

BISTRO: Binary Component Extraction and Embedding for Software Security Applications

Zhui Deng, Xiangyu Zhang, and Dongyan Xu

Department of Computer Science and CERIAS, Purdue University, West Lafayette, IN 47907
{deng14, xyzhang, dxu}@cs.purdue.edu

Abstract. In software security and malware analysis, researchers often need to directly manipulate binary program – benign or malicious – without source code. A useful pair of binary manipulation primitives are binary functional component *extraction* and *embedding*, for extracting a functional component from a binary program and for embedding a functional component in a binary program, respectively. Such primitives are applicable to a wide range of security scenarios such as legacy program hardening, binary semantic patching, and malware function analysis. Unfortunately, existing binary rewriting techniques are inadequate to support binary function carving and embedding. In this paper, we present BISTRO, a system that supports these primitives *without* symbolic information, relocation information, or compiler support. BISTRO preserves functional correctness of both the extracted functional component and the stretched binary program (with the component embedded) by patching them in a systematic fashion. We have implemented an IDA Pro-based prototype of BISTRO and evaluated it using real-world Windows software. Our results show the effectiveness of BISTRO, with each stretched binary incurring low time and space overhead. Furthermore, we demonstrate BISTRO’s capabilities in various security applications.

1 Introduction

In software security and malware analysis, researchers often need to manipulate binary code – benign or malicious – without source code and symbolic information. One pair of complementary binary manipulation primitives is to (1) extract a re-usable functional component from a binary program and (2) embed a value-added functional component in an existing binary program. We call these primitives *binary component extraction* and *embedding*. The primitives are useful in a wide range of software security scenarios. In *security hardening of legacy binaries*, binary component embedding enables the retrofitting of legacy or close-source software with a third-party functional component that performs a value-added security function such as access control. In *binary semantic patching*, binary programs from different vendors may leverage the same functional component. Suppose one vendor identifies a vulnerability in such a component and releases a patched version for its own program; whereas other vendors are not aware of the vulnerability or have not patched their products. We can apply binary component extraction to carve out the patched component from a patched program and replace the vulnerable version of the same component in an un-patched program using binary component embedding. In *malware analysis*, binary component extraction and embedding

supports “plug and play” of malicious functions extracted from malware captured in the wild. One can even “stitch” multiple extracted malware functions to compose a new piece of malware – a capability that might help enable strategic cyber defence.

Binary component extraction and embedding poses significant challenges. Brute force extraction and insertion of binary functions will most likely fail. Instead, both the extracted component and the target binary program need to be carefully transformed. For example, instructions in the target binary need to be shifted to create space for the embedded function; when a function is extracted from its origin binary, the instructions in it need to be re-positioned and re-packaged; accesses to global variables need to be re-positioned; function pointers need to be properly handled; and indirect jumps/calls need to have their target addresses recalculated. These problems are especially challenging when the binary component or the target binary program is *not relocatable*, which is often the case when dealing with legacy or malware binaries.

Despite advances in binary instrumentation and rewriting, existing techniques are inadequate to address the binary component extraction and embedding challenges. Dynamic binary instrumentation tools such as PIN, Valgrind, DynamoRIO [2] and QEMU perform instrumentation only when a binary program is executed on their infrastructures. They do not generate an instrumented, stand-alone binary for production runs. Static binary rewriting tools such as Diablo [4], Alto [5], Vulcan [30], and Atom [7] can generate instrumented, stand-alone binaries. However, they require symbolic information or that the binaries be generated by special compilers.

More lightweight techniques exist that do not require symbolic information or special compilers [8–13]. Among them, some create *trampolines* at the end of a target binary in which instrumentation is placed and then use control flow *detours* to access the trampolines [8–10]. The others duplicate the body of a target binary program in its virtual memory space and only the replica is instrumented. The original binary body is retained in its original position to provide a kind of control flow forwarding mechanism [11–13]. However, none of these techniques supports extraction of binary component or implanting an extracted component in another binary. Many of them cause substantial space/performance overhead. To the best of our knowledge, none of them has been successfully applied to large-scale Windows applications or kernel code. A more detailed comparison is presented in our technical report [1].

Recently, researchers proposed approaches that focus on identification, extraction and reuse of components from binaries. Inspector Gadget [29] performs dynamic slicing to identify and extract components from malware. The extracted component might have incomplete code path coverage due to the limitation of dynamic analysis. BCR [16] adopts a combination of static and dynamic approach to extract a function from a binary. However, it uses labels to represent jump/call targets, thus does not preserve the semantic of indirect jumps/calls. ROC [23] uses dynamic slicing to identify reusable functional components in a binary but does not extract them. These approaches do not aim to reuse extracted components for enhancing legacy binaries. Moreover, they cannot extract components from non-executable binaries (e.g., malware corpse) due to the use of dynamic analysis.

In this paper, we present BISTRO, a systematic approach to binary functional component extraction and embedding. BISTRO automatically performs the following:

(1) extracting a functional component, with its instructions and data section entries non-contiguously located in the virtual address space, from an original binary and (2) embedding a binary component of any size at any user-specified location in a target binary, without requiring symbolic information, relocation information, or compiler support. For both extraction and embedding, BISTRO preserves the functionalities of the target binary program and the extracted component by accurately patching them. BISTRO performs extraction and embedding efficiently and the “stretched” target binary after embedding only incurs small time and space overhead.

We have developed a prototype of BISTRO as an IDA-Pro [21] plugin. We have conducted extensive evaluation and case studies using real-world Windows applications (e.g., Firefox and Adobe Reader), kernel drivers, and malware. Our evaluation (Section 6) indicates BISTRO’s efficiency and precision in patching extracted components and target binaries. Moreover, the stretched target binary incurs small performance overhead (1.9% on average) and space overhead (10.9% on average). We have applied BISTRO to the following usage cases: (1) We carve out patched components from a binary and use them to replace their vulnerable versions in other applications, achieving binary semantic patching (Section 6.2); (2) We stitch malicious functions from an un-executable Conficker worm [14] sample and compose a new, executable malware (Section 6.3); and (3) We demonstrate the realistic threat of *trojan-ed* kernel drivers with malicious rootkit functions embedded in benign driver – using real-world drivers and rootkits¹.

2 Overview and Assumptions

An overview of BISTRO is shown in Figure 1. BISTRO has two key components: *binary extractor* and *binary stretcher*.




Fig. 1. Overview of BISTRO

The binary extractor is responsible for extracting a designated functional component c from an original binary Q . c includes both code and data of the functional component. The extractor does so by removing the unwanted code and data from Q and then collapsing the remaining data and code into a re-usable component c that occupies a contiguous virtual address region. More importantly, the instructions in c are properly patched for repositioning. We note that c can either be called as a library function or be embedded directly in another binary program.

The binary stretcher is responsible for stretching the target binary P to make “room” (holes in its address space) to embed a function component. As shown in Figure 1, the stretcher takes the target binary P and the to-be-embedded component c as input; stretches P , and patches the code in P to allow the embedding of c . The output of the stretcher is a “stretched” binary $P' = P + c$ ready for execution.

¹ Due to lack of space, Case (3) is presented in our technical report [1].

Summary of Enabling Techniques. Both the binary extractor and stretcher are based on the same binary stretching algorithm (Section 3). The overarching idea is to shift instructions for creating space (stretcher) or squeezing out unwanted space (extractor). The algorithm focuses on patching the control transfer and global data reference instructions by precisely computing the adjusted offsets. For instance, if a component with size $|c| = n$ is inserted, all the original instructions following the insertion point will be shifted by n bytes, and control transfers to any of the shifted instructions need to be incremented by n .

To address the challenge of handling indirect calls and call back functions invoked by external libraries, we develop another algorithm (Section 4.1) that stretches a subject binary at the original entries of functions that are potential targets of indirect calls, creating small holes (usually a few bytes) to hold a long jump instruction to forward any calls to those functions to their shifted locations. These holes must not be shifted by any stretching/shrinking operations. They always stay in their original positions and thus are called “*anchors*”. Our algorithm precisely takes into account these anchors when performing stretching/shrinking. To handle indirect jumps, we leverage an efficient perfect hashing scheme to translate jump targets dynamically. These techniques are used to patch indirect jumps/calls in both extracted components and target binaries.

Assumptions. We make the following assumptions (hence stating the *non-goals* of BISTRO): (1) The user, not BISTRO, will predetermine the semantic appropriateness of embedding component c in target program P . Furthermore, he/she will decide the specific location to insert c . This can be practically done by performing reverse engineering on P . For example, to harden P with some security policy enforcement mechanism based on control flow [6], the user can reconstruct the control flow graph of P , collect its dominance and post-dominance information, and decide proper locations to insert c . (2) The identification of c in the original binary Q is done a priori through manual or automated techniques, such as Inspector Gadget [29], binary slicing [15], binary differencing [31], and BCR [16]. While we will present our experience with functional component identification in our case studies (Section 6), the identification technique itself is *outside* the scope of this paper. (3) Binaries can be properly disassembled (e.g., by IDA-Pro) *before* being passed to BISTRO. This assumption is supported by the large number of real-world, off-the-shelf binaries in our experiments. Although we currently do not handle obfuscated or self-modifying binaries, we note that, in addition to IDA-Pro, other conservative disassembling [13, 35] and unpacking [34] tools can also be used as the pre-processor of BISTRO to handle more sophisticated binaries.

3 Basic Algorithm for Binary Extraction/Stretching

In this section, we present the basic algorithm (Algorithm 1) executed by both the binary extractor and stretcher of BISTRO. For the time being, we assume (1) there is no indirect control transfer and (2) global data is directly referenced in an instruction using its address. The algorithm takes the subject binary and a list of virtual address intervals called *snippets* representing (1) the holes to be created in the binary in the case of stretching *or* (2) the unwanted instruction/data blocks in the case of shrinking

(extraction). First, for each byte in the binary, the algorithm computes a mapping between its original index in the binary and its corresponding index after the snippets are inserted/removed. After that, the algorithm patches address operands in control transfer and global data reference instructions, and copy each byte to its mapped location according to the mapping.

Practical Challenges. To make BISTRO work for real-world large-scale software, we still need to overcome a number of practical challenges *not* addressed by Algorithm 1. Solutions to these challenges will be presented in the next few sections.

- The target of an indirect control transfer instruction (e.g., `call eax`) is computed during execution and takes different values depending on the execution path. Such an instruction cannot be patched by Algorithm 1.

Algorithm 1. Basic binary stretching/shrinking algorithm

Input: P – the subject binary; it has $size$ and $base_addr$ fields to represent its size when loaded into memory and base loading address, respectively.

M – a list of address intervals represent code/data to be inserted/removed, sorted increasingly by their location; each interval has $addr$, len and $type$ fields, denoting the location, size and type respectively. Type “INSERT” means inserting right before $addr$; “REMOVE” means the block starting at $addr$ is to be removed.

Output: P' – the stretched/shrunk binary.

```

1: function BASICSTRETCHING( $P$ ,  $M$ )
2:    $map \leftarrow$  ComputeMapping( $P$ ,  $M$ )
3:    $P' \leftarrow$  PatchTarget( $P$ ,  $map$ )
4: end function

5: function COMPUTEMAPPING( $P$ ,  $M$ )
6:    $offset \leftarrow 0$ 
7:    $m \leftarrow M.begin()$ 
8:   for  $i \leftarrow 0$  to  $P.size$  do
9:     if  $m.addr == P.base\_addr + i$  then
10:      if  $m.type == INSERT$  then
11:         $offset \leftarrow offset + m.len$ 
12:      else if  $m.type == REMOVE$  then
13:         $offset \leftarrow offset - m.len$ 
14:         $i \leftarrow i + m.len$ 
15:      end if
16:       $m \leftarrow M.next()$ 
17:    end if
18:     $map[i] \leftarrow i + offset$ 
19:  end for
20: return  $map$ 
21: end function

22: function PATCHTARGET( $P$ ,  $map$ )
23:    $P' \leftarrow \{nop, nop, \dots, nop\}$ 
24:   for  $i \leftarrow 0$  to  $P.size$  do
25:     if  $map[i] \neq \perp$  then
26:       if  $P[i]$  is instruction then
27:          $ins \leftarrow P[i]$ 
28:         for each data address operand  $op$  in  $ins$  do
29:            $target \leftarrow op.addr - P.base\_addr$ 
30:            $off \leftarrow map[target] - target$ 
31:            $op.addr \leftarrow op.addr + off$ 
32:         end for
33:         if  $ins$  is near call/jump then
34:            $target \leftarrow i + ins.len + ins.target$ 
35:            $off \leftarrow map[target] - target$ 
36:            $off' \leftarrow map[i + ins.len] - (i + ins.len)$ 
37:            $ins.target \leftarrow ins.target + off - off'$ 
38:         else if  $ins$  is far call/jump then
39:            $target \leftarrow ins.target - P.base\_addr$ 
40:            $off \leftarrow map[target] - target$ 
41:            $ins.target \leftarrow ins.target + off$ 
42:         end if
43:          $P'[map[i]] \leftarrow ins$ 
44:       else if  $P[i]$  is data then
45:          $P'[map[i]] \leftarrow P[i]$ 
46:       end if
47:     end if
48:   end for
49: return  $P'$ 
50: end function

```

- Function pointers may be present in data or in an instruction as an immediate operand. These function pointers might be passed as parameters to external libraries as callback functions. If a function is relocated due to stretching, the external library will call back to a wrong address. All these have to be properly handled to ensure correctness of binary stretching/shrinking.
- Accesses to global data may be via data pointers (e.g., `mov ebx, ptr_data; mov eax, [ebx+4]`). The addresses of data are not known until runtime. These instructions cannot be patched using Algorithm 1 either.

4 Handling Indirect Control Transfer

Handling indirect jumps and calls is one of the key challenges in the design of BISTRO. The difficulty is that the jump/call target cannot be known statically and thus is hard to patch. To understand the challenge, consider the example in Figure 2. On the left, there are three objects that are connected via pointers, with two of type B and one of type A. On the right, part of function `f00()` is presented. The function takes two parameters stored in `eax` and `ebx` denoting pointer values. These two pointers may be aliased to each other. If so, `ecx` at `0x4302B2` gets the value `0x400340` defined at `0x4302A0`, and then eventually the call instruction at `0x4302BD` acquires the function pointer `0x444142`. However, if the two pointer parameters are not aliased, the call instruction may get a completely different target, making statically patching it difficult.




Fig. 2. An example showing indirect call handling in binary stretching/extraction.

A naive solution is to identify and patch any constant value in the binary that appears to be a jump/call target. But this is not safe as such values may not be jump/call targets. Notice in the example, there is a null-terminated string “BAD” at address `0x40040A`. With the little endian representation in x86, this string has the same binary value as the function pointer at `0x400300`. Without type information, it is impossible to know whether the value is a string or function pointer. Failure to identify and patch a function pointer leads to broken control-flow, changing the semantics of the target binary. Misclassifying a string as a function pointer leads to undesirable changes to data. While it is plausible to leverage recent advances in binary type inference to type constants in a binary [17–20], the involvement of aliasing as in the example makes such analysis very difficult. In fact, IDA-Pro [21] failed to recognize the function pointer for this case.

If a binary has a relocation table and it does not perform any address space layout self-management such as through a packer, the relocation table will provide the positions of all constant values that are jump/call targets for BISTRO to patch them, thus lead to a sound and complete solution to binary stretching/shrinking. However, relocation table may be absent or contain bogus entries in legacy and malware binaries. Hence, for the rest of the paper, *we do not assume the presence of relocation tables in our design and evaluation*. Next, we describe how to handle indirect calls in Section 4.1 and indirect jumps in Section 4.2.

4.1 Handling Indirect Calls

Indirect calls are very common in modern binaries to leverage the flexibility of function pointers. We have discussed the difficulty of handling function pointers at the beginning

of Section 4. In fact, there is a more challenging situation, in which a binary may pass its function addresses to external library functions which call back the provided functions (e.g., a user function *cmp()* is provided as a parameter to an external library function *qsort()*). In this case, if a function entry has changed due to stretching or shrinking, its invocation sites are outside the body of the binary and thus beyond our control. It is difficult to patch call back function pointer parameters before they are passed on to libraries for two reasons. First, a function pointer might not directly appear as a parameter. It could be a member of a structure passed to an external library. It may even require several layers of pointer indirection to access its value. Patching that is challenging. Second, for many external library functions, we cannot assume the availability of their prototype definitions, it is hence difficult to know their parameter types.

To handle indirect calls including call back functions, we propose to stretch the target binary to make small holes at the entry point of each function that may be an indirect call target. These holes are called *anchors*; they should not be moved during stretching/shrinking. Inside an anchor, we place a jump instruction that jumps to its mapped new address in the stretched/shrunk binary, which is the new entry of the function. As such, we do not need to identify or patch any function pointers in the binary.

Since an anchor must be placed at a fixed address in the stretched binary, it could coincide with instructions that get shifted to that address. To ensure correctness, we put a jump right before an anchor to jump over it. We call the jump the *prefix* of an anchor.




Fig. 3. Stretching with Anchors. The shaded area in (b) is the 40-byte snippet inserted.

Consider the example in Figure 3(a). The call-back function *cmp()* is invoked inside *qsort()*. The entry address of function *cmp()* in the original binary is 0x400120. When we stretch without anchors (Figure 3(b)) in function *qsort()*, the indirect call to *cmp()* at 0xAF8614 will incorrectly go to 0x400120 in the shaded area. When we stretch with anchors (Figure 3(c)), an anchor containing the jump instruction will be placed at 0x400120. Any indirect call that goes to the original entry address of *cmp()*,

0x400120, will be redirected to the actual function body at the new entry address. The jump instruction preceding 0x400120 is its prefix.

Anchor-Based Algorithm. With the presence of anchors, fixing control flow transfer instructions becomes more challenging than in Algorithm 1. We hence devise a new algorithm (Algorithm 2). The idea is to divide the stretching/shrinking operation into two phases. In phase one, the subject binary program is stretched/shrunk using Algorithm 1 to create space for the inserted snippets or removed blocks. Then the stretched/shrunk binary is further stretched to insert anchors using a similar procedure. Separating the two phases substantially simplifies the interference of anchors.

Algorithm 2. Anchor-based stretching algorithm.

Input: P – the subject binary; it has $size$ and $base_addr$ fields to represent its size when loaded into memory and base loading address, respectively.
 M – a list of code/data snippets to be inserted/removed, sorted increasingly by their location; each snippet has $addr$, len and $type$ fields, denoting the location, size and type respectively.
 A – a list of anchors to be placed, sorted increasingly by their location; each anchor has $addr$ and len fields, denoting the location and the content size, respectively.

Output: $anchor_map$ – the mapping between the indices after placing snippets and their corresponding indices after anchors are placed.
 $prefixlen[a]$ – the prefix length of an anchor a .

```

1: function STRETCHINGWITHANCHOR( $P, M, A$ )
2:  $map \leftarrow \text{ComputeMapping}(P, M)$ 
3:  $P_t \leftarrow \text{PatchTarget}(P, map)$ 
4:  $anchor\_map \leftarrow \text{ComputeAcMapping}(P_t, A)$ 
5:  $P' \leftarrow \text{PatchTarget}(P_t, anchor\_map)$ 
6: end function

7: function COMPUTEACMAPPING( $P, A$ )
8:  $offset \leftarrow 0$ 
9:  $ac \leftarrow A.begin()$ 
10:  $i \leftarrow 0$ 
11: while  $i < P.size$  do
12:  $curaddr \leftarrow P.base\_addr + i + offset$ 
13: if  $ac.addr == curaddr$  then
14:  $prefix \leftarrow i - \text{SIZEOF}(JMP)$ 
15: if  $P[prefix]$  is not the start of an instruction then
16:  $prefix \leftarrow \text{start of instruction before } prefix$ 
17: end if
18:  $prefixlen[ac] \leftarrow i - prefix$ 
19:  $i \leftarrow prefix$ 
20:  $offset \leftarrow offset + ac.len + prefixlen[ac]$ 
21:  $ac \leftarrow A.next()$ 
22: else
23:  $anchor\_map[i] \leftarrow i + offset$ 
24:  $i \leftarrow i + 1$ 
25: end if
26: end while
27: return  $anchor\_map$ 
28: end function

```

Pruning Anchors. Potentially, we can create anchors for all function entries to guarantee that we never miss any necessary function call forwarding. However, this is not efficient. In fact, we only need to create anchors for the subset of functions that could be the possible target of some indirect call. Assuming a 32-bit machine, we construct the subset with the following criterion: *Any four-byte data value or any four-byte immediate operand in an instruction is considered a possible indirect call target, if it is equal to one of the function entries.* We obtain this subset by sequentially scanning data and code sections. Our pruning heuristic is very effective in practice. For example, the code section size of *gcc* in SPEC CPU 2000 benchmark suite is over 1MB, with over 2000 functions; after pruning, there are only 271 functions left that need anchors.

Embedding a Component with Anchors. If an extracted component contains a function that may be invoked by an indirect call in the component, BISTRO will create an anchor in the target binary at exactly the same address of the function entry in the component's original binary to allow proper forwarding. If the anchor conflicts with some

existing anchor in the target binary, BISTRO will integrate the two overlapping anchors into an arbitration function and redirect control flow to the function instead. The function further determines which real target it should forward the call to. The calls from the target binary and those from the to-be-embedded component are distinguished by setting a flag. The arbitration function uses the flag to decide the real forwarding target.

In some rare cases, the space between two function entries might not be enough to hold the anchors. In such cases, instead of using the jump instruction for redirection, we use a software interrupt instruction, which takes only one byte. When an indirect call reaches the old function entry, a software exception will be generated and intercepted by our exception handler, which will redirect the control flow to the new function entry.

4.2 Handling Indirect Jumps

Indirect jumps are different from indirect calls as the jump targets may not be function entries, but rather anywhere in the binary. If we adopt the anchor approach, there would be too many anchors needed. One might leverage some heuristics such as that indirect jumps usually receive their targets from jump tables and thus simply patch the jump table entries. However, this is unsafe because of the difficulty of determining jump table boundaries. A jump table may not be distinguishable from regular data. Hence, we propose a different approach. Specifically, we insert a code snippet right before each indirect jump to translate the jump target to its mapped address in the stretched binary at runtime, as shown in the example below.

<pre>jmp eax</pre>	→	<pre>mov eax, mapping[eax - old_base] add eax, new_base jmp eax</pre>
--------------------	---	---

Note that the example is just for illustration. In our implementation, we use perfect hashing for address lookup, which will be explained later, and preserve the flag register during translation. Since a complete byte-to-byte mapping is computed in Algorithm 1, any indirect jump target could be properly translated and handled by this method. Observe that additional instructions need to be added to perform translation. We can easily handle this by stretching the subject binary to accommodate those instructions.

Branch Target Set Pruning. Although the translation using a complete mapping guarantees safety, it also introduces significant memory overhead. Each byte in the original binary requires 4 bytes to represent its mapped address. In fact, we only need a subset of the mapping: the stretched/shrunk binary will be safe as long as the mapping contains translation for every possible indirect jump target.

We construct the set with the following criterion: *any four-byte data value or any four-byte immediate operand in an instruction is considered a possible indirect jump target, if the value falls in the range of some code section.* We further prune the set by removing the values that point to the middle of an existing instruction. Note that the strategy is safe for long/set jumps as their jump targets are acquired at runtime. This pruning strategy is very effective in practice. For example, the code section size of Adobe Reader X (AcroRd32.exe) is over 800KB, with over 260K instructions; after pruning, there are only 3635 possible branch targets left.

Perfect Hash Translator. The remaining challenge is to achieve fast translation. Note that after pruning, the jump target set becomes a sparse set in the address space. As a

compromise between memory consumption and runtime overhead, we choose to use perfect hashing for translation. A perfect hash function maps a set of keys to another set of integer values without any collision. It guarantees $O(1)$ translation time. We use gperf [22] to generate the perfect hash function for the jump target set and compile it into a linkable .obj file that can be embedded in the target binary through BISTRO.

A perfect hash function may require more space than the N keys to achieve $O(1)$ translation time. In practice, we find the size of generated perfect hash functions acceptable. For example, for the 3635 branch targets of Adobe Reader, the generated hash function is about 152KB, which is about 11% of the size of the Adobe Reader binary.

5 Handling Data References

Binary extraction/stretching may cause relocation of data entries, so we need to ensure the correctness of instructions referencing those data. We discuss how to address this problem from the perspectives of the target binary and the component to be embedded.

Compared to the component, the target binary is usually more complex and involves a lot of global data references. To handle this problem efficiently, we group data in the binary as continuous data blocks. If a data block might be indirectly accessed, we will make sure the block is not re-located to avoid patching data accesses, by wrapping the block in an anchor. Note that the number of data access instructions is much larger than the number of indirect jumps/calls. Otherwise, if the data block is only directly accessed, we allow it to be relocated (by Algorithm 1). We use the following criterion: *if the value of any four-byte data, or any four-byte immediate operand (in an instruction) that is not directly used as an address falls in the range of a data block, then this block might be indirectly accessed using data pointers and hence should not be re-located.*

In contrast, data entries extracted as part of the to-be-embedded component are most likely to be relocated. For example, if they are sparsely distributed in the address space, the BISTRO extractor (Section 2) will collapse them into a contiguous block, causing relocation. We adopt a method similar to the dynamic jump target translation scheme to translate data reference addresses. We add a comparison before translation to avoid translating stack or heap accesses. According to our experience, only 2% of dynamic memory references need to be translated. We further use offline static peephole scanning to identify references that surely access stack and avoid instrumenting them completely.

6 Evaluation

We have implemented BISTRO for Win32 PE binaries as an IDA-Pro plug-in. We have addressed a variety of engineering challenges such as virtual space layout rearrangement with a large embedded component, patching PE header, import and export tables, and re-generating relocation table. We omit the details due to space limitation.

6.1 Performance: Efficiency and Overhead

We first evaluate the performance of BISTRO by stretching (1) real-world Windows-based applications and (2) SPEC CPU 2000 binaries. Our experiments are done on a

Table 1. Performance results of stretching Windows software and SPEC CPU 2000 binaries

Binary	Instr. Count	Indirect Jumps	Indirect Calls	Call/Jump Targets: Anchors(%)	Data Blocks: Data Anchors(%)	File Size (KB)		Initial Mem. Image Size (KB)		Run Time (s)		Stretching Time (s)
						Orig: Stch'ed	growth(%)	Orig: Stch'ed	growth(%)	Orig: Stch'ed	overhead(%)	
SPEC CPU 2000 benchmarks												
164.gzip	19825	19	103	98: 23 (23.47%)	163: 1 (0.61%)	86.5: 98.5	13.87%	424: 440	3.77%	83.2: 84.6	1.68%	0.752
175.vpr	54595	53	106	229: 31 (13.54%)	404: 1 (0.25%)	232: 248.5	7.11%	248: 268	8.06%	64.5: 64.6	0.16%	0.755
176.gcc	337033	456	260	3855: 271 (7.03%)	2580: 14 (0.54%)	1264: 1393	10.21%	1348: 1480	9.79%	33.3: 33.9	1.8%	1.420
181.mcf	20566	36	103	144: 25 (17.36%)	100: 2 (2.00%)	76.5: 85.5	11.76%	100: 108	8%	40.2: 40.4	0.5%	0.685
186.crafty	65375	56	130	312: 29 (9.29%)	247: 1 (0.40%)	283: 298.5	5.48%	1344: 1360	1.19%	38.2: 38.9	1.83%	0.935
197.parser	44554	36	112	155: 27 (17.42%)	463: 1 (0.22%)	164: 173.5	5.79%	352: 360	2.27%	83.1: 83.5	0.48%	0.754
252.eon	114249	50	441	1659: 1253 (75.53%)	1455: 1 (0.07%)	499: 575	15.23%	592: 668	12.84%	42.7: 44.7	4.68%	0.950
253.perlbmk	164093	148	211	2166: 499 (23.04%)	1293: 6 (0.46%)	626: 743	18.69%	648: 764	17.9%	63.3: 67.9	7.27%	1.118
254.gap	129464	35	1357	816: 625 (76.59%)	1142: 1 (0.09%)	452.5: 492	8.73%	896: 936	4.46%	35.4: 37.2	5.08%	1.001
255.vortex	132034	66	145	446: 71 (15.92%)	738: 1 (0.14%)	561: 585	4.28%	588: 612	4.08%	50.6: 51.1	0.99%	1.050
256.bzip2	21360	36	101	145: 25 (17.24%)	150: 1 (0.67%)	87.5: 99	13.14%	172: 184	6.98%	73.4: 74.6	1.63%	0.714
300.twolf	64669	41	106	193: 30 (15.54%)	391: 2 (0.51%)	253: 263	3.95%	296: 304	2.7%	93.2: 93.6	0.43%	0.809
177.mesa	143679	211	552	2675: 473 (17.68%)	942: 8 (0.53%)	549.5: 652.5	18.74%	568: 672	18.31%	64.9: 65.6	1.08%	0.990
179.art	23353	38	103	149: 26 (17.45%)	103: 2 (1.94%)	85.5: 94.5	10.53%	104: 112	7.69%	32: 32.3	0.94%	0.690
183.equake	21824	38	101	146: 27 (18.49%)	116: 1 (0.86%)	88.5: 97	9.6%	104: 112	7.69%	26.1: 26.1	0%	0.720
188.ammp	61214	39	128	224: 70 (31.25%)	279: 1 (0.36%)	235.5: 245.5	4.25%	252: 264	4.76%	88.7: 88.3	1.92%	0.780
Average	-	-	-	-(24.80%)	-	-	10.09%	-	7.53%	-	1.90%	-
Real-world Windows-based Software												
putty	107220	57	662	942: 291 (30.89%)	93: 1 (1.08%)	444: 496	11.71%	472: 524	11.02%	-	-	0.865
gvim	561626	294	5111	3893: 1004 (25.79%)	5081: 22 (0.43%)	1950.5: 2150	10.23%	2008: 2212	10.16%	-	-	2.121
notepad++	272434	159	4302	4897: 2695 (55.03%)	3394: 7 (0.21%)	1584: 1864	17.68%	1660: 1940	16.87%	-	-	1.480
Adobe Reader	273710	146	2543	3635: 2160 (59.42%)	3037: 11 (0.36%)	1445.9: 1702.4	17.74%	1472: 1728	17.39%	-	-	1.556
Chrome	230234	82	1280	1842: 933 (50.65%)	930: 6 (0.65%)	1211: 1338	10.49%	1240: 1368	10.32%	-	-	1.391
Average	-	-	-	-(44.36%)	-(0.55%)	-	13.57%	-	13.15%	-	-	-

Dell Inspiron 15R laptop with Intel(R) Core(TM) i5-2410M 2.30GHz CPU and 4GB memory, running Windows 7 SP1. For the SPEC CPU 2000 benchmark suite, we use the “win32-x86-vc7” config file which includes all integer benchmark binaries and four floating-point benchmark binaries. We compile the benchmark suite using Visual Studio 2010, with full optimizations. To test BISTRO on non-relocatable binaries, we set “/DYNAMICBASE:NO” switch for the compiler to prevent it from generating relocatable binaries. The application binaries are readily available and we do not know about their compilers. Although the binaries of Adobe Reader and Chrome web browser carry relocation tables, we ignore them for testing our solutions for non-relocatable binaries.

We measure the following performance metrics: (1) space overhead – for both binary file and initial memory image – of a stretched binary compared with its original version, (2) runtime overhead of the stretched binary, and (3) time for BISTRO to stretch the binary. In particular, we are interested in the overhead incurred by BISTRO *itself*, not by the execution of the embedded components. As such, we embed a minimal component (a one-byte snippet) into each subject binary in our experiments. To create a “worst-case” scenario, we insert it at the beginning of each binary so that every byte in the binary gets shifted, which entails *all* indirect control transfer targets in the binary to be redirected. The measured overhead is hence the upper bound of overhead.

For each SPEC 2000 binary, we run both its original and stretched versions, and compare their execution time and file/initial image size. We do not measure the execution time of the Windows applications because they are all interactive. We experience no perceivable overhead when using their stretched versions.

The results are shown in Table 1. From the *Indirect Jumps* and *Indirect Calls*² columns, we observe that indirect calls are very common in application binaries, indicating that they might be C++ programs. Further investigation confirms our speculation, indicating BISTRO’s effectiveness for binaries compiled from C++ programs.

² We exclude indirect calls to external library functions through import address table (IAT), as these external targets are not handled by our redirection mechanisms.

Moreover, there are much less indirect jumps than indirect calls, indicating they are likely to have less impact on runtime overhead. Note that a small number of indirect jumps does not imply an equally small number of potential indirect jump targets. In fact, due to the difficulty of identifying jump table boundaries, we conservatively consider any constant in a binary that appears to be an instruction address as a potential jump target. The large number of potential jump targets and the low impact on performance justify our design choice of using the slightly more expensive but more flexible dynamic target translation scheme (Section 4.2), compared to the anchor scheme (Section 4.1).

The *Call/Jump Targets: Anchors* column shows the number of potential indirect call/jump targets, the number of anchors generated, and their comparison. Observe that the number of anchors created is small, compared to the size of the potential set. For binaries from C++ programs, due to the heavy use of virtual methods, it is not a surprise to see many anchors created. The *Data Blocks: Data Anchors* column shows that only less than 1% of all data blocks need to be preserved at their original locations using anchors. From the *File Size* columns, we can see BISTRO only increases the file size by 10.1% on average for SPEC programs, and 13.6% for application binaries. The overhead is dominated by the perfect hash tables. The *Initial Mem. Image Size* columns show the initial memory consumption when the binary is loaded into memory, which increases by only 7.5% on average for SPEC programs and 13.2% for application programs. Note that BISTRO does not cause any additional memory overhead during execution. The *Run Time* columns present the runtime overhead, which is only 1.9% on average. Except *eon*, *perlbmk* and *gap*, all SPEC binaries have less than 2% overhead. The last column *Stretching Time* shows the stretching time of BISTRO. The time is consistently short, implying that BISTRO can stretch a binary at runtime when it is loaded.

6.2 Case Study I: Binary-Level Semantic Patching Using BISTRO

Code reuse is a common practice in software development. One popular approach is to directly compile and statically link a piece of re-usable code with the target software – either directly in the executable or in some private library – to make the software self-contained, avoid compatibility problems, and improve performance. Indeed, developers of many popular programs (e.g., *chrome* and *firefox*) reuse code this way. The consequence is that programs reusing the same code may have the code placed at different locations in their address spaces. The reused code may not even have the same instructions if compiled by different compilers.

Table 2. Results of binary semantic patching using BISTRO

Vulnerability	Patch Extracted From	Vulnerable Application Patched	Original File Size (KB)	Patched File Size (KB) w. / w.o. Reloc	Semantic Patch Available	Vendor Patch Available
CVE-2010-1205	libpng 1.2.43 → 1.2.44 (rpng2-win.exe)	Firefox 3.6.6 (xu.dll)	11747.5	12371.5 / 13005	6/25/2010	7/20/2010
CVE-2011-3026	libpng 1.4.8 → 1.4.9 (rpng2-win.exe)	Zoner Photo Studio 15 (Zsl.dll)	8225.1	8502.1 / 9181.6	2/18/2012	N/A
SA47322 / CVE-2012-0025	IrfanView 4.30 → 4.32 (Fpx.dll)	XnView 1.99.5 (Xfpx.dll)	356	368 / 400	12/20/2011	N/A
		LeadTools 17.5 (ltdkdu.dll)	138.5	143 / 151	12/20/2011	N/A
SA47388	XnView 1.98.5 → 1.98.8 (Xfpx.dll)	IrfanView 4.35 (Fpx.dll)	432	448 / 508	3/12/2012	N/A
		LeadTools 17.5 (ltdkdu.dll)	372.5	428.5 / 493.5	3/12/2012	N/A
SA48772 / CVE-2012-0278	IrfanView 4.33 → 4.34 (Fpx.dll)	XnView 1.99.5 (Xfpx.dll)	356	368 / 400	4/13/2012	N/A
		LeadTools 17.5 (ltdkdu.dll)	138.5	142.5 / 150.5	4/13/2012	N/A
SA49091	XnView 1.98.8 → 1.99 (Xfpx.dll)	LeadTools 17.5 (ltdkdu.dll)	372.5	428.5 / 488.5	6/15/2012	N/A

However, code reuse via static linking introduces a security liability: When a piece of re-usable code contains a vulnerability, all programs that reuse the code will suffer from the same vulnerability. If these programs have been shipped in binary forms, the only way to fix the vulnerability is to release multiple binary patches – one for each program and by the corresponding vendor. However, not all vendors react to a vulnerability with equal timeliness and some may not even be aware of the vulnerability not in their own code. Thus it may be desirable for customers, who do *not* have source code access, to patch these programs without vendors’ involvement. Binary *syntactic patching*, which directly applies a patch for software A to software B sharing the same (vulnerable) code, will hardly work, because of the different locations of the code and the syntactic differences between the two code copies (due to different compilers used or different call/jump targets inside the copies).

In our first case study, we show that BISTRO can enable *binary semantic patching*. Assume that software A and B share a function f and the vendor of A has released a binary patch of f for a vulnerability. Let the patched program and the patched function be A' and f' , respectively. We will use BISTRO to extract f' from A' and embed it to B to replace the vulnerable version. Note that BISTRO is critical in ensuring the extracted f' is properly patched and the target binary B is properly stretched to contain f' .

We acquire a group of application binaries that leverage the same vulnerable component using public, vendor-provided information (e.g., which libraries are used in the software) or by finding similar binary snippets using the binary comparison tool *bindiff* [31]. Suppose at least one binary in the group, say A , has a patched version A' . Our goal is to extract a *semantic* patch out of A' and transplant it to patch the other vulnerable binaries $\{B_1, \dots, B_n\}$.

We collect 6 real-world vulnerabilities, with their CVE or Secunia IDs shown in Column 1 of Table 2. For each vulnerability, the vulnerable program(s) that has been patched by its vendor is shown in Column 2. The file names in braces represent the files that are patched. Column 3 shows a list of other un-patched programs with the same vulnerabilities. Column 6 shows the patch release date for the application in Column 2, i.e. the earliest date we can extract the semantic patch. Column 7 shows the date when the vendors for the software in Column 3 release their patches (N/A means no vendor patch is available yet). Most of the applications used in this case study are close-source (except *libpng* and *firefox*). Observe that most of the applications in Column 3 do not have vendor patches so far. For *firefox*, the new version (3.6.7) which patched the vulnerability was released – but with a one-month latency. With BISTRO, we can fix all these vulnerable applications as soon as one vendor releases the corresponding patch.

Failure of Syntactic Patching. We first verify that simple syntactic patching does not work – that is, using an existing binary differencing tool that generates and applies patches (e.g., *xdelta*, *bsdifff*, *bspatch*, etc.) will not properly patch $B_{1\dots n}$. For each vulnerability in Table 2, we use *bsdifff* to extract the syntactic difference between the pair of shared functions (f and f') in the versions in Column 2 as a patch, and use *bspatch* to apply it to the corresponding vulnerable applications in Column 3. None of the resultant binaries works. Further inspection shows that syntactic patches cannot properly fix the call/jump targets that are different among copies of the same reused code.

Patch Transplanting. We have developed a binary semantic patching tool based on BISTRO and *bindiff*. The identification of the vulnerable function f in A and $B_{1\dots n}$ and the patched function f' in A' is omitted here since it is not the main focus of this paper. Details are presented in our technical report [1]. We use BISTRO to extract f' from A' as the semantic patch for f . For each vulnerable binary B , we use BISTRO to cut out f and then stretch the resulting binary to implant f' at the same starting address of f . BISTRO ensures the correctness of both f' and the patched binary B' by properly stretching and patching control transfer instructions and data references. Our patching tool tries to avoid extracting dependent functions or global data entries of f' (i.e., functions being called and global data accessed by f') as much as possible by redirecting them to their counterparts in the target binary B . Since f' is a patched version of f , they likely share the same dependencies. For example, for each function invocation to function g' inside f' , if *bindiff* is able to identify the matching function g in B , our tool will automatically redirect the invocation in the extracted patch to g , without extracting g' . To be conservative, g and g' must be fully matched. Otherwise, g' will be extracted as part of the semantic patch.

We evaluate our patching tool on the subjects in Table 2. We apply our tool in two different ways to stress-test the robustness of BISTRO: first, we use the relocation information when it is present in the binary; second, we do not use relocation information at all. In both runs, the patching is successful: the patched applications work well and no longer suffer from the corresponding vulnerabilities. Columns 4 and 5 show the file size changes. We note that the patches are not large, each consisting of tens to hundreds of instructions. However, it is not straightforward to generate them independently because of the nature of the vulnerabilities being patched.

The first two vulnerabilities are in *libpng*, which is widely used in various software to read, write and render PNG images. The two vulnerable applications in Column 3 have *libpng* statically linked in their private DLLs (*xul.dll* and *Zxl.dll*). To patch these DLLs, we extract the semantic patch from *rpng2-win.exe*, a sample application in the *libpng* package. The remaining four vulnerabilities lie in *libfpx*, a library to handle the Flashpix (.fpx) image format. For the four vulnerabilities, only the first one was patched by the maintainer of *libfpx*; the other three were patched by individual developers who use *libfpx*. However, as shown in the table, individual developers only care about patching the *libfpx* code in their *own* applications. Using our binary semantic patching tool, users of the un-patched applications can transplant the patches and eliminate the vulnerabilities without the help of application developers.

6.3 Case Study II: Malware Stitching Using BISTRO

In the second case study, we demonstrate how BISTRO helps in the study of cyber attacks and counter-attacks. Specifically, we use BISTRO to compose a new, *executable* malware by stitching 3 separate functional components extracted from a *non-executable* sample of the Conficker worm [14]. It is an *unpacked* version without relocation information. Based on the published technical report of Conficker [14] and manual code inspection, we identify the code and data associated with the following 3 components:

- **DNS API hijacking.** This component prevent DNS query of the web sites in a blacklist by hijacking the functions *Query_Main*, *DNSQuery_A*, *DNSQuery_W* and

DNSQuery_UTF8 in *dnsapi.dll*. The result is that those web sites will no longer be accessible using their domain names.

- **Code injection.** To hijack the functions in *dnsapi.dll* used by a process (e.g., Internet Explorer), the malware must inject itself into the address space of the process. This component performs the injection. It takes the process identifier (PID) of the target process and the path of the malware as parameters.
- **Process identification.** This component gets a process' PID using its process name and provide the PID to the code injection component.

It takes us 60 minutes to manually identify the three components above. After that we use BISTRO to extract the components from the Conficker sample. We then create a dummy DLL to serve as the container of those components. Next, we use BISTRO to embed the 3 components into the empty DLL, right before the *DllMain()* function. After that, we add instructions to function *DllMain()* to invoke the inserted components. The invocation code first checks if the current process is the target process. If so, it will invoke the DNS API hijacking component to hijack the DNS query. If not, it will call the process identification component to find the PID of the target process, and then call the DLL injection component to inject itself into the target process for DNS API hijacking. The whole composition process takes us about 30 minutes.

To verify the functionality of the newly composed malware, we select two applications as our targets (in two experiment runs): Internet Explorer and FlashFXP (an FTP client). After being loaded, the malware injects itself into the target processes. Then, in the target application, we try to access web site *avast.com*, which is blacklisted by Conficker [14]. Interestingly, the access was not blocked at first (namely, the malware did not succeed). After debugging, we found that it was due to a bug in Conficker's original code: the hijacked *DNSQuery_W()* has one unnecessary instruction which sets a wrong return value. We point out that *we would not have spotted the problem, had we not made these components executable and observed their runtime behavior*. After removing this instruction using BISTRO, both IE and FlashFXP are successfully compromised: they can no longer access *avast.com* due to a DNS query error.

7 Discussion

BISTRO cannot work on self-modifying, self-checking or obfuscated binaries. Self-modifying binaries generate instructions dynamically during runtime, which could not be statically patched using BISTRO. Self-checking binaries use checksum or other integrity checks to detect changes made to their code by BISTRO. Obfuscated binaries in many cases cannot be properly disassembled. However, we note that all other static binary rewriting/instrumentation techniques face the same challenge.

Our anchor and branch target set pruning criteria assume the constants in a binary represent a superset of all possible indirect control transfer targets. This assumption should hold for binaries generated by common compilers. One exception is position independent code (PIC), which obtains addresses at runtime and uses them to compute indirect control transfer targets. All PIC we encountered has the form of making a call and then obtaining the return address from the stack (e.g., *call \$+5; pop eax*), which

is the address of the instruction right after the call. We identify all such instructions and insert snippets to adjust the addresses to their mapped addresses. Also, special compilers or hand-written binaries might violate our assumption. For example, in the instruction sequence *mov eax, Target; add eax, 5; jmp eax*, the actual target is *Target + 5* instead of the constant *Target*; our pruning heuristic will miss the actual target. For such binaries, we can choose not to prune the anchor set or the branch target set.

Currently, BISTRO only supports Win32 PE binaries. However, the design is general, without relying on specific features of Win32 PE.

8 Related Work

The most related work is discussed in Section 1 (with details in [1].) In this section, we discuss other related work in the general area of binary manipulation. They fall into three categories: (1) static binary rewriting, (2) dynamic binary rewriting, and (3) binary component identification, extraction and reuse.

Static Binary Rewriting. Static binary rewriting is widely applied to many scenarios, such as in-lined reference monitors [33], software fault isolation [24, 25, 6, 26], binary instrumentation [10, 9, 11, 5, 7, 4], binary obfuscation [36, 37] and retrofitting security in legacy binaries [27, 12]. Most of these rewriters require the binary to be compiled by specific compilers, or contains symbolic information.

PEBIL [11], REINS [33], STIR [13] and SecondWrite [12] are recently developed rewriters targeting stripped binaries. However, they all aim at rewriting a single binary, so they all keep the original code and data sections in place. In contrast, BISTRO supports “transplanting” binary components from one or more binaries to a target binary, which requires rewriting and combining multiple binaries. Keeping original code and data sections in place may result in address space conflicts and hence is not an option for BISTRO. Detour-based techniques [10, 8, 9] are lightweight and can work on stripped binaries. However, they cannot patch non-trivial jumps/calls that are repositioned.

Dynamic Binary Rewriting. Dynamic binary rewriters [2, 3, 28] are generally more robust as they do not require specific compilers or symbolic information. It is possible to apply them to conduct binary stretching and transplanting. However, we choose to use a static approach mainly because of the following two reasons: (1) Dynamic binary rewriters usually have much higher run time overhead than static ones. (2) It is more difficult to deploy a instrumented binary using dynamic approaches, as the rewriter itself must be deployed along with the binary.

Binary Component Identification, Extraction and Reuse. Recently, researchers proposed to identify, extract and reuse components from binaries for security applications [29, 16, 23]. Kolbitsch et al. proposed Inspector Gadget [29], which performs dynamic slicing on a malware binary to identify and extract the slice pertinent to a specific malicious functionality, and wrap the slice into a stand-alone binary that could be reused later to execute the malicious functionality. Inspector Gadget is able to extract component from self-modifying code, which is not supported by BISTRO due to the limitation of static binary manipulation. Using dynamic slicing, Inspector Gadget also avoids the problem of handling indirect calls/jumps in BISTRO as all call/jump targets

are directly known in the slice. However, the slice may not cover all possible code paths, which could result in incorrect execution when the user provides an input that would lead to a code path which is not included in the slice. Compared to Inspector Gadget, BISTRO statically extracts the component from the binary, which involves handling of indirect calls/jumps but provides better code path coverage.

Caballero et al. proposed BCR [16] to identify and extract a function from a binary using a combination of static and dynamic analysis. The extracted function, in the format of disassembly, is wrapped in a C file to be reused. BCR statically disassembles the designated function starting at its entry point; when encountering indirect call/jumps, BCR utilizes dynamic execution trace to find the call/jump targets. During the extraction, BCR rewrites all calls/jumps to use labels. Using labels implies that indirect call/jump can only have one target, which may not always hold in practice. Although BCR specially handles indirect jumps that use jump tables, there are other forms of multiple-target indirect calls/jumps such as function pointers and vttables. Compared to BCR, BISTRO preserves the original semantic of indirect calls/jumps when performing component extraction, hence does not suffer from this problem.

Neither Inspector Gadget nor BCR could extract components from non-executable binaries (as in Section 6.3) because they are based on dynamic analysis. This is a very common case in malware analysis, where a given malware sample may not run due to various reasons (e.g., missing dependent libraries, missing inputs). In such a case, BISTRO can still perform component extraction statically. Moreover, neither Inspector Gadget nor BCR supports reusing extracted components to enhance legacy binaries (as in Section 6.2), as they lack the capability of embedding instructions that invoke the components into the target binary. BISTRO is able to handle such a scenario by performing both binary component extraction and embedding.

Lin et al. proposed ROC [23] which uses dynamic slicing to identify reusable functional components in a binary. Different from BISTRO, ROC only invokes the identified components from the *same* binary; it does not support extracting a component for reuse in a *different* binary.

9 Conclusions

We have developed a new pair of binary program manipulation primitives called BISTRO for extracting and re-packaging a functional component from a binary program; and for embedding a functional component in a target binary program, respectively. We address the challenges of patching control transfer instructions and data references to preserve the semantics of both the extracted component and the stretched binary program, especially indirect calls and jumps. BISTRO incurs low runtime overhead (1.9% on average) and small space overhead (11% on average). The extraction and embedding operations are highly efficient, with less than 1.5s for most cases. We have applied BISTRO to two security application scenarios, demonstrating its efficiency, precision, and versatility.

Acknowledgements. We thank the anonymous reviewers and our shepherd, Pratyusa Manadhata, for their insightful comments. This research has been supported by DARPA under Contract 12011593. Any opinions, findings, and conclusions in this paper are those of the authors only and do not necessarily reflect the views of DARPA.

References

1. Deng, Z., Zhang, X., Xu, D.: BISTRO: Binary Component Extraction and Embedding for Software Security Applications. CERIAS Technical Report TR 2013-3, Purdue University (June 2013)
2. Bruening, D.: Efficient, transparent, and comprehensive runtime code manipulation. Ph.D. dissertation. MIT (2004)
3. Bellard, F.: Qemu, a fast and portable dynamic translator. In: USENIX ATC 2005 (2005)
4. De Sutter, B., De Bus, B., De Bosschere, K.: Link-time binary rewriting techniques for program compaction. In: TOPLAS 2005 (2005)
5. Muth, R., Debray, S., Watterson, S., De Bosschere, K.: Alto: a link-time optimizer for the compaq alpha. In: SPE 2001 (2001)
6. Abadi, M., Budi, M., Erlingsson, U., Ligatti, J.: Control-flow integrity principles, implementations, and applications. In: TISSEC 2009 (2009)
7. Eustace, A., Srivastava, A.: Atom: A flexible interface for building high performance program analysis tools. In: USENIX ATC 1995 (1995)
8. Buck, B., Hollingsworth, J.K.: An api for runtime code patching. In: IJHPCA 2000 (2000)
9. Romer, T., Voelker, G., Lee, D., Wolman, A., Wong, W., Levy, H., Bershad, B., Chen, B.: Instrumentation and optimization of win32/intel executables using etch. In: USENIX Windows NT Workshop (1997)
10. Hunt, G., Brubacher, D.: Detours: Binary interception of win32 functions. In: USENIX Windows NT Symposium (1999)
11. Laurenzano, M., Tikir, M., Carrington, L., Snaveley, A.: Pebil: Efficient static binary instrumentation for linux. In: ISPASS 2010 (2010)
12. O'Sullivan, P., Anand, K., Kotha, A., Smithson, M., Barua, R., Keromytis, A.D.: Retrofitting security in COTS software with binary rewriting. In: Camenisch, J., Fischer-Hübner, S., Murayama, Y., Portmann, A., Rieder, C. (eds.) SEC 2011. IFIP AICT, vol. 354, pp. 154–172. Springer, Heidelberg (2011)
13. Wartell, R., Mohan, V., Hamlen, K., Lin, Z.: Binary stirring: Self-randomizing instruction addresses of legacy x86 binary code. In: CCS 2012 (2012)
14. Porras, P., Saidi, H., Yegneswaran, V.: Conficker c analysis. SRI International (2009)
15. Johnson, N., Caballero, J., Chen, K., McCamant, S., Poosankam, P., Reynaud, D., Song, D.: Differential slicing: Identifying causal execution differences for security applications. In: IEEE S&P 2011 (2011)
16. Caballero, J., Johnson, N., Mccamant, S., Song, D.: Binary code extraction and interface identification for security applications. In: NDSS 2010 (2010)
17. Balakrishnan, G., Reps, T.: Analyzing memory accesses in x86 executables. In: Duesterwald, E. (ed.) CC 2004. LNCS, vol. 2985, pp. 5–23. Springer, Heidelberg (2004)
18. Slowinska, A., Stancescu, T., Bos, H.: Howard: A dynamic excavator for reverse engineering data structures. In: NDSS 2011 (2011)
19. Lee, J., Avgerinos, T., Brumley, D.: Tie: Principled reverse engineering of types in binary programs. In: NDSS 2011 (2011)
20. Lin, Z., Zhang, X., Xu, D.: Automatic reverse engineering of data structures from binary execution. In: NDSS 2010 (2010)
21. Hex-Rays, Ida pro disassembler, <http://www.hex-rays.com/products/ida/index.shtml>
22. Schmidt, D.: Gperfl: a perfect hash function generator. More C++ gems (2000)
23. Lin, Z., Zhang, X., Xu, D.: Reuse-oriented camouflaging trojan: Vulnerability detection and attack construction. In: DSN 2010 (2010)

24. Wahbe, R., Lucco, S., Anderson, T., Graham, S.: Efficient software-based fault isolation. *OS Review* (1994)
25. McCamant, S., Morrisett, G.: Evaluating sfi for a cisc architecture. In: *USENIX Security 2006* (2006)
26. Erlingsson, Ú., Abadi, M., Vrable, M., Budiú, M., Necula, G.: Xfi: Software guards for system address spaces. In: *OSDI 2006*(2006)
27. Prasad, M., Chiueh, T.: A binary rewriting defense against stack based buffer overflow attacks. In: *USENIX ATC 2003* (2003)
28. Scott, K., Kumar, N., Velusamy, S., Childers, B., Davidson, J., Soffa, M.: Retargetable and reconfigurable software dynamic translation. In: *CGO 2003* (2003)
29. Kolbitsch, C., Holz, T., Kruegel, C., Kirda, E.: Inspector gadget: Automated extraction of proprietary gadgets from malware binaries. In: *IEEE S&P 2010* (2010)
30. Srivastava, A., Edwards, A., Vo, H.: Vulcan: Binary transformation in a distributed environment. *Tech. Rep., Microsoft Research* (2001)
31. Flake, H.: Structural comparison of executable objects. In: *DIMVA 2004* (2004)
32. Falliere, N., Murchu, L., Chien, E.: W32. stuxnet dossier. White paper, Symantec Corp., *Security Response* (2011)
33. Wartell, R., Mohan, V., Hamlen, K., Lin, Z.: Securing untrusted code via compiler-agnostic binary rewriting. In: *ACSAC 2012* (2012)
34. Dinaburg, A., Royal, P., Sharif, M., Lee, W.: Ether: malware analysis via hardware virtualization extensions. In: *CCS 2008* (2008)
35. Nanda, S., Li, W., Lam, L., Chiueh, T.: BIRD: binary interpretation using runtime disassembly. In: *CGO 2006* (2006)
36. Moser, A., Kruegel, C., Kirda, E.: Limits of static analysis for malware detection. In: *ACSAC 2007* (2007)
37. Popov, I., Debray, S., Andrews, G.: Binary obfuscation using signals. In: *USENIX Security 2007* (2007)