

# Obfuscation of Executable Code to Improve Resistance to Static Disassembly \*

Cullen Linn      Saumya Debray  
Department of Computer Science  
University of Arizona  
Tucson, AZ 85721.  
{linnc, debray}@cs.arizona.edu

## ABSTRACT

A great deal of software is distributed in the form of executable code. The ability to reverse engineer such executables can create opportunities for theft of intellectual property via software piracy, as well as security breaches by allowing attackers to discover vulnerabilities in an application. The process of reverse engineering an executable program typically begins with disassembly, which translates machine code to assembly code. This is then followed by various decompilation steps that aim to recover higher-level abstractions from the assembly code. Most of the work to date on code obfuscation has focused on disrupting or confusing the decompilation phase. This paper, by contrast, focuses on the initial disassembly phase. Our goal is to disrupt the static disassembly process so as to make programs harder to disassemble correctly. We describe two widely used static disassembly algorithms, and discuss techniques to thwart each of them. Experimental results indicate that significant portions of executables that have been obfuscated using our techniques are disassembled incorrectly, thereby showing the efficacy of our methods.

## Categories and Subject Descriptors

K.6.5 [Management of Computing and Information Systems]: Security and Protection—*unauthorized access*

## General Terms

Security

## Keywords

disassembly, code obfuscation

## 1. INTRODUCTION

Advances in program analysis and software engineering technology in recent years have led to significant improvements in tools for program analysis and software development. Unfortunately, this same technology can, in many cases, be subverted to reverse engineer software systems with the goal of discovering vulnerabilities,

\*This work was supported in part by the National Science Foundation under grants EIA-0080123 and CCR-0113633.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CCS'03, October 27–31, 2003, Washington, DC, USA  
Copyright 2003 ACM 1-58113-738-9/03/0010 ...\$5.00.

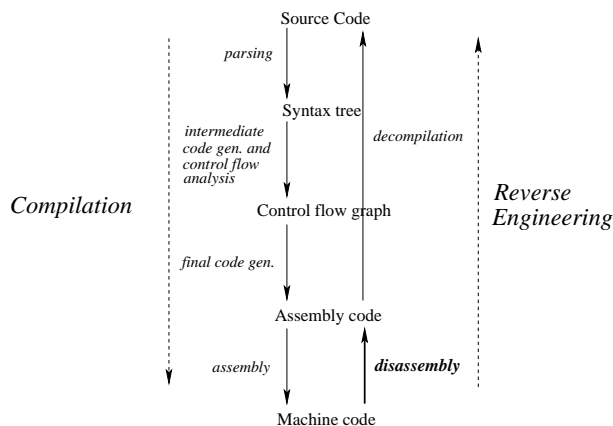


Figure 1: The processes of compilation and reverse engineering

making unauthorized modifications, or stealing intellectual property. These all require an ability to take an executable program and reconstruct its high-level structure to some extent. For example, to identify vulnerabilities in a software system, a hacker has to be able to figure out how it works and where it may be attacked. Similarly, to steal a piece of software with an embedded copyright notice or software watermark, a pirate must reconstruct enough of its internal structure to be able to identify and delete the copyright notice or watermark without affecting the functionality of the program.

The problem can be addressed by maintaining the software in encrypted form and decrypting it as needed during execution [1], or using specialized hardware (e.g., see [16]). While effective, such approaches have the disadvantages of high performance overhead or loss of flexibility because the software can no longer be run on stock hardware. An alternative approach, which we focus on, is to use code obfuscation techniques to enhance software security [9, 10, 11, 12, 28]. The goal is to deter attackers by making the cost of the reconstructing the high-level structure of the program prohibitively high.

The processes of compilation and reverse engineering are illustrated in Figure 1. Compilation refers to the translation of a source-language program to machine code; it consists of a series of steps, each producing successively lower-level program representations. Reverse engineering is the dual process of recovering higher-level structure and semantics from a machine code program. Broadly speaking, we can divide the reverse engineering process into two parts: *disassembly*, which produces assembly code from machine code; and *decompilation*, which reconstructs the higher-level se-

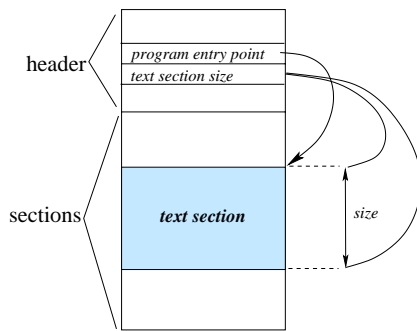


Figure 2: The structure of an executable file

mantic structure of the program from the assembly code. Most of the prior work on code obfuscation and tamper-proofing focus on various aspects of decompilation. For example, a number of researchers suggest relying on the use of difficult static analysis problems, e.g., involving complex Boolean expressions, pointers, or indirect control flow, to make it harder to construct a precise control flow graph for a program [3, 12, 20, 26, 27].

The work described in this paper, by contrast, focuses on the disassembly process. Our goal is to increase the difficulty of statically disassembling a program. This is independent of, and complementary to, current approaches to code obfuscation. It is independent of them because our techniques can be applied regardless of whether or not any of these other obfuscating transformations are being used. It is complementary to them because, by making a program harder to disassemble accurately, we add yet another barrier to recovering high-level semantic information about a program.

## 2. BACKGROUND: DISASSEMBLY

A machine code file typically consists of a number of different sections, e.g., *text*, *read-only data*, etc., that contain various sorts of information about the program, together with a header describing these sections. Among other things, the header contains information about the program entry point, i.e., the location in the file where the machine instructions begin (and where program execution begins), and the total size or extent of these instructions<sup>1</sup> (see Figure 2) [15]. Disassembly refers to the process of recovering a sequence of assembly code instructions from such a file, e.g., in a textual format readable by a human being.

Broadly speaking, there are two approaches to disassembly: *static disassembly*, where the file being disassembled is examined by the disassembler but is not itself executed during the course of disassembly; and *dynamic disassembly*, where the file is executed on some input and this execution is monitored by an external tool (e.g., a debugger) to identify the instructions that are being executed. Static disassembly has the advantage of being able to process the entire file all at once, while dynamic disassembly only disassembles a “slice” of the program, i.e., those instructions that were executed for the particular input that was used. Another advantage of static disassembly is that it takes time proportional to the size of the program, while the time taken by dynamic disassembly is typically proportional to the number of instructions executed by the program at runtime—the former tends to be considerably less than

<sup>1</sup>This applies to most file formats commonly encountered in practice, including Unix *a.out*, ELF, COFF, and DOS EXE files. The information about the entry point and code section size is implicit in DOS COM files.

the latter (often by several orders of magnitude), making static disassembly considerably more efficient than dynamic disassembly.

This paper focuses on static disassembly. There are two generally used techniques for this: *linear sweep* and *recursive traversal* [22]. The remainder of this section sketches each of them.

### 2.1 Linear Sweep

The linear sweep algorithm begins disassembly at the input program’s first executable byte, and simply sweeps through the entire text section disassembling each instruction as it is encountered:

```

global startAddr, endAddr;
proc DisasmLinear(addr)
begin
  while (startAddr ≤ addr < endAddr) do
    I := decode instruction at address addr;
    addr += length(I);
  od
end

proc main()
begin
  startAddr := address of the first executable byte;
  endAddr := startAddr + text section size;
  DisasmLinear(ep);
end

```

This method is used by programs such as the GNU utility *objdump* [19] as well as a number of link-time optimization tools [8, 18, 24].

The main weakness of this algorithm is that it is prone to disassembly errors resulting from the misinterpretation of data that is embedded in the instruction stream. Only under special circumstances, e.g., when an invalid opcode is encountered, can the disassembler become aware of such disassembly errors.

### 2.2 Recursive Traversal

The problem with the linear sweep algorithm is that, because it does not take into account the control flow behavior of the program, it cannot “go around” data (e.g., alignment bytes, jump tables, etc.) embedded in the instruction stream, and mistakenly interprets them as executable code. An obvious fix would be to take into account the control flow behavior of the program being disassembled in order to determine what to disassemble. Intuitively, whenever we encounter a branch instruction during disassembly, we determine the possible control flow successors of that instruction, i.e., addresses where execution could continue, and proceed with disassembly at those addresses (e.g., for a conditional branch instruction we would consider the branch target and the fall-through address):

```

global startAddr, endAddr;
proc DisasmRec(addr)
begin
  while (startAddr ≤ addr < endAddr) do
    if (addr has been visited already) return;
    I := decode instruction at address addr;
    mark addr as visited;
    if (I is a branch or function call)
      for each possible target t of I do
        DisasmRec(t);
      od
    return;
    else addr += length(I);
  od
end

```

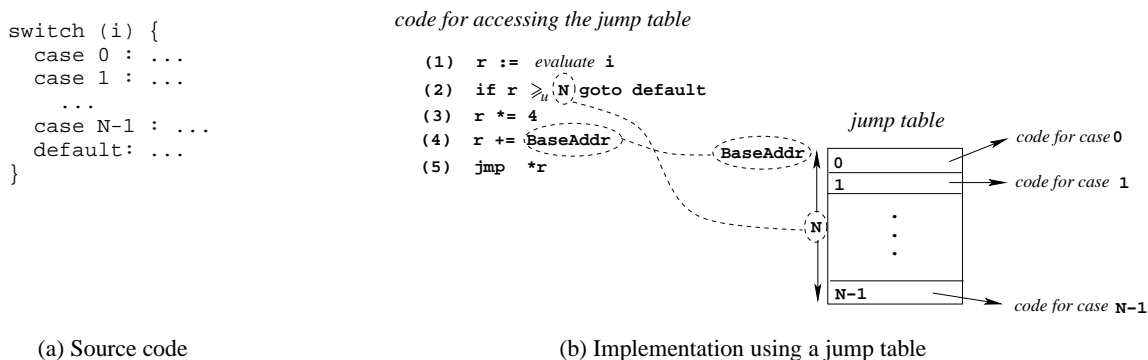


Figure 3: A example of a C switch statement and its implementation using a jump table

```

proc main()
begin
  startAddr := program entry point;
  endAddr := startAddr + text section size;
  DisasmRec(startAddr);
end

```

Variations on this basic approach to disassembly, which we term *recursive traversal*, are used by a number of binary translation and optimization systems [4, 23, 25].

A virtue of this algorithm is that, by following the control flow behavior of the program being processed, it is able to “go around” and thus avoid disassembly of data embedded in the text section. Its main weakness is that its key assumption, that we can precisely identify the set of control flow successors of each control transfer operation in the program, may not always hold in the case of indirect jumps. Imprecision in determining the set of possible targets of such a jump will result either in a failure to disassemble some reachable code (if the set of targets is underestimated) or erroneous disassembly of data (if the set of targets is overestimated).

Some researchers have proposed *ad hoc* extensions to the basic algorithm outlined above to handle common cases of indirect jumps. As an example, one of the most common uses of indirect jumps involves *jump tables*, a construct used by compilers to implement C-style **switch** statements [2]. This is illustrated in Figure 3. The jump table itself is a contiguous array of  $N$  code addresses, corresponding to the  $N$  cases in the **switch** statement. The code to access the jump table evaluates the index expression  $i$ ; checks to see whether this expression falls within the bounds of the table, i.e., whether  $0 \leq i < N$ ; adds the scaled value of the index expression to the base address of the table to obtain the address of the  $i^{\text{th}}$  entry in the table; then jumps indirectly through this location. The check of whether the index expression falls within the table bounds can be accomplished using a single unsigned comparison (denoted by  $\geq_u$  in instruction (2) in Figure 3(b)) [2]. To determine the possible target addresses of an indirect jump through a jump table, a disassembler needs to know the base address of the table and its extent, i.e., the values of `BaseAddr` and `N` in Figure 3(b). This can be done by scanning back from the indirect jump instruction to find the instruction that adds the scaled index to the base address (instruction (4) in Figure 3(b)), whence the base address can be extracted; and the unsigned compare of the index (instruction (2) in Figure 3(b)), whence the table size can be determined. Once this has been done, disassembly can continue at each target identified from the  $N$  table entries starting at location `BaseAddr` [6].

Code that is reachable only through indirect control transfers may not be found using the basic algorithm above. To handle

this problem, some systems, e.g., the UQBT binary translation system [5], resort to “speculative disassembly.” The idea is to process undisassembled portions of the text segment that appear to be code, in the expectation that they might be the targets of indirect function calls; a “speculative bit” is set when this is done, and speculative disassembly of a particular region of memory is abandoned if an invalid instruction is encountered.

### 3. THWARTING DISASSEMBLY

In order to thwart a disassembler, we have to somehow confuse, as much as possible, its notion of where the instruction boundaries in a program lie. This section discusses some ways in which this can be achieved. We first discuss a phenomenon that we had not expected: that of disassembly errors that “repair” themselves within a relatively short distance. This is followed by a discussion of a general technique we use to inject “junk bytes” into the instruction stream to introduce disassembly errors. After this we discuss specific details of the way in which this is done to confuse the two disassembly algorithms discussed in the previous section.

#### 3.1 Self-Repairing Disassembly

On some instruction sets—most notably, that of the Intel IA-32 architecture—the instruction structure is such that, very often, the disassembly process is self-repairing: even when a disassembly error occurs (e.g., due to the disassembly of data), the disassembler eventually ends up re-synchronizing with the actual instruction stream. This is illustrated by the example in Figure 4, which shows a typical byte sequence in memory, together with the actual disassembly (on the left), and the disassemblies we obtain if the disassembler is off by 1, 2, or 3 bytes, on the right. When the disassembly is initially off by a single byte, the disassembler produces two erroneous instructions but is back in synchrony with the original disassembly by the second instruction in the actual disassembly sequence. A similar phenomenon occurs when the disassembler is initially off by two bytes: it resynchronizes with the second instruction in the actual disassembly after producing a single incorrectly disassembled instruction. If the disassembler is initially off by three bytes, it generates three incorrectly disassembled instructions but resynchronizes by the third instruction in the actual disassembly.

Obviously, the actual resynchronization behavior on a particular program will depend on its particular distribution of instructions. In practice, however, we have found that disassembly errors usually resynchronize quite quickly—often within just one or two instructions beyond the point at which the disassembly error occurred. Efforts to confuse disassembly have to take this self-repairing aspect of disassembly into account.

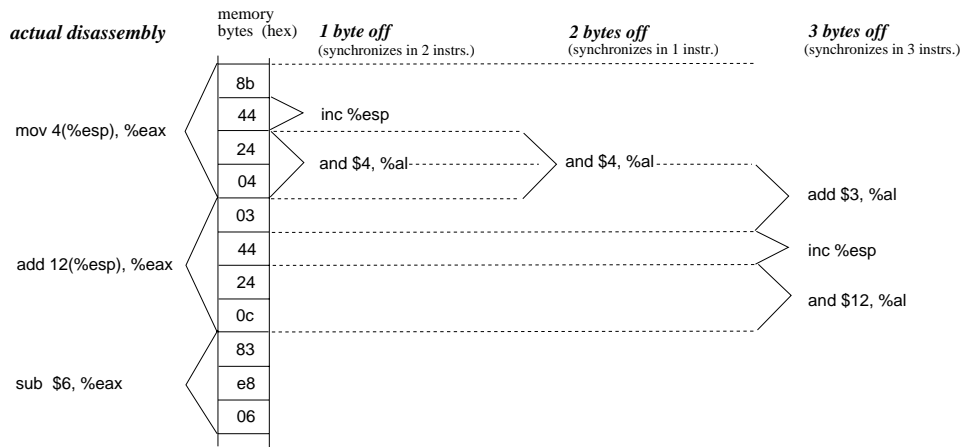


Figure 4: An example of self-repairing disassembly

### 3.2 Junk Insertion

We can introduce disassembly errors by inserting “junk” bytes at selected locations in the instruction stream where the disassembler is likely to expect code. (An alternative approach involves partially or fully overlapping instructions, e.g., see [7]: this is discussed in Section 5.) It is not difficult to see that any such junk bytes must satisfy two properties. First, in order to actually confuse the disassembler, the junk bytes must be partial instructions, not complete instructions. Second, in order to preserve program semantics, such partial instructions must be inserted in such a way that they are unreachable at runtime. To this end, define a basic block as a *candidate block* if it can have such junk bytes inserted before it. In order to ensure that any junk so inserted is unreachable during execution, a candidate basic block cannot have execution fall through into it. In other words, the basic block immediately before a candidate block must end in an unconditional control transfer, e.g., an unconditional jump or a return from a function. Candidate blocks can be identified in a straightforward way by scanning the basic blocks of the program after their final memory layout has been determined.

As mentioned in Section 3.1, the static disassembly process very often manages to “re-synchronize” itself after a disassembly error. Once a candidate block  $B$  has been identified, we have to determine what junk bytes to insert before it so as to confuse the disassembler as much as possible, i.e., delay this re-synchronization for as long as possible. To do this, we take a particular  $n$ -byte instruction  $I$  (our current implementation considers a 6-byte bitwise-OR instruction, but it is easy to extend this to other instructions), and determine how far away this re-synchronization would occur if the first  $k$  bytes of  $I$  were to be inserted immediately before the candidate block  $B$ , for each  $k$ ,  $0 < k < n$ . To determine the re-synchronization point, for each such  $k$  we simulate disassembly for the candidate block, assuming that the disassembler encounters the first  $k$  bytes of  $I$  at the beginning of  $B$ , then continuing with the byte sequence comprising the machine-level encodings of the instructions actually in  $B$ . Using this approach we determine the value  $k_{max}$  of  $k$  for which the re-synchronization distance is maximized, and insert the first  $k_{max}$  bytes of  $I$  immediately before block  $B$ .

### 3.3 Thwarting Linear Sweep

As observed in Section 2.1, linear sweep disassembly is generally unable to distinguish data embedded in the text section. We can exploit this weakness by inserting “junk” bytes at selected lo-

cations in the instruction stream, as discussed in Section 3.2. One point to note here is that since the simulation of disassembly scans forward from each candidate to determine the number of “junk” bytes to be inserted there, it is important to ensure that such decisions made for one candidate are not subsequently invalidated by the insertion of junk into subsequent candidates. To avoid such effects, we consider candidate blocks in reverse order when inserting junk.

With the approach described thus far, we find that we are typically able to attain a “confusion factor” of about 15% on average—i.e., 15% of the instructions in a program are incorrectly disassembled (confusion factors are discussed in more detail in Section 4). The reason that it is not higher is that candidates for the insertion of junk bytes cannot have execution fall through into them: the preceding block has to end in an unconditional control transfer. We have found that, in programs obtained from a typical optimizing compiler, candidate blocks tend to be around 30 instructions apart on average.<sup>2</sup> This distance, combined with the self-repairing nature of disassembly, means that when disassembly goes wrong after the insertion of junk before a candidate, it typically manages to re-synchronize before the next candidate is encountered. We increase the number of candidates by a transformation called *branch flipping*. The idea is to invert the sense of conditional jumps, by converting code of the form

$$b_{cc} \text{ Addr}$$

where  $cc$  represents a condition, e.g., ‘eq’ or ‘ne’, to

$$b_{\overline{cc}} L'$$

$$j_{mp} \text{ Addr}$$

$L'$ :

where  $\overline{cc}$  is the complementary condition to  $cc$ , e.g., a ‘beq . . .’ is converted to a ‘bne . . .’. The basic block at  $L'$  now becomes a candidate. With this transformation, the distance between candidate blocks drops to about 12 instructions on average, and the confusion factor rises to about 37%. Yet another measure that can be taken to increase candidates for junk insertions is *call conversion* which raises instruction confusion to about 42%. This method is discussed in more detail in section 3.4.2.

<sup>2</sup>These data reflect the SPECint-95 benchmark suite compiled using *gcc* at optimization level  $-O3$ . The averages given here were computed as geometric means.

### 3.4 Thwarting Recursive Traversal

The main strength of the recursive disassembly algorithm—its ability to deal intelligently with control flow and thereby disassemble around data embedded in the text segment—also turns out to be a weakness that we can take advantage of to confuse the disassembly process. There are two (related) aspects of recursive traversal that we can exploit. The first is that when it encounters a control transfer, disassembly continues at those locations that are deemed to be the possible control transfer targets. In this context, disassemblers typically assume that commonly encountered control transfers, such as conditional branches and function calls, behave “reasonably.” For example, a conditional branch is assumed to have two possible targets: the branch target and the fall through to the next instruction. Similarly, a function call is assumed to return to the instruction immediately following the call instruction.

The second aspect of recursive traversal is that identifying the set of possible targets of indirect control transfers is difficult. Recursive traversal disassemblers therefore generally resort to *ad hoc* techniques, such as examining bounds checks associated with jump tables, or disassembling speculatively, to handle commonly encountered situations involving indirect jumps.

Below we discuss different ways in which these characteristics can be exploited to confuse recursive traversal disassembly.

#### 3.4.1 Branch Functions

The assumption that a function returns to the instruction following the call instruction can be exploited using what we term *branch functions*. The idea is illustrated in Figure 5. Given a finite map  $\varphi$  over locations in a program

$$\varphi = \{a_1 \mapsto b_1, \dots, a_n \mapsto b_n\}$$

a branch function  $f_\varphi$  is a function that, whenever it is called from one of the locations  $a_i$ , causes control to be transferred to the corresponding location  $b_i$ ,  $1 \leq i \leq n$ . Given such a branch function  $f_\varphi$ , we can replace  $n$  unconditional branches in a program,

```
a1 : jmp b1
...
a2 : jmp b2
...
an : jmp bn
```

by calls to the branch function:

```
a1 : call fφ
...
a2 : call fφ
...
an : call fφ
```

The code for the branch function is responsible for determining the target location  $b_i$  based on the location  $a_i$  it was called from, then branching to the appropriate  $b_i$ . Moreover, it has to do this in such a way that the program state is that which would have been encountered at the location  $b_i$  in the original code with unconditional branches. Note that a branch function does not behave like “normal” functions, in that it typically does not return to the instruction following the call instruction, but instead branches to some other location in the program that depends, in general, on where it was called from.<sup>3</sup>

Branch functions serve two distinct purposes. The first is to obscure the flow of control in the program: by sufficiently obscuring

<sup>3</sup>This can, however, have adverse performance implications on some architectures, e.g., the Intel Pentium, by interfering with the branch prediction and/or return stack buffer mechanisms.

the computation of the target address  $b_i$  within the branch function, we can make it difficult for an attacker to reconstruct the original map  $\varphi$  it realizes. The second is to create opportunities for misleading a disassembler: since a disassembler will typically continue disassembly at the instruction following the call instruction, we can introduce errors in the disassembly by inserting junk bytes at the point immediately after each ‘call  $f_\varphi$ ’ instruction in a manner similar to that discussed in Section 3.3.

Branch functions can be implemented in a number of ways. For example, a straightforward implementation might use the return address to look up a table, via a simple linear or binary search, to determine the target address. The disadvantage with such schemes is that they are relatively straightforward to reverse engineer.

A more sophisticated implementation of branch functions has the callee pass, as an argument to the branch function, the offset from the instruction immediately after it (whose address is passed to the branch function as the return address) to the target  $b_i$ . The branch function simply adds the value of its argument to the return address, so that the return address becomes the address of the original target  $b_i$ . The code for this, on the Intel IA-32 architecture, might be as follows:<sup>4</sup>

```
xchg %eax, 0(%esp) # I1
add %eax, 8(%eax) # I2
pop %eax # I3
ret # I4
```

Instruction I1 exchanges the contents of register `%eax` with the word at the top of the stack, effectively saving the contents of `%eax` and at the same time loading the displacement to the target (passed to the branch function as an argument on the stack) into `%eax`. Instruction I2 then has the effect of adding this displacement to the return address. I3 restores the previously saved value of `%eax`, and I4 then has the effect of branching to the address computed by the function.

Our current implementation uses a variation on this idea that uses perfect hashing [14, 17] and is, we believe, harder to reverse engineer. Once the final code layout has been determined and we know the mapping  $\varphi = \{a_1 \mapsto b_1, \dots, a_n \mapsto b_n\}$  we want the branch function to implement, we create a perfect hash function  $h_\varphi$ :

$$h_\varphi : \{a_1, \dots, a_n\} \mapsto \{1, \dots, n\}$$

We then construct a table  $T$  in the data section of the binary, that lists offsets for each  $(a_i, b_i)$  pair, as follows:

$$T[h_\varphi(a_i)] \leftarrow b_i - a_i \quad [\text{for each } i].$$

Upon invocation the branch function proceeds as follows:

- (i) apply the perfect hash function  $h_\varphi$  to its return address  $a$  to compute a perfect hash value  $h_\varphi(a)$ ;
- (ii) use the table  $T$  to obtain the offset  $T[h_\varphi(a)]$  to the target;
- (iii) add this offset to the return address;
- (iv) return.

Since the resulting code is quite a bit more expensive, in execution cost, than the single branch instruction in the original program, we use execution profile information to apply the transformation only to code that is not “hot,” i.e., that is not frequently executed; the details are discussed in Section 4.

<sup>4</sup>If any of the condition code flags is live at the call point, they have to be saved by the caller just before the call, and restored at the target.




Figure 5: Branch functions

The complexity of a branch function’s implementation, and the way in which it is accessed, offer an interesting tradeoff between execution speed, on the one hand, and difficulty of reverse engineering, on the other. For example, we can choose different branch function implementations for jump instructions depending on their execution frequencies: frequently executed jump instructions might be directed to a lightweight branch function, less frequently executed ones to a more complex branch function, and so on.

### 3.4.2 Call Conversion

A variation of the branch function scheme described in section 3.4.1 can be used to extend the candidates for junk insertions to include those basic blocks directly following *call* instructions as well. Recall that the reason junk bytes can typically not be inserted after call instructions is that control returns to the address directly after the last byte of a call instruction upon completion of an invoked function. This being the case, if junk bytes were inserted after a “well-behaved” call instruction then it would be possible for control to reach the junk bytes and therefore violate the constraints of junk insertions described in section 3.2.

One solution to this problem is to reroute call instructions through a specialized branch function that branches to the intended target function via perfect hashing, as in the standard branch function, but then returns to some predetermined offset from the original call instruction (i.e., the offset to the real successor instruction that lies beyond some number of junk bytes). Using this method we are able to obscure control flow information by making function entry points more difficult to decipher while also increasing the potential to mislead the disassembler.

### 3.4.3 Opaque Predicates

The assumption that a conditional branch has two possible targets can be exploited by disguising an unconditional branch as a conditional branch that happens, at runtime, to always go in one direction—i.e., either it is always taken, and never falls through; or it is never taken, and always falls through. This technique relies on using predicates that always evaluate to either the constant *true* or the constant *false*, regardless of the values of their inputs: such predicates are known as *opaque predicates* [12]. Other researchers have discussed techniques for synthesizing opaque predicates; their ideas translate in a straightforward way to our context, so we do not discuss this issue further.

Once an unconditional branch has been replaced by a conditional branch that uses an opaque predicate, we have a location—either the branch target or the fall through, depending on whether the opaque predicate is always false or always true—that appears to be a legitimate continuation for execution from the conditional branch but, in fact, is not. We can then insert junk bytes at this point, as discussed earlier, to mislead the disassembly.

### 3.4.4 Jump Table Spoofing

In addition to simply inserting junk bytes at the fake target of an opaquely directed conditional branch, we can also insert artificial jump tables to mislead recursive traversal disassembly. We refer to this technique as *jump table spoofing*.

Recall that, as discussed in Section 2.2, recursive traversal disassemblers may attempt to use the bounds check for a jump table to identify its size, and thereby determine the set of possible targets of an indirect jump through a jump table. We can exploit this to mislead a disassembler by introducing a jump table that is unreachable at runtime. The jump table code can be made unreachable either by using a conditional branch that uses an opaque predicate to jump around it, or by using an opaque expression—which is very similar to an opaque predicate, expect that the value need not be simply a truth value—whose value is guaranteed to fail the bounds check. The code addresses in this jump table can now be set to “junk addresses”—text segment addresses that do not correspond to actual instructions—and thereby cause disassembly errors.

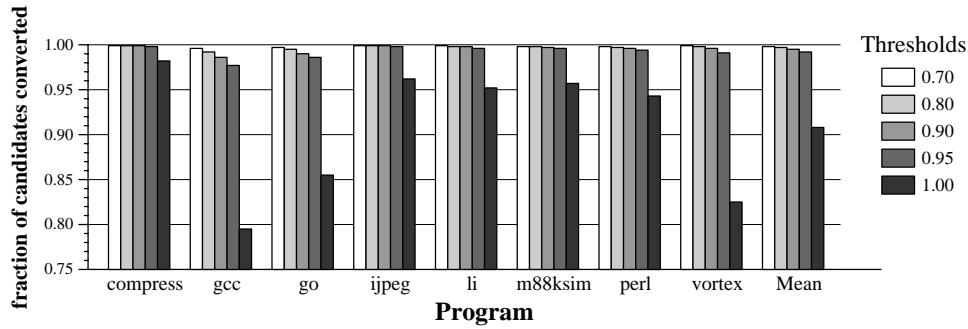
A variation on this idea is to take an unconditional jump to an address  $\ell$  and convert it to an indirect jump through a jump table where the address  $\ell$  appears as the  $k^{\text{th}}$  table entry. The table is indexed by the value of an opaque expression that always evaluates to  $k$ . However, the bounds check for the table uses a table size  $m > k$ , leading the disassembler to believe that the jump table contains  $m$  entries. Only one of these  $m$  entries—namely, the  $k^{\text{th}}$  entry—contains a real code address: the other entries contain junk addresses.

## 3.5 Implementation Status

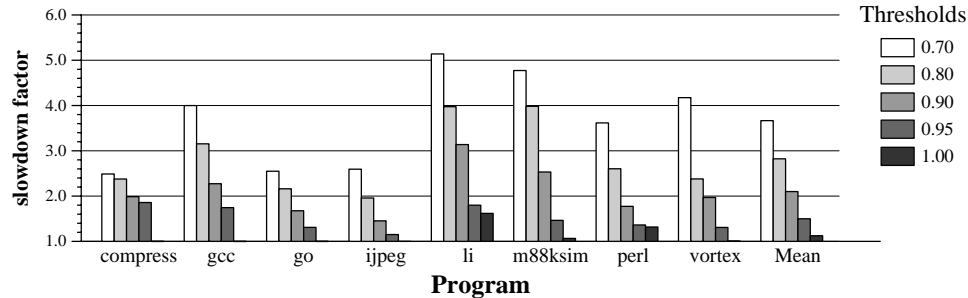
We have implemented our ideas using PLTO, a binary rewriting system developed for Intel IA-32 executables [21]. The system reads in statically linked executables,<sup>5</sup> disassembles the input binary, and constructs a control flow graph. This control flow graph is then processed in one of two ways. If the user specifies that profiling is to be carried out, instrumentation code is inserted to generate an edge profile when the resulting binary is executed. If, on the other hand, the user requests obfuscating transformations to be carried out, the system reads in edge profile information if available, carries out branch flipping to increase the number of candidate blocks (Section 3.3), applies various obfuscating transformations, and writes out the resulting executable.

The transformations currently implemented in the system are junk insertion (Section 3.2) and transformation of unconditional jumps and call instructions to the respective branch function calls

<sup>5</sup>The requirement for statically linked executables is a result of the fact that PLTO relies on the presence of relocation information to distinguish addresses from data. The Unix linker `ld` refuses to retain relocation information for executables that are not statically linked.



(a) Fraction of jumps converted to branch function calls



(b) Slowdown in execution speed

Figure 6: Effect of “hot code threshold” on branch function conversion and execution speed

(Sections 3.4.1 and 3.4.2). We expect to have additional transformations, such as jump table spoofing (Section 3.4.4), implemented in the near future.

## 4. EXPERIMENTAL EVALUATION

We evaluated the efficacy of our techniques using the SPECint-95 benchmark suite. Our experiments were run on an otherwise unloaded 2.4 GHz Pentium IV system with 1 GB of main memory running RedHat Linux 8.0. The programs were compiled with *gcc* version *egcs-2.91.66* at optimization level `-O3`. The programs were profiled using the SPEC training inputs and these profiles were used to identify any hot spots during our transformations. The final performance of the transformed programs were then evaluated using the SPEC reference inputs. Each execution time reported was derived by running seven trials, removing the highest and lowest times from the sampling, and averaging the remaining five.

We experimented with three different “attack disassemblers” to evaluate our techniques. The first of these is the GNU *objdump* utility which employs a straight-forward linear sweep algorithm. The second, which we wrote ourselves, is a recursive disassembler that incorporates a variation of speculative disassembly (see Section 2). In addition we also provide the recursive disassembler with extra information about the address and size of each jump table in the program as well as the start and end address of each function. The results obtained from this disassembler therefore serve as a lower bound estimate of the extent of obfuscation achieved. Our third disassembler is IDA Pro [13], a commercially available disassembly tool that is generally regarded to be among the best disassemblers available.

For each of these, the efficacy of obfuscation was measured by computing “confusion factors” for the instructions, basic blocks,

and functions. Intuitively, the confusion factor measures the fraction of program units (instructions, basic blocks, or functions) in the obfuscated code that were incorrectly identified by a disassembler. More formally, let  $A$  be the set of all *actual* instruction addresses, i.e., those that would be encountered when the program is executed, and  $P$  the set of all *perceived* instruction addresses, i.e., those addresses produced by a static disassembly.  $A - P$  is the set of addresses that are not correctly identified as instruction addresses by the disassembler. We define the *confusion factor*  $CF$  to be the fraction of instruction addresses that the disassembler fails to identify correctly:<sup>6</sup>

$$CF = |A - P| / |A|.$$

Confusion factors for functions and basic blocks are calculated analogously: a basic block or function is counted as being “incorrectly disassembled” if any of the instructions in it is incorrectly disassembled. The reason for computing confusion factors for basic blocks and functions as well as for instructions is to determine whether the errors in disassembling instructions are clustered in a small region of the code, or whether they are distributed over significant portions of the program.

As mentioned in Section 3.4.1, we transform jumps to branch function calls only if the jump does not occur in a “hot” basic block. The first questions we have to address, therefore, are: how are hot basic blocks identified, and what is the effect of different choices of what constitutes a “hot” block on the extent of obfuscation achieved and the performance of the resulting code? To identify the “hot,” or

<sup>6</sup>We also considered taking into account the set  $P - A$  of addresses that are erroneously identified as instruction addresses by the disassembler, but rejected this approach because it “double counts” the effects of disassembly errors.

PROGRAM	Confusion factor (%)								
	LINEAR SWEEP (OBJDUMP)			RECURSIVE TRAVERSAL			COMMERCIAL (IDA PRO)		
	Instructions	Basic blocks	Functions	Instructions	Basic blocks	Functions	Instructions	Basic blocks	Functions
<i>compress95</i>	43.93	63.68	100.00	30.04	40.42	75.98	75.81	91.53	87.37
<i>gcc</i>	34.46	53.34	99.53	17.82	26.73	72.80	54.91	68.78	82.87
<i>go</i>	33.92	51.73	99.76	21.88	30.98	60.56	56.99	70.94	75.12
<i>jpeg</i>	39.18	60.83	99.75	25.77	38.04	69.99	68.54	85.77	83.94
<i>li</i>	43.35	63.69	99.88	27.22	38.23	76.77	70.93	87.88	84.91
<i>m88ksim</i>	41.58	62.87	99.73	24.34	35.72	77.16	70.44	87.16	87.16
<i>perl</i>	42.34	63.43	99.75	27.99	39.82	76.18	68.64	84.62	87.13
<i>vortex</i>	33.98	55.16	99.65	23.03	35.61	86.00	57.35	74.55	91.29
<i>Geo. mean</i>	39.09	59.34	99.75	24.76	35.69	74.43	65.45	81.40	84.97

Figure 7: Efficacy of obfuscation: confusion factors ( $\theta = 1.0$ )

“frequently executed,” basic blocks, we start with a (user-defined) fraction  $\theta$  ( $0.0 < \theta \leq 1.0$ ) that specifies what fraction of the total number of instructions executed at runtime should be accounted for by “hot” basic blocks. For example,  $\theta = 0.8$  means that hot blocks should account for at least 80% of all the instructions executed by the program. More formally, let the *weight* of a basic block be the number of instructions in the block multiplied by its execution frequency, i.e., the block’s contribution to the total number of instructions executed at runtime. Let *tot\_instr\_ct* be the total number of instructions executed by the program, as given by its execution profile. Given a value of  $\theta$ , we consider the basic blocks  $b$  in the program in decreasing order of execution frequency, and determine the largest execution frequency  $N$  such that

$$\sum_{b: \text{freq}(b) \geq N} \text{weight}(b) \geq \theta \cdot \text{tot\_instr\_ct}.$$

Any basic block whose execution frequency is at least  $N$  is considered to be hot.

The effect of varying the hot code threshold  $\theta$  on performance (both obfuscation and speed) is shown in Figure 6. Figure 6(a) shows the fraction of candidates that are converted to branch function calls at different thresholds; this closely tracks the overall confusion factors achieved. Figure 6(b) shows the concomitant degradation in performance. It can be seen, from Figure 6(a), that most programs have a small and well-defined hot spot, and as a result varying the threshold from a modest 0.70 to a value as high as 1.0 does not dramatically affect the number of candidates converted. The benchmark that is affected the most is *gcc*, and even here over 79% of the candidates are converted at  $\theta = 1.0$ . On average, about 91% of the candidates are converted at  $\theta = 1.0$ . However, as illustrated in Figure 6(b), varying the hot code threshold has a significant effect on execution speed. For example, at  $\theta = 0.70$  the programs slow down by a factor of 3.67 on average, with the *li* benchmark experiencing the largest slowdown, by a factor of 5.14. However, as  $\theta$  is increased the slowdown drops off quickly, to a factor of 3.14 at  $\theta = 0.9$  and 1.62 at  $\theta = 1.0$ . In summary, choosing a threshold  $\theta$  of 1.0 still results in most of the candidate blocks in the program being converted to branch function calls without excessive performance penalty. For the purposes of this paper, therefore, we give measurements for  $\theta = 1.0$ .

Figure 7 shows the efficacy of our obfuscation transformations for both of the disassembly methods discussed in Section 2. The confusion factors achieved for linear sweep disassembly are quite modest: on average, 39% of the instructions, 59% of the basic blocks, and nearly 100% of the functions are incorrectly disassembled. For recursive traversal, the confusion factors are somewhat lower because in this case the disassembler can understand and deal with control flow somewhat better than with linear sweep and as a

PROGRAM	EXECUTION TIME (SECS)		
	Original ( $T_0$ )	Obfuscated ( $T_1$ )	Slowdown ( $T_1/T_0$ )
<i>compress95</i>	34.49	34.44	1.00
<i>gcc</i>	23.27	23.23	1.00
<i>go</i>	53.17	53.08	1.00
<i>jpeg</i>	40.13	40.15	1.00
<i>li</i>	26.50	42.91	1.62
<i>m88ksim</i>	28.18	30.02	1.07
<i>perl</i>	28.62	37.71	1.32
<i>vortex</i>	48.84	49.05	1.00
<i>Geo. mean</i>			1.13

Figure 8: Effect of obfuscation on execution speed ( $\theta = 1.0$ )

last resort actually reverts to linear sweep for the speculative disassembly of undisassembled code. Nevertheless, we find that, on average, over 25% of the instructions in the program incur disassembly errors. As a result, over 35% of the basic blocks and close to 74% of the functions, on average, are incorrectly disassembled using this disassembly method. This is achieved at the cost of a 13% penalty in execution speed (see Figure 8).

The recursive traversal data reported in Figure 7 are actually quite conservative since these were gathered using our own recursive disassembler which, as mentioned before, is supplied with extra information to avoid unduly optimistic results. To evaluate the efficacy of our techniques in a more realistic situation, we used a commercial disassembly tool, IDA Pro version 4.3x [13], which is widely considered to be the most advanced disassembler available. The results of this experiment are reported in Figure 7. It can be seen that this tool fails on most of the program: close to 65% of the instructions, and about 85% of the functions in the program, are disassembled incorrectly. Part of the reason for this high degree of failure is that IDA Pro only disassembles addresses that (it believes) can be guaranteed to be instruction addresses. This has two effects: first, large portions of the code that are reached by branch function addresses are simply not disassembled, being presented instead to the user as a jumble of hex data; and second, the location immediately following a branch function call is treated as an address to which control returns, and this causes some junk bytes to be erroneously disassembled. Overall, this shows that our techniques are effective even against state-of-the-art disassembly tools.

Finally, Figure 9 shows the impact of obfuscation on code size, both in terms of the number of instructions (which increases, for example, due to branch flipping), as well as the number of bytes occupied by the text section. The latter includes the effects of the new instructions inserted as well as all junk bytes added to the program. Overall, it can be seen that there is a 20% increase in the



PROGRAM	NO. OF INSTRUCTIONS			TEXT SECTION SIZE (BYTES)		
	Original ( $I_0$ )	Obfuscated ( $I_1$ )	Change ( $I_1/I_0$ )	Original ( $S_0$ )	Obfuscated ( $S_1$ )	Change ( $S_1/S_0$ )
<i>compress95</i>	74787	92137	1.231	265985	311095	1.169
<i>gcc</i>	327133	387289	1.183	1128273	1290419	1.143
<i>go</i>	124424	145953	1.173	468537	525232	1.121
<i>jpeg</i>	105766	127012	1.200	363169	419535	1.155
<i>li</i>	89309	109652	1.227	310301	363801	1.172
<i>m88ksim</i>	104211	127358	1.222	368798	430845	1.168
<i>perl</i>	137947	169054	1.225	484194	566935	1.170
<i>vortex</i>	174960	204230	1.167	592076	672795	1.136
<i>Geo. mean</i>			1.204			1.154

Figure 9: Effect of obfuscation on code size ( $\theta = 1.0$ )

total number of instructions, and a 15% increase in the size of the text section of the resulting executables.

The techniques described here apply to a wide variety of architectures. The insertion of partial instructions to confuse disassembly, as discussed in Section 3.2, is applicable to variable-length instruction sets, such as those on the widely used Intel Pentium and Motorola 680x0, as well as the StrongArm (together with the Thumb 16-bit instruction encoding) and other mixed-mode architectures such as the MIPS32/MIPS16. Branch functions and jump table spoofing can be used on any architecture.

## 5. RELATED WORK

The only work we are aware of that addresses the problem of making executable programs harder to disassemble is by Cohen, who proposes overlapping adjacent instructions to fool a disassembler [7]. We are not aware of any actual implementations of this proposal. We implemented this idea as well as a number of variations on the basic scheme, but found the results disappointing: the resulting confusion factors were typically less than 1%. The reason for this is that in order to overlap two adjacent instructions  $I$  and  $J$ , we have to satisfy several conditions, among them:

- (i) execution cannot fall through from  $I$  to  $J$ ; and
- (ii) the trailing  $k$  bytes of  $I$  must be identical with the leading  $k$  bytes of  $J$  for some  $k > 0$ .

There tend to be relatively very few candidates satisfying these criteria (e.g., the largest number of overlaps we achieved was for the *gcc* benchmark, where we found only 27 overlaps out of 360,152 instructions; by contrast, our approach can use 11,205 candidates before branch flipping and call conversion on this program, and 56,925 candidates after branch flipping and call conversion). Variations on this theme, e.g., by judicious insertion, immediately before the instruction  $J$ , of no-ops or dead code that satisfy the second condition above, do not seem to help matters significantly either. This scarcity of candidates for overlapping, together with the self-repairing property of disassembly errors discussed in Section 3.1, results in poor confusion factor numbers using this approach.

There is a considerable body of work on code obfuscation that focuses on making it harder for an attacker to decompile a program and extract high level semantic information from it [3, 12, 20, 26, 27, 28]. Typically, these authors rely on the use of computationally difficult static analysis problems, e.g., involving complex Boolean expressions, pointers, or indirect control flow, to make it harder to construct a precise control flow graph for a program. Of the references cited, only Wroblewski focuses specifically on obfuscation of executable programs [28]. Our work is orthogonal to these proposals, and complementary to them. We aim to make a program

harder to disassemble correctly, and to thereby sow uncertainty in an attacker's mind about which portions of a disassembled program have been correctly disassembled and which parts may contain disassembly errors. If the program has already been obfuscated using any of these higher-level obfuscation techniques, our techniques add an additional layer of protection that makes it even harder to decipher the actual structure of the program.

Even greater security may be obtained by maintaining the software in encrypted form and decrypting it as needed during execution, as suggested by Aucsmith [1]; or using specialized hardware, as discussed by Lie *et al.* [16]. Such approaches have the disadvantages of high performance overhead (in the case of runtime decryption in the absence of specialized hardware support) or a loss of flexibility because the software can no longer be run on stock hardware.

## 6. CONCLUSIONS

A great deal of software is distributed in the form of executable code. Such code is potentially vulnerable to reverse engineering, in the form of disassembly followed by decompilation. This can allow an attacker to discover vulnerabilities in the software, modify it in unauthorized ways, or steal intellectual property via software piracy. This paper describes and evaluates techniques to make executable programs harder to disassemble. Our techniques are seen to be quite effective: applied to the widely used SPECint-95 benchmark suite, they result in disassembly errors over much of the program; the best commercially available disassembly tool fails to correctly disassemble over 65% of the instructions, and 85% of the functions, in the obfuscated binaries.

## Acknowledgements

We are grateful to Christian Collberg and Gregg Townsend for very helpful discussions and comments.

## 7. REFERENCES

- [1] D. Aucsmith. Tamper-resistant software: An implementation. In *Information Hiding: First International Workshop: Proceedings*, volume 1174 of *Lecture Notes in Computer Science*, pages 317–333. Springer-Verlag, 1996.
- [2] R. L. Bernstein. Producing good code for the case statement. *Software—Practice and Experience*, 15(10):1021–1024, October 1985.
- [3] W. Cho, I. Lee, and S. Park. Against intelligent tampering: Software tamper resistance by extended control flow obfuscation. In *Proc. World Multiconference on Systems, Cybernetics, and Informatics*. International Institute of Informatics and Systematics, 2001.

- [4] C. Cifuentes and K. J. Gough. Decompilation of binary programs. *Software—Practice and Experience*, 25(7):811–829, July 1995.
- [5] C. Cifuentes and M. Van Emmerik. UQBT: Adaptable binary translation at low cost. *IEEE Computer*, 33(3):60–66, March 2000.
- [6] C. Cifuentes and M. Van Emmerik. Recovery of jump table case statements from binary code. *Science of Computer Programming*, 40(2–3):171–188, July 2001.
- [7] F. B. Cohen. Operating system protection through program evolution, 1992. <http://all.net/books/IP/evolve.html>.
- [8] R. S. Cohn, D. W. Goodwin, and P. G. Lowney. Optimizing Alpha executables on Windows NT with Spike. *Digital Technical Journal*, 9(4):3–20, 1997.
- [9] C. Collberg and C. Thomborson. Software watermarking: Models and dynamic embeddings. In *Proc. 26th. ACM Symposium on Principles of Programming Languages (POPL 1999)*, pages 311–324, January 1999.
- [10] C. Collberg and C. Thomborson. Watermarking, tamper-proofing, and obfuscation – tools for software protection. Technical Report TR00-03, The Department of Computer Science, University of Arizona, February 2000.
- [11] C. Collberg, C. Thomborson, and D. Low. Breaking abstractions and unstructuring data structures. In *Proc. 1998 IEEE International Conference on Computer Languages*, pages 28–38.
- [12] C. Collberg, C. Thomborson, and D. Low. Manufacturing cheap, resilient, and stealthy opaque constructs. In *Proc. 25th. ACM Symposium on Principles of Programming Languages (POPL 1998)*, pages 184–196, January 1998.
- [13] DataRescue sa/nv, Liège, Belgium. IDA Pro. <http://www.datarescue.com/idadabase/>.
- [14] M. L. Fredman, J. Komlós, and E. Szemerédi. Storing a sparse table with  $O(1)$  worst case access time. *Journal of the ACM*, 31(3):538–544, July 1984.
- [15] J. R. Levine. *Linkers and Loaders*. Morgan Kaufman Publishers, San Francisco, CA, 2000.
- [16] D. Lie, C. Thekkath, M. Mitchell, P. Lincoln, D. Boneh, J. Mitchell, and M. Horowitz. Architectural support for copy and tamper resistant software. In *Proc. 9th. International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS-IX)*, pages 168–177, November 2000.
- [17] K. Mehlhorn and A. K. Tsakalidis. Data structures. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science, Volume A: Algorithms and Complexity (A)*, pages 301–341. MIT Press, 1990.
- [18] R. Muth, S. K. Debray, S. Watterson, and K. De Bosschere. `alto`: A link-time optimizer for the Compaq Alpha. *Software—Practice and Experience*, 31:67–101, January 2001.
- [19] Objdump. *GNU Manuals Online*. GNU Project—Free Software Foundation. [http://www.gnu.org/manual/binutils-2.10.1/html\\_chapter/binutils\\_4.html](http://www.gnu.org/manual/binutils-2.10.1/html_chapter/binutils_4.html).
- [20] T. Ogiso, Y. Sakabe, M. Soshi, and A. Miyaji. Software obfuscation on a theoretical basis and its implementation. *IEEE Trans. Fundamentals*, E86-A(1), January 2003.
- [21] B. Schwarz, S. K. Debray, and G. R. Andrews. `Plto`: A link-time optimizer for the Intel IA-32 architecture. In *Proc. 2001 Workshop on Binary Translation (WBT-2001)*, 2001.
- [22] B. Schwarz, S. K. Debray, and G. R. Andrews. Disassembly of executable code revisited. In *Proc. IEEE 2002 Working Conference on Reverse Engineering (WCRE)*, pages 45–54, October 2002.
- [23] R. L. Sites, A. Chernoff, M. B. Kirk, M. P. Marks, and S. G. Robinson. Binary translation. *Communications of the ACM*, 36(2):69–81, February 1993.
- [24] A. Srivastava and D. W. Wall. A practical system for intermodule code optimization at link-time. *Journal of Programming Languages*, 1(1):1–18, March 1993.
- [25] H. Theiling. Extracting safe and precise control flow from binaries. In *Proc. 7th Conference on Real-Time Computing Systems and Applications*, December 2000.
- [26] C. Wang, J. Davidson, J. Hill, and J. Knight. Protection of software-based survivability mechanisms. In *Proc. International Conference of Dependable Systems and Networks*, July 2001.
- [27] C. Wang, J. Hill, J. Knight, and J. Davidson. Software tamper resistance: Obstructing static analysis of programs. Technical Report CS-2000-12, 12 2000.
- [28] G. Wroblewski. *General Method of Program Code Obfuscation*. PhD thesis, Wroclaw University of Technology, Institute of Engineering Cybernetics, 2002.