for a small number of runs, without exploring paths.

Consider the code snippet from mcf shown in Figure 7. Variable arc is a buffer that contains many pointer fields. As it is copied to new_arc at line 3, the pointer fields in arc and new_arc are linearly correlated. However, X-Force misses such correlations as it does not trace into memcpy () at line 2. This could lead to missing dependences such as that between lines 4 and 5; and also bogus dependences. For example, the read * (new_arc->tail) at line 5 must falsely depend on some write that happened earlier.


Fig. 8: Overall result of malware analysis.

TABLE II: Experiment with $m c f$.

| Item | Execution Time (s) |  | Memory Dependence |  |  |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | PMP | X-Force | ground | PMP |  |  | Naive |  |  | X-Force |  |  |
|  |  |  |  | found | fp | fn | found | fp | fn | found | fp | fn |
| test | 0.0305 | 1.987 | 1847 | 1847 | 0 | 0 | 1848 | 5 | 4 | 1858 | 28 | 17 |
| train | 0.0348 | 2.578 | 2065 | 2065 | 0 | 0 | 2069 | 13 | 9 | 2088 | 45 | 22 |
| ref | 0.0609 | 4.390 | 2062 | 2062 | 0 | 0 | 2068 | 14 | 8 | 2080 | 37 | 19 |

## C. Malware Analysis

We use 400 malware samples. Half of them are acquired from VirusTotal under an academic license, and the other half fall into the set of malware used in the Padawan project. Note that the authors of Padawan cannot share their samples due to licensing limitations. Hence, we crawled the Internet for these samples based on a set of hash values provided by the Padawan's authors through personal communication. Many samples could not be found and are hence elided. The 400 samples cover up-to-date malware of different families captured from year 2016 to 2018 . We compare the malware analysis result of PMP with that of Cuckoo [2] (a well-known sandbox for automatic malware analysis), Padawan [8] (an academic multi-architecture ELF malware analysis platform), Habo [10] (a commercial malware analysis platform used by VirusTotal for capturing behaviors of ELF malware samples) as well as X-Force [32].

In order to compare our technique with the state-of-the-art anti-evasion measures, we implemented two popular anti-evasion methods [19] (i.e. system time fast-forwarding and anti-virtualization-detection) as extensions to Cuckoo. We name the extended system Cuckoo ${ }^{++}$. Specifically in the first method, we modify the kernel to make the system clock much faster (e.g., 100 times faster), mainly for the following two reasons. First, a malware analysis VM often has a very short uptime since it restarts for each malware execution. As such, advanced malware may check the system uptime to determine the presence of sandbox VM. Second, advanced malware samples often sleep for a period of time before executing their payload (in order to defeat dynamic analysis). In the other method, we intercept file system operations to conceal the artifacts produced by virtual machine (e.g., /sys/class/dmi/id/product_name and /sys/class/dmi/id/sys_vendor).
The detailed comparison results are shown in Appendix C. Note that the malware behaviors of Padawan are provided by its authors. We set up an execution environment similar to Padawan (Ubuntu 16.04 with Linux kernel version 4.4) for

TABLE III: Analysis on malware samples used for case study.

| Case | ID | Cuckoo | Habo | Padawan | Cuckoo++ | X-Force | PMP |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 1 | 031 | 12 | 17 | 12 | 12 | 283 | 301 |
| 2 | 004 | 27 | 29 | 28 | 27 | 32 | 216 |
| 3 | 225 | 49 | 49 | 166 | 165 | 183 | 220 |
| 4 | 309 | 153 | 169 | 292 | 221 | 274 | 705 |

the other tools, including PMP, X-Force, Habo, Cuckoo and Cuckoo ${ }^{++}$, so that the results can be comparable. We set 5 minutes timeout for each malware sample.

Result Summary. Figure 8 presents the overall result of malware analysis. Specifically, the number of unique system call sequences exposed by different tools are show in Figure 8 a. To avoid considering similar system call sequences that have only small differences on argument values as different sequences, we consider sequences that have more than $90 \%$ similarity as identical. As we can see that the executions with anti-evasion measures enabled (i.e., Cuckoo ${ }^{++}$and Padawan) expose more system call sequences than the native executions (i.e., Cuckoo and Habo), but disclose fewer than the forced execution methods (i.e., X-Force and PMP). On average, PMP reports $220 \%, 243 \%, 150 \%, 151 \%$ and $98 \%$ more system call sequences over Cuckoo, Habo, Cuckoo ${ }^{++}$, Padawan and XForce, respectively. Details can be found in Appendix C.

The comparison of execution speed and length of path schemes between PMP and X-Force are shown in Figure 8b and Figure 8c respectively. Note that Cuckoo and Padawan only runs each sample once (instead of multiple executions on different path schemes as force execution tools do). Hence we do not compare their execution speeds and length of path scheme. On average, PMP is 9.8 times faster than X-Force and yields path schemes with the length 1.5 times longer than X-Force. The longer the path scheme, the deeper the code was explored. The second case studies in this subsection show that with the longer path schemes, PMP can expose some malicious behavior in deep program paths that could not be exposed by X-Force.

Case Studies. Next, we use four case studies from different malware families to illustrate the advantages of PMP.
Case1: 1e19b857a5f5a9680555fa9623a88e99. It is a ransom malware that uses UPX packer [11] to pack its malicious payload in order to evade static analysis. Figure 9a shows a constructed code snippet to demonstrate part of its malicious logic. It mmaps a writable and executable memory area (line 2), then unpacks itself (line 3) and transfers control

```
int main(int argc, char **argv) {
    void *code_area = map_exec_write_mem();
    upx_unpack(code_area);
    transfer_control(code_area, argc, argv);
}
    void code_area(int argc, char **argv) {
    if (!is_cmdline_valid(argc, argv)) exit();
    char *action = argv[1], *key = argv[2];
    delete_self();
    if (strcmp(action, encrypt) == 0) {
        for (FILE *file: traverse_directory()) {
            FILE *encrypted_file = encrypt(file, key);
            replace_file(encrypted_file, file);
        }
    }
}
```

(a) simplified code.
mmap ( $0 \times 400000$, ,PROT_EXEC|PROT_READ|PROT_WRITE, )
unlink("/root/Malware/1e19b857a5f5a9680555fa9623a88e99") open ("/etc", O_RDONLY|O_DIRECTORY|O_CLOEXEC)
getdents64(0,)
open("/etc/passwd", O_RDONLY)
open ("/etc/passwd.encrypted", O_WRONLY|O_CREAT, 0666 )
. unlink("/etc/passwd")
(b) captured system call sequence.

Fig. 9: Case 1: the ransom malware sample.
(line 4) to the unpacked payload (lines 7-17). The malicious payload checks the validity of command line parameters (line 8) and deletes itself from the file system (line 10). If the command line parameter specifies the encrypt action, the malware traverses the file system to replace each file with its encrypted copy (lines 13-14).
The comparison of different tools on this malware is shown in the second row of Table III. Triggering payload requires the correct command line parameters. Hence directly running the malware using Cuckoo, Habo, Cuckoo ${ }^{++}$and Padawan fail to expose the malicious behavior. Both X-Force and PMP expose the payload. Figure 9b shows the captured system call sequence. Observe the unlink syscall b that removes the malware itself and the encryption and removal of "/etc/passwd" by syscalls e-g.

Case2: 03cfe768a8b4ffbe0bb0fdef986389dc. It is a bot malware that receives command from a remote server. Figure 10a shows the simplified code of its processing logic. It checks whether a file exists that indicates the right execution environment (line 2) and whether the remote server is connectable (line 4). If both conditions are satisfied, the malware communicates with the remote server. The remote server will validate the identity of the malware by its own communication protocol (lines 4-7). If the validation is successful, a command received from the remote server will be executed on the victim machine (lines 8-9).

The comparison of different tools on this malware is shown in the third row of Table III. The malicious payload of this malware sample is hidden in a deeper path, which requires a much longer path scheme. Figure 10b shows the path scheme enforced by PMP to expose the malicious behaviors. The length is 28 , which is larger than the longest path scheme that is enforced by X-Force within the 5 minutes limit. These forced branches are to get through the ID validation protocol.

```
int main(int argc, char **argv) {
    if (!files_exist("/tmp/ReV1112")) exit(0);
    if (!connectable("ka3ek.com")) exit(0);
    Info *info = get_system_info();
    Greet *greet = get_validation(info);
    Reply *reply = compute_reply(greet);
    Cmd *cmd = get_command(reply);
    if (!cmd) exit(0);
    execute_cmd(cmd);
}
```

(a) simplified code.

| 40492b:T | 404 aec : T | 404e07:T | 401f3f: F | 401ee3:T |
| :---: | :---: | :---: | :---: | :---: |
| 404fdc: F | 404fea:T | 405118: F | 40513a:F | 405144: F |
| 40517b: F | 40517f: F | 40523e: F | 405254:T | 40523e: F |
| 405254:T | 40523e:F | 405254:T | 40523e:F | 405254:T |
| 40523 e : F | 405254: F | 4044be:T | 4044e9:F | 40454b: F |
| 404565:T | 404596:T | 404794: F |  |  |

(b) path scheme.

Fig. 10: Case 2: the bot malware sample.
Case3: 14b788d4c5556fe98bd767cd10ac53ca. It is an enhanced variant of Mirai, which is equipped with a timebased cloaking technique. Figure 11 shows a simplified version of its code snippet. At line 4, it checks whether the system uptime is short, which indicates a potential analysis environment. If the system uptime is long enough, it checks whether there exists any initialization script in the "/etc/init.d" directory (line $8)^{2}$. If both conditions are satisfied, the malware sample adds itself to an initialization script for launching at system reboot.
Cuckoo and Habo cannot expose the aforementioned behaviors. Cuckoo ${ }^{++}$and Padawan can expose the traversal of the "/etc/init.d" directory (line 6), by passing though the uptime check via fast-forwarding system time and using a long-running VM snapshot, respectively. However, they cannot expose the modification of initialization script (line 9), due to the failure of the initialization script check, as the default OS environment does not have any initialization script. PMP and X-Force can expose both behaviors by forcing the branch results.

Case4: 8ab6624385a7504e1387683b04c5f97a. This is a sniffer equipped with a vm-detection-based cloaking technique. Figure 12 shows a simplified version of its code snippet. If a VM environment is detected, the malware sample deletes itself and exits (lines 2-3). Otherwise, it enters a sniffing loop, which randomly selects an intranet IP address and a known vulnerability and checks whether the host with the IP contains the vulnerability (lines 5-7). If so, the information about the vulnerable host is sent to the server and the payload is sent to the vulnerable host (lines 8-9).

Cuckoo and Habo cannot expose the aforementioned behaviors. Cuckoo ${ }^{++}$and Padawan can expose the network communication to the selected IP address, since they are enhanced to conceal VM-generated artifacts. However, they cannot expose sending the vulnerable host information and payload, since the analysis environment is often offline and there may not exist a vulnerable host on the intranet. PMP can expose both behaviors. X-Force can expose both in theory

[^1]```
int main(int argc, char **argv) {
    struct sysinfo info;
    sysinfo(&info);
    if (info.uptime < 128) exit(0);
    DIR *dir = opendir("/etc/init.d");
    while (struct dirent *ent = readdir(dir)) {
        char name = ent->d_name;
        if (name[0] == 'S' && is_num(name[1]))
            add_to_init_script("/etc/init.d/S99");
    }
}
```

Fig. 11: Case 3: the enhanced variant of Mirai.
but fails within the timeout limit due to its substantially larger runtime cost.

## D. Time Distribution

We measure the runtime overhead of different components. The distribution is shown in Appendix B. As we can see that most of the time ( $84 \%$ ) is spent on code execution, while only $13 \%$ and $3 \%$ of time are spent on memory preplanning and path exploration, respectively. In memory preplanning, $2 \%, 5 \%, 69 \%$ and $24 \%$ of time are spent on PAMA preparation, initialization of global variables, local variables and heap variables. Observe that PAMA preparation takes very little time as most work is done offline.

## V. Related Work

Forced Execution. Most related to our work is X-Force [32]. The technical differences between the two were discussed in the introduction section. As shown by our results, PMP is 84 times faster than X-Force, has 6.5X, and $10 \%$ fewer FPs and FNs of dependencies, respectively, and exposes $98 \%$ more payload in malware analysis. Following X-Force, other forced-execution tools are developed for different platforms, including Android runtime [33] and JavaScript engine [25], [21]. Compared to these techniques, PMP targets x86 binaries and addresses the low level invalid memory operations. Additionally, PMP is based on novel probabilistic memory pre-planning instead of demand driven recovery.
Memory Randomization. Memory randomization has been leveraged for different purposes, such as reducing vulnerability to heap-based security attacks through randomizing the base address of heap regions [14] and randomly padding allocation requests [15]. DieHard [13] tolerates memory errors in applications written in unsafe languages through replication and randomization. It features a randomized memory manager that randomizes objects in a "conceptual heap" whose size is a multiple of the maximum real size allowed. PMP shares a similar probabilistic flavor to DieHard. The difference lies in that PMP pre-plans the memory by pre-allocation and filling the pre-allocated space and variables with crafted values. In addition, PMP aims to survive memory exceptions caused by forced-execution whereas DieHard is for regular execution.

Malware Analysis. The proliferation of Malware in the past decades provide strong motivation for research on detecting, analyzing and preventing malware, on various platforms such as Windows [16], [23], Linux [19], [20], as well as Web

```
char *data = read_file("/sys/class/dmi/id/product_name");
if (contains(data, "VirtualBox", "VMware"))
        remove_self_and_exit();
while (1) {
        char *ip = select_intranet_ip(ip_list);
        char *vuln = select_known_vuln(vuln_list);
        if (connect_and_check(ip, vuln)) {
        send_info_to_server(ip, vuln);
        send_payload(ip, vuln);
    }
}
```

Fig. 12: Case 4: the sniffer malware sample.
browsers [24], [22]. Traditional malware analysis fall into two categories: signature-based scanning and behavioral-based analysis. The former [12], [28] detects malware by matching extracted features with known signatures. Although commonly used by anti-malware industry, signature-based approaches are susceptible to evasion through obfuscation. To address this, behavioral-based approaches [34], [26], [17] execute a subject program and monitor its behavior to observe any malicious behavior. However, traditional behavioral-based approaches are limited to observing code that is actually executed.

Anti-targeted Evasion. Modern sophisticated malware samples are equipped with various cloaking techniques (e.g., stalling loop [27] and VM detection [6]) to evade detection. To fight against evasion, unpacking techniques [18], [29] are applied to enhance signature-based scanning, and dynamic anti-evasion methods [26], [30] are developed to hide dynamic features of analysis environment such as execution time and file system artifacts. These techniques are very effective for known targeted evasion methods. Compared to these techniques, PMP is more general. More importantly, PMP and forced execution type of techniques allow exposing payload guarded by complex conditions that are irrelevant to cloaking.

## VI. Conclusion

We develop a lightweight and practical force-execution technique that features a novel memory pre-planning method. Before execution, the pre-planning stage pre-allocates a memory region and initializes it (and also variables in the subject binary) with carefully crafted values in a random fashion. As a result, our technique provides strong probabilistic guarantees to avoid crashes and state corruptions. We apply the prototype PMP to SPEC2000 and 400 recent malware samples. Our results show that PMP is substantially more efficient and effective than the state-of-the-art.

## ACKNOWLEDGMENT

The authors would like to thank the anonymous reviewers and Dr. William Robertson (the PC contact) for their constructive comments. Also, the authors would like to express their thanks to VirusTotal and the authors of Padawan for their kindness in sharing malware samples and the analysis results. The Purdue authors were supported in part by DARPA FA8650-15-C-7562, NSF 1748764, 1901242 and 1910300, ONR N000141410468 and N000141712947, and Sandia National Lab under award 1701331. The UVA author was supported in part by NSF 1850392.

## REFERENCES

[1] Clickless powerpoint malware installs when users hover over a link. https://blog.barkly.com/powerpoint-malware-installs-when-users-hover-over-a-link.
[2] Cuckoo. https://cuckoosandbox.org/.
[3] Cybersecurity statistics. https://blog.alertlogic.com/10-must-know-2018-cybersecurity-statistics/.
[4] Evil clone attack. https://gbhackers.com/evil-clone-attack-legitimate-pdf-software.
[5] Fileless malware. https://www.cybereason.com/blog/fileless-malware.
[6] Linux anti-vm. https://www.ekkosec.com/blog/2018/3/15/linux-anti-vm-how-does-linux-malware-detect-running-in-a-virtual-machine-.
[7] Mirai malware. https://en.wikipedia.org/wiki/Mirai_(malware).
[8] Padawan. https://padawan.s3.eurecom.fr/about.
[9] Qemu user emulation. https://wiki.debian.org/QemuUserEmulation.
[10] Tencent habo. https://blog.virustotal.com/2017/11/malware-analysis-sandbox-aggregation.html.
[11] Upx. https://upx.github.io/.
[12] Virustotal. https://www.virustotal.com/gui/home/upload.
[13] Emery D. Berger and Benjamin G. Zorn. Diehard: Probabilistic memory safety for unsafe languages. In Proceedings of the 27th ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI '06. ACM, 2006.
[14] Sandeep Bhatkar, Daniel C. DuVarney, and R. Sekar. Address obfuscation: An efficient approach to combat a board range of memory error exploits. In Proceedings of the 12th Conference on USENIX Security Symposium - Volume 12, SSYM'03. USENIX Association, 2003.
[15] Sandeep Bhatkar, R. Sekar, and Daniel C. DuVarney. Efficient techniques for comprehensive protection from memory error exploits. In Proceedings of the 14th Conference on USENIX Security Symposium Volume 14, SSYM'05. USENIX Association, 2005.
[16] Leyla Bilge, Davide Balzarotti, William Robertson, Engin Kirda, and Christopher Kruegel. Disclosure: detecting botnet command and control servers through large-scale netflow analysis. In Proceedings of the 28th Annual Computer Security Applications Conference. ACM, 2012.
[17] Ahmet Salih Buyukkayhan, Alina Oprea, Zhou Li, and William Robertson. Lens on the endpoint: Hunting for malicious software through endpoint data analysis. In International Symposium on Research in Attacks, Intrusions, and Defenses. Springer, 2017.
[18] Binlin Cheng, Jiang Ming, Jianmin Fu, Guojun Peng, Ting Chen, Xiaosong Zhang, and Jean-Yves Marion. Towards paving the way for large-scale windows malware analysis: generic binary unpacking with orders-of-magnitude performance boost. In Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security. ACM, 2018.
[19] Emanuele Cozzi, Mariano Graziano, Yanick Fratantonio, and Davide Balzarotti. Understanding linux malware. In Proceedings of the 39th IEEE Symposium on Security and Privacy, 2018.
[20] Yanick Fratantonio, Antonio Bianchi, William Robertson, Engin Kirda, Christopher Kruegel, and Giovanni Vigna. Triggerscope: Towards detecting logic bombs in android applications. In 2016 IEEE symposium on security and privacy (SP). IEEE, 2016.
[21] Xunchao Hu, Yao Cheng, Yue Duan, Andrew Henderson, and Heng Yin. Jsforce: A forced execution engine formalicious javascript detection. In Xiaodong Lin, Ali Ghorbani, Kui Ren, Sencun Zhu, and Aiqing Zhang, editors, Security and Privacy in Communication Networks. Springer International Publishing, 2018.
[22] Alexandros Kapravelos, Chris Grier, Neha Chachra, Christopher Kruegel, Giovanni Vigna, and Vern Paxson. Hulk: Eliciting malicious behavior in browser extensions. In 23rd \{USENIX\} Security Symposium ( $\{$ USENIX $\}$ Security 14), 2014.
[23] Amin Kharraz, William Robertson, Davide Balzarotti, Leyla Bilge, and Engin Kirda. Cutting the gordian knot: A look under the hood of ransomware attacks. In International Conference on Detection of Intrusions and Malware, and Vulnerability Assessment. Springer, 2015.
[24] Amin Kharraz, William Robertson, and Engin Kirda. Surveylance: automatically detecting online survey scams. In 2018 IEEE Symposium on Security and Privacy (SP). IEEE, 2018.
[25] Kyungtae Kim, I Luk Kim, Chung Hwan Kim, Yonghwi Kwon, Yunhui Zheng, Xiangyu Zhang, and Dongyan Xu. J-force: Forced execution on javascript. In Proceedings of the 26th International Conference on World Wide Web, WWW '17. International World Wide Web Conferences Steering Committee, 2017.
[26] Clemens Kolbitsch, Paolo Milani Comparetti, Christopher Kruegel, Engin Kirda, Xiaoyong Zhou, and Xiaofeng Wang. Effective and efficient malware detection at the end host. In USENIX 2009, 18th Usenix Security Symposium, 2009.
[27] Clemens Kolbitsch, Engin Kirda, and Christopher Kruegel. The power of procrastination: detection and mitigation of execution-stalling malicious code. In Proceedings of the 18th ACM conference on Computer and communications security. ACM, 2011.
[28] Christopher Kruegel, Engin Kirda, Darren Mutz, William Robertson, and Giovanni Vigna. Polymorphic worm detection using structural information of executables. In International Workshop on Recent Advances in Intrusion Detection. Springer, 2005.
[29] Lorenzo Martignoni, Mihai Christodorescu, and Somesh Jha. Omniunpack: Fast, generic, and safe unpacking of malware. In 23rd Annual Computer Security Applications Conference (ACSAC 2007), 2007.
[30] Kirti Mathur and Saroj Hiranwal. A survey on techniques in detection and analyzing malware executables. International Journal of Advanced Research in Computer Science and Software Engineering, 3(4), 2013.
[31] Tipp Moseley, Dirk Grunwald, Daniel A Connors, Ram Ramanujam, Vasanth Tovinkere, and Ramesh Peri. Loopprof: Dynamic techniques for loop detection and profiling. In Proceedings of the 2006 Workshop on Binary Instrumentation and Applications (WBIA), 2006.
[32] Fei Peng, Zhui Deng, Xiangyu Zhang, Dongyan Xu, Zhiqiang Lin, and Zhendong Su. X-force: Force-executing binary programs for security applications. In Proceedings of the 23rd USENIX Security Symposium, 2014.
[33] Zhenhao Tang, Juan Zhai, Minxue Pan, Yousra Aafer, Shiqing Ma, Xiangyu Zhang, and Jianhua Zhao. Dual-force: Understanding webview malware via cross-language forced execution. In Proceedings of the 33rd ACM/IEEE International Conference on Automated Software Engineering, ASE 2018. ACM, 2018.
[34] Heng Yin, Dawn Song, Manuel Egele, Christopher Kruegel, and Engin Kirda. Panorama: Capturing system-wide information flow for malware detection and analysis. In Proceedings of the 14th ACM Conference on Computer and Communications Security, CCS '07. ACM, 2007.

## Appendix

## A. Spec2000 Benchmark

| Benchmark | source lines | binary size | \# insn | \# block | \# func |
| :---: | :---: | :---: | :---: | :---: | :---: |
| 164.gzip | 8,643 | 143,760 | 7,650 | 707 | 61 |
| 175.vpr | 17,760 | 435,888 | 32,218 | 2,845 | 255 |
| 176.gcc | 230,532 | $4,709,664$ | 378,261 | 36,931 | 1,899 |
| 181.mcf | 2,451 | 62,968 | 2,977 | 213 | 24 |
| 186.crafty | 21,195 | 517,952 | 42,084 | 4,433 | 104 |
| 197.parser | 11,421 | 367,384 | 24,584 | 2,911 | 297 |
| 252.eon | 41,188 | $3,423,984$ | 40,119 | 7,963 | 615 |
| 253.perlbmk | 87,070 | $1,904,632$ | 133,755 | 12,933 | 717 |
| 254.gap | 71,461 | $1,702,848$ | 91,608 | 9,020 | 458 |
| 255.vortex | 67,257 | $1,793,360$ | 109,739 | 16,970 | 624 |
| 256.bzip2 | 4,675 | 108,872 | 6,859 | 577 | 63 |
| 300.twolf | 20,500 | 753,544 | 57,460 | 4,280 | 167 |

## B. Time Distribution



## C. Details of Malware Analysis Result

|  | Cuckoo | Habo | Padawan | Cuckoo++ | X-Force | PMP |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| Avg. | 41.65 | 38.88 | 53.15 | 53.28 | 67.40 | 133.36 |







[^0]:    ${ }^{1}$ Eight is an empirical choice and works well in our evaluation. The number and the distribution of canaries are configurable.

[^1]:    ${ }^{2}$ An initialization script has a file name that starts with ' S ', followed by a number indicating the priority.

