

Protecting Software through Obfuscation: Can It Keep Pace with Progress in Code Analysis?

SEBASTIAN SCHRITTWIESER, St. Pölten University of Applied Sciences, Austria
STEFAN KATZENBEISSER, Technische Universität Darmstadt, Germany
JOHANNES KINDER, Royal Holloway, University of London, United Kingdom
GEORG MERZDOVNIK, SBA Research, Vienna, Austria
EDGAR WEIPPL, SBA Research, Vienna, Austria

Software obfuscation has always been a controversially discussed research area. While theoretical results indicate that provably secure obfuscation in general is impossible, its widespread application in malware and commercial software shows that it is nevertheless popular in practice. Still, it remains largely unexplored to what extent today's software obfuscations keep up with state-of-the-art code analysis, and where we stand in the arms race between software developers and code analysts. The main goal of this survey is to analyze the effectiveness of different classes of software obfuscation against the continuously improving de-obfuscation techniques and off-the-shelf code analysis tools.

The answer very much depends on the goals of the analyst and the available resources. On the one hand, many forms of lightweight static analysis have difficulties with even basic obfuscation schemes, which explains the unbroken popularity of obfuscation among malware writers. On the other hand, more expensive analysis techniques, in particular when used interactively by a human analyst, can easily defeat many obfuscations. As a result, software obfuscation for the purpose of intellectual property protection remains highly challenging.

Categories and Subject Descriptors: D.2.0 [Software Engineering]: General—*Protection mechanisms*

General Terms: Security

Additional Key Words and Phrases: software obfuscation, program analysis, reverse engineering, software protection, malware

ACM Reference Format:

Sebastian Schrittwieser, Stefan Katzenbeisser, Johannes Kinder, Georg Merzdovnik, and Edgar Weippl, 2015. Protecting Software through Obfuscation: Can It Keep Pace with Progress in Code Analysis? *ACM Comput. Surv.* 49, 1, Article 4 (April 2016), 40 pages.
DOI : <http://dx.doi.org/10.1145/0000000.0000000>

1. INTRODUCTION

The development of code obfuscation techniques is driven by the desire to hide the specific implementation of a program from automatic or human-assisted analysis of exe-

The research was funded under Grant 826461 (FIT-IT) and COMET K1 by the FFG – Austrian Research Promotion Agency, the Austrian Federal Ministry of Science, Research and Economy and the National Foundation for Research, Technology and Development, CASED, and CRISP. Author's addresses: S. Schrittwieser, Josef Ressel Center for Unified Threat Intelligence on Targeted Attacks, Department Computer Science & Security, St. Pölten University of Applied Sciences, Austria; S. Katzenbeisser, Security Engineering Group, Technische Universität Darmstadt, Germany; J. Kinder, Department of Computer Science, Royal Holloway, University of London, United Kingdom; G. Merzdovnik, SBA Research, Vienna, Austria; Edgar Weippl, SBA Research, Vienna, Austria.

Permission to make digital or hard copies of part or all 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 show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© 2016 ACM 0360-0300/2016/04-ART4 \$15.00

DOI : <http://dx.doi.org/10.1145/0000000.0000000>

cutable code. Its roots date back to the early days of programming when the main purpose of obfuscation was to place hidden functionality into programs—often just for hiding a secret message to surprise users who reveal it by coincidence (e.g., Easter eggs). Particularly creative ways to obscure functionality were even rewarded in competitions like the International Obfuscated C Contest, which has been held since 1984¹.

By the late 1980s, malware became a major reason for the steady refinement of code obfuscation techniques. For instance, in 1986 the Brain computer virus, which is believed to be the first computer virus for MS-DOS, obfuscated the functionality of its code by intercepting reads of the virus binary to return innocuous code [Avoine et al. 2007]. These early occurrences of software obfuscation marked the beginning of an arms race between developers and analysts. Since then, software protection researchers have been developing more and more sophisticated obfuscation techniques to hide behavior of code, while analysts have been using increasingly complex code analysis techniques to defeat obfuscations.

The underground economy and its demand for stealthy malware is still one of the major driving forces behind the development of obfuscation techniques. The other leading application area of code obfuscation is the protection of commercial software against reverse engineering [Grover 1992]. The motivation for reverse engineering can be diverse: an analyst might be interested in extracting some secret information from the program code that should not be revealed to the user. This secret might be a cryptographic key, a sophisticated algorithm considered a trade secret, or credentials for a remote service. Typical examples include media players that store secret keys or incorporate proprietary algorithms like the CSS (content scramble system) cipher to play back copy-protected content. Another motivation for reverse engineering is the modification of software to change its behavior. An analyst might want to use hidden program functionality, uncover firmware features that are disabled in low-cost devices, or interface commercial software with own products; all these actions will likely interfere with the business model of the software vendor.

Techniques originally developed for malware have experienced a comeback in the field of attack prevention. Since the early 1990's polymorphic engines had been widely used in malware development to generate differently looking versions of malicious code to evade signature-based detection [Nachenberg 1997; Song et al. 2010]. Software diversity is a related concept for delivering syntactically different but semantically identical versions of applications to different users [Franz 2010; Davi et al. 2012]. This way, automated attacks developed against one instance of a program are thus less likely to work against a different obfuscated version [Forrest et al. 1997; Anckaert et al. 2004, 2006].

On the analyst's side, research in the area of code analysis progressed significantly over the past years. Disassemblers, which extract the executable assembly code from binaries, became increasingly sophisticated, implementing complex heuristics or expensive static analysis methods to reconstruct code from a potentially obfuscated binary as precisely as possible. In addition, research tools such as JAKSTAB [Kinder and Veith 2008] or BAP [Brumley et al. 2011] directly analyze binary code to precisely reconstruct control flow and reason about the behavior of the code. Systems for dynamic malware analysis, such as Anubis [Bayer et al. 2006], allow a detailed analysis of the runtime behavior of a piece of code.

Since the groundbreaking work of Collberg et al. [1997] on a taxonomy of obfuscation techniques, a vast number of code obfuscation schemes have been proposed in the literature. Nevertheless, the security and effectiveness of these techniques remain controversially discussed topics. Indeed, early theoretical results for obfuscation that

¹<http://www.ioccc.org/1984/> (accessed February 25, 2015)

is secure in a cryptographic sense were negative. Barak et al. [2001] showed that it is impossible to construct a universal obfuscator that is applicable to any program; positive results are known for very restricted classes of functionality such as point functions [Lynn et al. 2004; Wee 2005; Canetti and Dakdouk 2008]. Other works on theoretical aspects of code obfuscation are due to Dalla Preda and Giacobazzi [2005], Goldwasser and Rothblum [2007], and Giacobazzi [2008]. While addressing the theoretical understanding of obfuscation, these results have little to say about the practical effectiveness of obfuscation. Recently, Garg et al. [2013] presented a first candidate for general-purpose obfuscation (*indistinguishability obfuscation*). While it is unclear how useful this progress is to hide secrets in a program in practical settings, it still brought new momentum to the discussion and encouraged follow-up work [Brakerski and Rothblum 2014; Barak et al. 2014; Bitansky et al. 2014].

Code obfuscation is widely employed in practice: for example, almost all newly discovered malware comes with some form of obfuscation to hide its functionality, and commercial software products such as Skype or DRM (Digital Rights Management) engines use obfuscation techniques as part of their protection portfolio. Consequently, a plethora of obfuscation techniques have been described in the literature (e.g., [Collberg et al. 1997; Wang et al. 2000; Dedić et al. 2007; Jakubowski et al. 2007]); most of them try to “raise the bar” for reverse engineering attempts, but do not come with a rigorous security analysis, let alone a proof. Some attempts have been made to quantify the additional complexity that is added to an executable by employing code obfuscation techniques (for example, see [Madou et al. 2006]); however it remains unclear whether such notions indeed capture all security properties of obfuscation correctly. While this is clearly unsatisfactory from a theoretic point of view, it is probably the best one can achieve with current knowledge.

Still, it remains largely unexplored to what extent today’s code obfuscation techniques can keep up with the progress in code analysis and where we stand in the arms race between developers and analysts. The main goal of this survey is thus to provide a comprehensive picture of the state of the art in code obfuscation and of its effectiveness against de-obfuscation techniques and code analysis tools that are available today.

To this end, we first build a classification of analysis scenarios in Section 2; in particular, we pair each analysis method with a specific analysis goal to define a set of *scenarios* that form our basic model of the code analyst. Subsequently, we describe and categorize existing code obfuscation techniques in Section 3. There is a large body of literature on obfuscation and analysis, and we do not claim to survey it in all its breadth; instead, we try to paint a representative picture of the state of the art in both fields.

In all of our analysis scenarios, we assume that the analyst receives binary or byte code. Thus, we limit ourselves to obfuscations that *affect* the binary or byte code representation of a program. We make no distinction whether the obfuscation is *applied* before, during, or after compilation, however. We decided to leave obfuscation techniques targeting the readability of source code (such as the removal of comments or renaming of variables) out of scope, even though we acknowledge the importance of such obfuscations for languages such as JavaScript.

Section 4 discusses the state of the art in program analysis methods and particularly focuses on their capabilities and limits. Section 5 uses the described classifications to assess the security of code obfuscation techniques against the different scenarios, taking into account recently published attacks as well as recent off-the-shelf program analysis tools. We rely on published results where possible and apply our own judgment where we could not find literature addressing a particular scenario. Finally, Section 6 concludes the paper.

2. ANALYSIS SCENARIOS

In this section, we detail our methodology and introduce the analysis scenarios we use throughout the rest of the paper to assess the strength of obfuscation techniques against specific adversaries with specific goals. This allows us to characterize the protection a developer may want to achieve at a finer granularity than the black-box notion typically used in the formal analysis of software obfuscation [Barak et al. 2001].

2.1. Methodology of this Survey

We classify the feasibility of real-life analysis scenarios on programs in the presence of code obfuscation, based on a careful analysis of the literature on program analysis and reverse engineering. For our classification, we use a number of analysis scenarios consisting of an analysis technique and a goal that an analyst wants to achieve.

We assume that a developer implements the strongest possible obfuscator of one particular class of obfuscations. We do not consider second order effects, i.e., combinations of different types of obfuscations. A systematic analysis of how obfuscation techniques compose with one another would be challenging already just due to the sheer number of possible combinations. Still, it certainly makes an interesting and relevant topic for future work, since many commercial obfuscators employ (and indeed recommend to use) multiple obfuscations at the same time.

We make a similar assumption on the side of the analyst and evaluate the strength of obfuscations against four code analysis categories (introduced in Section 2.2) individually. Approaches that involve the staged application of multiple techniques are represented by the category of their first analysis stage; for instance, pattern matching on disassembled code requires an at least partially successful disassembly by static analysis. Hybrid techniques that unite aspects of multiple categories are sorted into the category they share the most properties with. Our strongest category involves a human analyst who may use multiple techniques of different types; our model here assumes that the analyst will apply techniques from either category and switch to a different one if they do not make any more progress towards the analysis goal.

As stated above, our evaluation of the effectiveness of different classes of software obfuscation schemes against the defined analysis scenarios is based on a literature review. For each combination of an analysis scenario and a specific class of obfuscations (see Section 3), we sought out literature specifically addressing that combination. Sometimes we could directly report published results of analyzing particular obfuscations, sometimes we had to make our own inferences from the available information. Our results are not derived from a formal analysis and are necessarily open to interpretation. Nevertheless, we believe that our scenarios capture the relevant analysis context and that our classification presents an accurate snapshot of the current state in the arms race between code obfuscation and analysis.

2.2. Code Analysis Categories

We categorize code analysis techniques in four general classes. An analyst can use different analyses depending on their goal and their available resources and time. For example, a human reverse engineer who tries to understand a piece of code of a competitor may afford spending time and effort on highly complex and time-consuming analyses; in contrast, an anti-virus vendor, who has to timely analyze hundreds of thousands of different malware samples each day, may be required to resort to very lightweight and thus limited analysis techniques that at the same time ensure a low false positive rate.

Pattern matching. Pattern matching is the simplest and fastest form of code analysis; it is a syntactic analysis performed on the program binary or byte code. Techniques

in this category include identification of static sequences of instructions, regular expressions, or machine learning-based classifiers for binary data (see Section 4.1).

Automated static analysis. Static analysis inspects a program (on binary or byte code level) without actually executing it. In contrast to syntactic pattern matching, static analysis reasons about the program semantics. Simple forms of static analysis include disassemblers that interpret branch targets. Static analysis is frequently used to reconstruct high-level information about programs, such as their control flow graph (see Section 4.2).

Automated dynamic analysis. Dynamic analysis runs a program to observe its actions and/or collect its flow of information. Dynamic analysis has the advantage of being able to very precisely reason about the program behavior along the observed traces. However, the data gathered from one or several runs of a program does not necessarily allow to draw conclusions about the behavior of the entire program (see Section 4.3).

Human-assisted analysis. As the most capable “analysis”, we consider a human analyst who performs a tool-assisted exploration of a piece of code; the tools can be based on all of the other analysis types. Typically, this process is referred to as reverse engineering, where the analyst aims at understanding the program structure and behavior with the help of a variety of tools (see Section 4.4).

2.3. Analyst’s Aims

We now systematically categorize and characterize the possible motivations of an analyst for analyzing software, in order of increasing complexity. We believe that almost all analyst goals observed in practice fit in one of these four general categories. As a running example to illustrate the goals throughout the section, we use a program implementing a cryptographic algorithm with an embedded secret key. Furthermore, we present other examples for each category to demonstrate the practicability of this classification.

Finding the location of data. The analyzer wishes to retrieve some data embedded in the program in its original, non-obfuscated representation from the obfuscated program. In our running example, the analyzer may want to extract the secret cryptographic key from the obfuscated program to be able to decrypt data in a different context than provided by the application (e.g., to circumvent DRM policies). Other typical examples that fall into this category are the extraction of licensing keys, certificates, credentials for remote services, and device configuration data.

Finding the location of program functionality. The analyzer aims at identifying the entry point of a particular function within an obfuscated program. In our running example, the analyzer may want to find the entry point of a cryptographic algorithm in the obfuscated program to subsequently analyze it. Another aim could be to find the exact location of a copy protection mechanism (such as a check for the presence of a hardware dongle or the validation of a licensing key) in order to circumvent it. Furthermore, finding the code representation of a particular functionality of a program can be useful for manual reverse engineering on small areas of the program. More generally, one may ask the question whether a program implements a particular functionality at all (such as the AES encryption algorithm) or simply whether a program is malicious or not.

Another related aim of the analyst might be to modify the behavior of a program in a particular way (e.g., bypassing a copy protection mechanism). However, we consider this to be out of scope of this work, because it falls into the field of tamper-proofing

rather than obfuscation (e.g., [Horne et al. 2002] and [Chang and Atallah 2002]). Still, *finding the location of program functionality* is a fundamental prerequisite of this aim.

Extraction of code fragments. The analyzer aims at extracting a piece of code including all possible dependencies that implement a particular functionality from the obfuscated program. In our running example, the analyzer's aim can be the extraction of the cryptographic algorithm in order to build a custom decryption routine. Note that for this purpose it is not necessary to fully understand the code, just using it in a new application may be enough. This approach has been used for breaking DRM implementations. Instead of understanding how exactly the decryption routine embedded into a player works, it is simply extracted and included in a counterfeit player, which decrypts the digital media without enforcing the contained usage policies. Another aim of the analyst may be the extraction of fragments of commercial software of competitors. In some analysis scenarios code does not need to be extracted, but a functional component of an executable is accessed at runtime (*in-situ reuse*) [Miles et al. 2012]. This technique is comparable to the concept of return oriented programming, which reuses code fragments to create new functionality [Shacham 2007].

Understanding the program. The analyzer wishes to fully understand a code fragment or even the entire obfuscated program. This requires that the analyst must be able to remove the applied obfuscation techniques and gain full understanding of the original, non-obfuscated program or a non-trivial fragment of it. In our running example, the analyzer may want to understand how a proprietary cipher embedded into the obfuscated program works in order to start cryptanalysis attempts. Other motivations for trying to understand a program can be the desire of an analyst to find vulnerabilities, to correct flaws in software for which the source code is not available, and create new programs that are compatible with proprietary software. Finally, a major driving force for human-assisted reverse engineering is to gain understanding of proprietary implementation details such as file formats or protocols, which often constitutes intellectual property theft.

2.4. Scenarios

By combining each code analysis technique with each of the analyst's aims, we arrive at the analysis scenarios that form the basis of our survey.

Not all combinations are reasonable according to our definitions. While pure pattern matching can clearly help to locate a code fragment, its extraction will then either (a) require additional static analysis for dependency analysis and thus fall into another category, or (b) be trivial and thus identical in difficulty. Similarly, pattern matching does not lend itself to code understanding, so we leave out both combinations and consider a total of 14 scenarios. An overview of the literature that describes code analysis techniques based on these scenarios is provided in Table I. In the following, we explain all 14 scenarios in more detail.

Locating data through pattern matching. A matching algorithm is used to determine the existence and location of data that conforms to a pattern specification. Patterns for the automated identification of data inside a program describe the structure of the data such as its length or data type, its environment in the form of the code surrounding it, or even properties such as its entropy.

Locating code through pattern matching. Particular code fragments are detected by pattern matching, e.g., on instruction sequences with or without wildcards. When pattern matching is applied to the outcome of a static or dynamic analysis (e.g., on disassembled or normalized code [Christodorescu et al. 2007]), we classify the analysis as static or dynamic, respectively.

Locating data through static analysis. A static analysis interprets the semantics of a program to locate particular data. For example, the reconstruction of the possible arguments of function calls can reveal a cryptographic key that is used to decrypt DRM-protected media or sent to an external device, or it could determine a token that is transmitted over the network.

Locating code through static analysis. A static analysis that determines code that could be executed at runtime. Examples are disassemblers that interpret control flow instructions using heuristics or formal methods. The control flow graph often allows to recognize particular structures; e.g., the control flow characteristics of a cryptographic algorithm can reveal its location in a binary.

Extracting code through static analysis. A static analysis technique (such as def-use chains) analyzes the dependencies of a piece of code to compute a self-contained slice that can be extracted and run on its own. Most static analysis techniques will result in extracting a superset of the required code (see Section 4.2).

Understanding code through static analysis. This scenario captures static de-obfuscation techniques that are able to transform fragments or all of the obfuscated code into a low- or high-level representation from which a human analyst can understand the program functionality with reasonable effort.

Locating data through dynamic analysis. A dynamic analysis observes the data accessed or stored by a program at runtime and uses these observations to locate particular data, e.g., in the parameters of system calls or particular memory locations.

Locating code through dynamic analysis. Dynamic analysis reveals the behavior of a program and its interaction with the environment (e.g., system calls) at runtime. In this scenario, a particular functionality is located through its specific runtime behavior.

Extracting code through dynamic analysis. A dynamic analysis locates a piece of code and also determines its dependencies along the observed traces. Because the code extracted from one run may not contain everything required for another run of the program, the resulting control flow graph is a subset of all possible execution paths (see Section 4.3).

Understanding code through dynamic analysis. Analogous to the scenario of understanding code through static analysis, this scenario targets automated dynamic de-obfuscation techniques that are able to transform the obfuscated code into a representation from which a human analyst can understand the program's functionality with reasonable efforts.

Human assisted analysis (four scenarios). Static and/or dynamic approaches assisted by a human analyst aim at getting full understanding of a particular aspect of the program. This aspect can be data in its pure form (data de-obfuscation), the location of code implementing a particular functionality or its dependencies, as well as the entire program itself.

3. SOFTWARE OBFUSCATION

In this section we briefly describe various code obfuscation schemes from the literature and classify them into the three categories *data obfuscation*, *static code rewriting*, and *dynamic code rewriting*. Many of the described obfuscation techniques appeared first in malware samples. Consequently, it is hard to pay tribute to the original source; we did this wherever possible. More details on early techniques are given by Collberg et al. [1997] and Collberg and Nagra [2009].

Table I. Literature on code analysis categorized based on the 14 scenarios.

	Pattern Matching	Automatic Static	Automatic Dynamic	Human Analysis
Locating Data	[Moser et al. 2007b]	[Shamir and Van Someren 1999] [Kinder and Veith 2008] [Kinder 2012]	[Cozzie et al. 2008] [Lin et al. 2010] [Slowinska et al. 2011] [Zhao et al. 2011]	[Jacob et al. 2003] [Link et al. 2004] [Billet et al. 2005] [Wyseur and Preneel 2005] [Goubin et al. 2007] [Piazalunga et al. 2007] [Wyseur et al. 2007] [Michiels et al. 2009] [Saxena et al. 2009] [Wyseur 2009] [De Mulder et al. 2010]
Locating Code	[Newsome et al. 2005] [Moser et al. 2007b] [Tang and Chen 2007] [Griffin et al. 2009] [Dalla Preda et al. 2011]	[Flake 2004] [Harris and Miller 2005] [Bruschi et al. 2006a] [Bruschi et al. 2006b] [Chouchane and Lakhota 2006] [Dalla Preda et al. 2006] [Walenstein et al. 2006] [Bilar 2007] [Karnik et al. 2007] [Nagarajan et al. 2007] [Popov et al. 2007] [Gao et al. 2008] [Coogan et al. 2009] [Treadwell and Zhou 2009] [Tsai et al. 2009] [Jacob et al. 2012] [Kinder 2012] [Bourquin et al. 2013]	[Deprez and Lakhota 2000] [Madou et al. 2005] [Zhang and Gupta 2005] [Royal et al. 2006] [Wilde and Scully 2006] [Moser et al. 2007a] [Brumley et al. 2008] [Sharif et al. 2008] [Song et al. 2008] [Guizani et al. 2009] [Li et al. 2009] [Sharif et al. 2009] [Webster and Malcolm 2009] [Comparetti et al. 2010] [Debray and Patel 2010] [Yin and Song 2010] [Coogan et al. 2011] [Gröbert et al. 2011] [Calvet et al. 2012] [Zeng et al. 2013]	[Madou et al. 2006b] [Madou et al. 2006c] [Roles 2009] [Quist and Liebrock 2009]
Extracting Code	<i>invalid scenario</i>	[Sneed 2000] [Christodorescu et al. 2007]	[Christodorescu et al. 2007] [Leder et al. 2009] [Sharif et al. 2009] [Caballero et al. 2010] [Kolbitsch et al. 2010] [Zeng et al. 2013]	[Ning et al. 1993] [Canfora et al. 1994] [Field et al. 1995] [Cimitile et al. 1996] [Lanubile and Visaggio 1997] [Canfora et al. 1998] [Danicic et al. 2004] [Fox et al. 2004] [Danicic et al. 2005]
Understanding Code	<i>invalid scenario</i>	[Rugaber et al. 1995] [Majumdar et al. 2006] [Raber and Laspe 2007] [Guillot and Gazet 2010]	[Udupa et al. 2005]	[Biondi and Desclaux 2006] [Eagle 2008] [Myska 2009] [Kholia and Wegrzyn 2013]

3.1. Data obfuscation

Code obfuscation techniques of this category modify the form in which data is stored in a program to hide it from direct analysis. Usually, data obfuscation also requires the program code to be modified, so that the original data representation can be reconstructed at runtime. Many data obfuscation techniques were first described by Collberg et al. [1998a].

Reordering of data. Variables can be split into two or more pieces in order to make it more difficult for an analyst to identify them. The mapping between an actual value of a variable and its split representation is managed by two functions. One is executed

at obfuscation time, the other one reconstructs the original value of a variable from its split parts at runtime. For example, Boolean variables can be obfuscated by splitting them into multiple Boolean values. At runtime the variable's actual value is retrieved by performing a specific Boolean operation (such as a logical XOR) over the parts of the variable. Other data types such as integers and string variables can be obfuscated in a similar way. In contrast to variable splitting, variable merging combines two or more variables into one.

To obfuscate the structure of an array it can be *split* into two or more subarrays. Conversely, multiple arrays can be *merged* into one. *Folding* (increasing the number of dimensions of the array) and *flattening* (decreasing the number of dimensions) are similar techniques which can be used for obfuscating data stored in arrays.

Obfuscating the structure of data by reordering its components to decrease locality (logically related items are physically close in the binary) is another fundamental obfuscation technique [Collberg et al. 1998a]. For example, this obfuscation is often applied to cryptographic keys stored within commercial software.

A low-level implementation of data reordering for obfuscation was introduced by Anckaert et al. [2009]. By redirecting memory accesses through a software-based dispatcher, the order of data in memory can be shuffled periodically, making its identification and analysis more difficult.

Encoding. Static data (such as strings) within binaries contains useful information for an analyst. Under this obfuscation technique data is converted to a different representation with some special encoding function to mitigate the need of storing the static data in cleartext in the binary. At runtime, the inverse function is used to decode the data. A variant of data encoding that works via mixed-mode computation over Boolean-arithmetic algebras was introduced by Zhou et al. [2007].

Converting static data to procedures. This obfuscation method replaces static data with a function that calculates the data at runtime. For example, a string object can be built at runtime, so that an analyst is not able to extract its value by examining the binary.

An extreme form of this obfuscation method is white-box cryptography. Its basic idea is to merge a secret key with elements of the cipher (e.g., the S-boxes), so that the key cannot be found in the binary anymore. The first white-box implementations of DES [Chow et al. 2003a] and AES [Chow et al. 2003a] were proposed by Chow et al., followed by other approaches [Link and Neumann 2005; Wyseur and Preneel 2005; Bringer et al. 2006]. However, all published white-box algorithms have been broken so far (for more details, please refer to Section 5).

3.2. Static code rewriting

A static rewriter is similar to an optimizing compiler, as it modifies program code during obfuscation, but allows its output to be executed without further run-time modifications. Strictly speaking, all data obfuscation techniques described above would also fall into the category of static code rewriting. However, as the obfuscation targets are distinct (data vs. binary code) and require specific obfuscation techniques, we use separate categories for data obfuscation and static code rewriting. In malware obfuscation the term *metamorphism* describes automated mutation of binary code through static code rewriting techniques applied to its disassembled representation.

Replacing instructions. Any behavior of a program can be implemented in multiple ways, and instructions or sequences of instructions can be replaced with syntactically different, yet semantically equivalent code. For example, on the Intel x86 platform the instructions `MOV EAX, 0` and `XOR EAX, EAX` are equivalent and can be replaced with

each other. More complex obfuscations of this type include the replacement of CALLs with a combination of PUSH and RET instructions [Lakhotia et al. 2010]. De Sutter et al. [2009] replaced infrequently used opcodes with blocks of more frequently used ones to reduce the total number of different opcodes used in the code and to normalize their frequency. Similarly, shellcode for application exploits can be transformed into innocuous-looking sequences of input characters by tools such as SCMorphism². Mason et al. [2009] demonstrated how shellcode can even be encoded into a representation that looks similar to English prose. Related concepts are used to hide network traffic by transforming an encrypted payload into a format that resembles common network protocols such as HTTP [Dyer et al. 2013].

Potentially malicious code can also be hidden in side effects of innocent looking sequences of instructions [Schrittwieser et al. 2013]. A side effect can be any effect on the state of the underlying machine that is not covered by the analysis model (e.g., the state of the flags register). Similar concepts have also been discussed in the context of shellcode obfuscation³.

Opaque predicates. A predicate (Boolean-valued function) is opaque if its outcome is known to the obfuscator at obfuscation time, but difficult to determine for a de-obfuscator [Collberg et al. 1997, 1998b]. Opaque predicates are used to make static reverse engineering more complex by introducing an analysis problem which is difficult to solve without running the program. The prime example for the use of opaque predicates is the obfuscation of a program’s control flow graph by adding conditional jumps that are dependent on the result of opaque predicates.

To prevent an analyst from identifying opaque predicates through their static behavior across a large number of program executions, refined concepts for opaque predicates were developed. Palsberg et al. [2000] introduced the concept of *dynamic opaque predicates*. It uses a set of correlated opaque predicates that all evaluate to the same result in one run, but they may all evaluate to the same different result in other runs of the program. Majumdar and Thomborson [2006] described *temporally unstable opaque predicates* for which evaluations at multiple points inside the program lead to different results.

Inserting dead code. The term “dead code” refers to code blocks which are not or simply cannot be reached in the control flow graph and thus never get executed [Collberg et al. 1997]. Inclusion of such code can make the analysis of a program more time consuming as it increases the amount of code that has to be analyzed. For making the identification of dead code more difficult, opaque predicates that always resolve to either true or false can be used.

Inserting irrelevant code. Cohen [1993] described the concept of irrelevant (“garbage”) code. Sequences of instructions that do not have an effect on the execution of a program can be inserted into the code to make analysis more complex. The most simple form of irrelevant code are NOP instructions which do not modify the program’s state. In contrast to dead code, irrelevant code can be reached by the control flow of the program and is executed at runtime, however, without any effect on the program state.

Reordering. Similarly to data structures, also expressions and statements can be reordered to decrease locality, whenever the order does not affect the program behavior. Such techniques were originally introduced for code optimization [Bacon et al. 1994] but also apply in the context of obfuscation.

²<http://www.kernelhacking.com/rodrigo/scmorphism/HowItWorks.txt> (accessed June 07, 2015)

³<http://www.securityfocus.com/archive/82/327100/2009-02-24/1> (accessed June 07, 2015)

This concept can be taken even further to move parts of the code or functionality into different modules or programs, like it was done in the Stuxnet malware [Matrosov et al. 2010].

Loop transformations. Many loop transformations have been designed to improve the performance and space usage of loops [Bacon et al. 1994]. Some of them increase the complexity of the code and can therefore be used for obfuscation purposes. *Loop tiling* was originally designed to optimize the cache behavior of code. It breaks up the iteration space of a loop and creates inner loops that fit in the cache. In *loop unrolling*, originally developed to improve performance, the body of the loop is replicated one or more times to reduce the number of loop iterations. The *loop fission* method splits a loop into two or more loops with the same iteration space and spreads the loop body over these new loops.

Function splitting/recombination. Function cloning describes the concept of splitting the control flow in two or more different paths that look different to the analyst, while they are in fact semantically equivalent. Another transformation type merges the bodies of two or more (similar) functions. The new method has a mixed parameter list of the merged functions and an extra parameter that selects the function body to be executed.

The related idea of overlapping functions—where the binary code of one function ends with bytes that also define the beginning of another function—is commonly used by compilers for optimization purposes and can also be used to confuse a disassembler. A similar but more sophisticated concept was introduced by Jacob et al. [2007]: two independent code blocks are interweaved in a way that, depending on the entry and exit points of the merged code, different functionality is executed.

Aliasing. Inserting spurious aliases (i.e., pointers to memory locations) can make code analysis more complex, as the number of possible ways for modifying a particular location in memory increases [Horwitz 1997; Ramalingam 1994]. These pointer-references can also be used as indirections to complicate the reconstruction of the control flow graph of a program in static analysis scenarios [Wang et al. 2001].

Name scrambling. Modifying identifier names such as the ones of variables and methods and replacing them with random strings is a prime example for source code obfuscation, which is not covered in our work. While binary code usually does not contain identifier names any more, byte code preserves some of the identifier names. For example, Java byte code contains class, field, and method names. By substituting expressive names with random strings, semantic information that can be important for a human analyst is removed.

Control flow obfuscation. This class of obfuscations aims at obfuscating the program control flow graph. The control flow flattening obfuscation completely obscures the links between basic blocks. Wang et al. [2000] described *chenxification*, which puts the basic blocks of a program into a large switch-statement (called dispatcher) that decides where to jump next based on an opaque variable. Control flow flattening using a central dispatcher was also described by Chow et al. [2001]. A similar concept by Linn and Debray [2003] uses a so-called *branch function* to obfuscate the targets of CALL instructions. All calls are forced to pass through the branch function, which directs the control flow to the actual target based on a call table. Popov et al. [2007] proposed to replace control transfer instructions by traps that cause signals. The signal handling code then performs the originally intended control flow transfer. Further control obfuscation techniques were described by László and Kiss [2009], Cappaert and Preneel [2010], Coppens et al. [2013], and Schrittwieser and Katzenbeisser [2011].

Parallelized code. While originally being a code optimization technique, parallelizing code also became popular in the context of code obfuscation, since parallelized code tends to be harder to understand than sequential code [Collberg et al. 1997]. Adding dummy processes to a program or parallelizing sequential code blocks that do not depend on each other [Wolfe et al. 1995] increases the complexity of analysis.

Removing library calls. Calls to libraries of programming languages (particularly ones with a high level of abstraction) offer useful information to an analyst. Because they are called by their name, they cannot be obfuscated. By replacing standard libraries with custom versions, these calls can be removed and thus their functionality obfuscated. Another variant of this obfuscation method is to link libraries statically into the application or to combine many small libraries into a few large ones.

Breaking relations. This technique aims at obfuscating relations between components of a program such as the structure of the call graph or the inheritance structure of an object-oriented program. For example, classes can be split up (*factoring*); similarly, common features of independent classes that do not have common behavior can be moved into a new parent class (*false refactoring*). Sosonkin et al. [2003] presented concepts for Java bytecode obfuscation through class coalescing, class splitting and object type hiding. A related approach by Sakabe et al. [2005] takes advantage of concepts in object-oriented languages such as polymorphism to obfuscate the relation between objects. Moreover, the idea of class hierarchy flattening to remove all inheritance relations from object-oriented programs was introduced in the literature [Foket et al. 2013, 2014].

3.3. Dynamic code rewriting

The main characteristic of code obfuscation schemes in this category is that the executed code differs from the code that is statically visible in the executable.

Packing/Encryption. Various malware obfuscation approaches analyzed in the literature follow the concept of *packing*, which hides malicious code by encoding or encrypting it as data that cannot be interpreted by static analysis. An unpacking routine turns this data back into machine-interpretable code at runtime. By changing the encryption/encoding keys, packed program code can easily be rewritten upon distribution to complicate simple pattern matching analysis (*polymorphism* [Nachenberg 1997]).

The concept of *packing* is also used for benign software. Both the reduction of storage requirements through compression and the aim for obfuscating the code of an application to deter program analysis are key motivations for the adaption of packing technologies. For these application areas, a large number of commercial packers such as VMProtect⁴, ASPack⁵, Armadillo⁶, Execryptor⁷, Enigma⁸, PECompact⁹, and Themida¹⁰ as well as open-source tools (e.g., UPX¹¹ and Yoda¹²) exist. Most of these tools are also popular with malware authors to hide the maliciousness of their code [Brosch and Morgenstern 2006].

⁴<http://vmpsoft.com> (accessed February 25, 2015)

⁵<http://www.aspack.com> (accessed February 25, 2015)

⁶<http://www.siliconrealms.com/armadillo.php> (accessed February 25, 2015)

⁷<http://www.strongbit.com/execryptor.asp> (accessed February 25, 2015)

⁸<http://enigmaprotector.com> (accessed February 25, 2015)

⁹<http://bitsum.com/pecompact> (accessed February 25, 2015)

¹⁰<http://www.oreans.com/themida.php> (accessed February 25, 2015)

¹¹<http://upx.sourceforge.net> (accessed February 25, 2015)

¹²<http://yodap.sourceforge.net> (accessed February 25, 2015)

The concept of packing for software protection was also discussed in academia. Capaert et al. [2006] introduced a modified form of packing, where code can be decrypted at fine granularity right before execution using a key derived from other code sections. Wu et al. [2010] introduced a polymorphism based concept called *mimimorphism* which encodes data into a representation that looks like program code. A taxonomy of packer based obfuscation schemes as well as similar techniques was presented by Mavrogiannopoulos et al. [2011]. Recently, Roundy and Miller [2013] presented a survey on obfuscation techniques used in malware packers.

Dynamic code modification. In this technique similar functions are obfuscated by providing a general template in memory that is patched right before its execution [Collberg et al. 1997]. Static analysis techniques fail to analyze the program, as its functionality is available at runtime only. Other concepts of dynamic code modification [Kanzaki et al. 2003; Madou et al. 2006a] implement the idea of correcting intentionally erroneous code at runtime right before execution.

Environmental requirements. Riordan and Schneier [1998] proposed the concept of environmental key generation, in which a cryptographic key is not statically stored in a binary but constructed from *environmental* data collected from within the computing environment. Only if a specific environmental condition is met (called activation environment), the program is able to generate the key and execute its code. Outside the activation environment the program does not reveal its secrets to an analyst.

Similar concepts can be applied to code as well. Sharif et al. [2008] proposed a malware obfuscation scheme that makes the code conditionally dependent on an external trigger value. Without knowledge of this specific value, the triggered behavior is concealed from dynamic analysis. Similar techniques are widely used in malware.

Hardware-assisted code obfuscation. Hardware tokens can be used to improve the strength of other code obfuscation techniques [Fu et al. 2007; Zhuang et al. 2004; Bitansky et al. 2011]. The basic idea is to create a hardware-software binding by making the execution of the software dependent on some hardware token. Without this token, analysis of the software will fail, because important information (e.g., targets of indirect jumps) is not available.

Hardware-based isolation mechanisms for trusted computation such as Intel SGX (Software Guard Extensions) allow an application to prevent other applications and even the operating system kernel from accessing certain memory regions [Anati et al. 2013]. Such mechanisms seem well-suited to protect code and data from runtime inspection and tampering and may in the future become an impediment to dynamic analysis.

Virtualization. Virtualization describes the concept of converting the program's functionality into byte code for a custom virtual machine (VM) interpreter that is bundled with the program [King and Chen 2006; Ghosh et al. 2010]. Virtualization can also be combined with *polymorphism* by implementing custom virtual machine interpreters and payloads for each instance of the program [Anckaert et al. 2006]. Vrba et al. [2010] proposed the combination of fine-granular encryption and virtualization to hide VM code from analysis. Collberg et al. [1997] described a variant of this concept under the term *table interpretation*. A similar concept by Monden et al. [2004] uses a finite state machine-based interpreter to dynamically map between instructions and their semantics.

Anti-debugging and -disassembly techniques. This obfuscation category includes techniques that actively oppose analysis attempts via disassembly or debugging. For example, attached debuggers can be detected based on timing and latency analysis or

the identification of code modifications caused by software breakpoints. Another technique is the execution of undocumented instructions in order to confuse a code analysis tool or a human analyst [Brand 2010].

4. CODE ANALYSIS

We consider four broad classes of analysis against programs protected by obfuscation: pattern matching, automated static analysis, automated dynamic analysis, and manual reverse engineering by a human who has access to automated tools.

Note that there is no universally accepted definition for static or dynamic analysis; in fact, we believe that it is impossible to draw strict boundaries between any of the analysis classes, so we apply our best judgment to arrive at a meaningful classification of existing approaches to analyzing programs. Madou et al. [2005] argue that considering resilience against only particular types of analysis can leave a technique open to circumvention by hybrid approaches.

4.1. Pattern Matching

With the term “pattern matching” we refer to fast automated techniques for finding and unveiling known structures in binary programs. For instance, this could be artifacts introduced by obfuscation tools or signatures of well known libraries and code samples. In contrast to our definition of static and dynamic analysis, pattern matching is purely syntactic and does not reason about the program semantics. The types of patterns to look for can range from simple byte strings to regular expressions or other languages allowing wildcards. We also consider machine learning techniques that essentially treat programs as byte strings to be pattern matching.

One example for pattern matching is the Fast Library Identification and Recognition Technology (F.L.I.R.T.) in IDA¹³. It uses signatures to identify library functions with the goal of simplifying reverse engineering. Function signatures of well known libraries are stored in a database; when disassembling a new binary, its bytes are checked against known signatures. Recognized functions are named and annotated using the stored information to aid a human reverse engineer in understanding the functionality of the code.

Usually, the patterns to look for have been generated ahead of time or are curated by human analysts. Differential pattern matching approaches work without a priori knowledge and instead compare static or dynamic artifacts, such as control flow graphs [Flake 2004; Nagarajan et al. 2007] or instruction traces [Zhang and Gupta 2005]. As mentioned when discussing our methodology in Section 2.1, we consider this type of pattern matching a multi-stage analysis with an initial static or dynamic phase. Note that a simple translation of byte patterns to instruction opcodes in the manner of a linear sweep disassembler would still fall under pure pattern matching; we draw the boundary to static analysis where a disassembler interprets possible jump targets.

4.2. Static Analysis

Static analysis is widely used for optimizing code, finding or proving the absence of bugs, and reverse engineering. In its broadest sense, static analysis refers to any program analysis that is performed just by inspecting executable code or a disassembled representation of a program of interest but without ever executing it on a real or virtual machine. In this survey, we distinguish static analysis from syntactic pattern matching as describing analyses that reason about the program semantics. For a detailed introduction to static program analysis, we refer the reader to the classic textbook by Nielson et al. [1999].

¹³https://www.hex-rays.com/products/ida/tech/flirt/in_depth.shtml (accessed February 25, 2015)

Static analysis deduces information about a program by reasoning about the possible executions of the target program. A static analysis is considered to be *sound* if it is guaranteed to include *all* possible executions of a program in its judgment. For example, a static analysis that attempts to generate a control flow graph from a binary program is sound if the resulting graph contains at least all those control transfers that will ever occur at runtime. Because of undecidability, static analysis can only achieve soundness by *over-approximating*, i.e., by generalizing the concrete program behavior and accepting that this generalization will also include behavior that does not occur in real executions. To continue the example, a sound control flow graph could contain edges that will never be taken at runtime. Over-approximation is the reason why static analysis can report false positives when it is used for bug finding or malware detection.

In principle, the precision of a static analysis can be increased to reduce the number of false positives, but such added precision comes with a corresponding increase in cost. As a compromise, many practical static analysis systems choose to be *unsound*, i.e., they do not guarantee to cover all possible behaviors. This may introduce false negatives in addition to false positives, but is often an acceptable trade-off to make systems useful in practice.

An example of the trade-offs in precision, performance, and soundness can be seen in the various approaches to reconstruct the control flow of binaries: a linear sweep disassembler that disassembles byte after byte starting at the entry point performs a low-precision static analysis that is unsound because it can miss code reached through branches [Schwarz et al. 2002] (in fact, since its interpretation of the instruction semantics is trivial, we classify linear sweep as pattern matching, see above). A recursive traversal disassembler improves accuracy by trying to identify and follow the location of branches during analysis, but it will still leave out many indirect branches whose targets are computed at runtime. Disassembly by abstract interpretation performs a sound static analysis of the binary code to also systematically resolve indirect branches whose targets are computed at runtime [Kinder et al. 2009]. The computational cost depends on the exact type of analysis performed but can be significant.

Analysis platforms such as CODESURFER/x86 [Balakrishnan and Reps 2004], JAKSTAB [Kinder and Veith 2008], or BAP [Brumley et al. 2011] provide the necessary infrastructure and abstractions to apply higher level static analysis to binaries. Typically, low-level details of instruction encoding and semantics are abstracted by such platforms, and custom analyses can work on a simplified low-level language instead of machine code.

From a conceptual standpoint, obfuscation decreases the precision a particular static analysis can achieve, i.e., obfuscation introduces additional sources of over-approximation [Giacobazzi and Mastroeni 2012]. Recall from Section 3.2 that control flow flattening forces all control flow to pass through a dispatcher. Because the dispatcher is executed many times on all paths, a static analyzer that generalizes the program behavior at the location of the dispatcher will be very imprecise. To maintain precision, a static analyzer can compensate by generalizing the behavior at a finer granularity [Kinder 2012].

In general, whenever the details of a particular obfuscation technique are known, an analysis can be crafted to almost completely eliminate the effect of the obfuscation [Krügel et al. 2004]. However, despite existing work on automated refinement of the precision of general static analysis [Das et al. 2002; Bardin et al. 2011], tailoring analyzers to cope with obfuscation schemes is still a largely manual process.

Besides the conceptual precision requirements, there are also significant practical challenges for statically analyzing obfuscated code. All existing tools make some assumptions about the behavior or structure of the executable code, which can be broken by obfuscations. Many static analyzers fail even in the presence of simple obfuscations

used regularly by malware [Moser et al. 2007b]. Especially tools that depend on an initial disassembly phase, such as CODESURFER/X86 [Balakrishnan and Reps 2004] or early semantic malware detectors [Christodorescu et al. 2005; Kinder et al. 2005], are vulnerable to syntactic obfuscations. Such obfuscations target disassemblers that rely mostly on heuristics to discover all executable code [Linn and Debray 2003]. Removing the separate disassembly phase and working on code directly improves resilience against these simpler obfuscation schemes [Kinder et al. 2009; Brumley et al. 2011]. Nevertheless, especially dynamic obfuscations are difficult to handle for static analyzers. For example, no static tool to date is able to reliably deal with self-modifying code, even though one could, in theory, imagine the static analysis compensating by analyzing the possible runtime state of the code section.

However, static analysis of obfuscated code can still give a reverse engineer valuable information and lead to a better understanding of the overall binary and further guide additional de-obfuscation steps. Our analysis in Section 5 considers static analysis tools at the current state of the art, including minor adaptations for each particular scenario.

4.3. Dynamic Analysis

Dynamic analysis refers to observing the behavior of deployed and running systems, and it is today an important part of forensic analysis of malware [Egele et al. 2012]. It analyzes real executions of a program, either online (at runtime) or offline (over a recorded trace). In particular, any form of software testing is a dynamic analysis. Just like static analysis, dynamic analysis is also performed with respect to particular properties of interest. For example, a dynamic analysis may record system call invocations or executed instructions, or trace the flow of “tainted” data that is received from the network.

Dynamic analysis is conceptually dual to static analysis: a sound dynamic analysis considers a subset of all execution traces of a program and is therefore *under-approximate*. Each bug, warning, or suspicious behavior found is then guaranteed to also occur during at least one real execution. On the flip side, a dynamic analysis may miss behavior when the number of possible traces in a program is too large to be exhaustively tested. In the general case this is unavoidable due to undecidability.

While static analysis trades off precision against cost, dynamic analysis trades off *coverage* against cost. Each execution of the program under test corresponds to an additional trace covered. An automated dynamic analysis will typically try to cover traces that exhibit diverse behaviors to get a representative profile of the program. If a program exhibits a particular type of behavior only under very specific circumstances, it may never be observed. This is especially problematic for malware, which may respond only to certain “triggers” [Crandall et al. 2006; Kolbitsch et al. 2011]. This challenge is closely related to the problem of automated test case generation, and approaches can be roughly classified into blackbox, graybox, and whitebox depending on how they treat the program under test.

Blackbox testing assumes no knowledge of the program and is usually done via random input generation. Since it is very hard to achieve meaningful coverage through blind enumeration of inputs, effective fuzz testing tools employ a *graybox* approach and use domain-specific knowledge about the input format to generate inputs that have different meaning (e.g., image files with different sizes). Coverage-guided fuzzers such as AFL¹⁴ mutate an initial input following a variety of heuristics. *Whitebox* testing also uses the program structure to guide input generation. Tools based on symbolic execution [King 1976] and dynamic test generation (also referred to as concolic execu-

¹⁴<http://lcamtuf.coredump.cx/afl/> (accessed June 07, 2015)

tion in the literature) have been particularly successful [Godefroid et al. 2005; Cadar et al. 2006; Sen et al. 2005]. They use a constraint solver to generate inputs systematically, such that one input will cover exactly one unique control flow path through the program. However, since the number of control flow paths can be infinite (due to loops), even symbolic execution will remain incomplete and thus under-approximate in a finite amount of time. A survey by Schwartz et al. [2010] provides a good introduction to symbolic execution and explains the related challenges in more detail.

In practice, systematic exploration methods such as symbolic execution are additionally limited by their underlying constraint solver which is used to identify valid control flow paths. A new control flow path that exhibits hitherto unseen behavior can only be triggered if the constraint solver is able to find an input that drives execution down that path. If the constraint is unsupported (e.g., floating point) or simply too difficult, the symbolic execution engine is unable to cover the path. As a best effort solution, the engine can then resort to random testing (fuzzing) of input parameters [Godefroid et al. 2008]. As with static analysis, there are also dynamic techniques that compromise on soundness. For instance, it is possible to directly manipulate registers or memory to trigger different branches or execute particular areas of code [Madou et al. 2005; Moser et al. 2007a]. While this may put the program into a state that would never occur under normal circumstances, it will cause a particular path to be executed without requiring any constraint solving, which may just be enough for the particular reverse engineering task at hand.

A significant practical advantage of dynamic over static approaches in the analysis of obfuscated code is that it can be applied to binaries with relative ease. In fact, it is often simpler to dynamically analyze binaries than source code, because traces recorded at runtime show addresses of instructions in the binary and not just (obfuscated) source code information. However, the use of anti-debugging techniques can oppose dynamic analysis attempts.

Because dynamic analysis monitors what is actually executed at runtime, obfuscations cannot fully conceal the behavior of a program. For instance, a dynamic analysis can trace self-modifying code just like regular code if it records the opcode and operands of the current instruction in addition to the value of the program counter [Thakur et al. 2010].

As for static analysis, Section 5 lists what is possible using the current state of dynamic analysis tools, with only relatively minor technical adaptations.

4.4. Human Assisted Reverse Engineering

We use this very broad category to cover any kind of analysis that a skilled human reverse engineer can perform with the help of any state-of-the-art tool. The ability for creative problem solving and adaptation makes humans much more efficient in dealing with obfuscations than fully automated techniques. In contrast to a purely automated approach, however, humans can be misled by clues that suggest structure, such as type names or inheritance relationships among classes.

Disassemblers and decompilers alike went through major improvements over the past two decades [Cifuentes and Gough 1995; Schwarz et al. 2002; Schwartz et al. 2013]. The de facto industry standard for disassembling and reverse engineering, IDA Pro, now includes a powerful decompiler for the x86, x64, and ARM processor architectures¹⁵. Academic research has made significant progress as well. The Boomerang Project¹⁶ is an attempt to develop an open source decompiler. Early results from a study [Emmerik and Waddington 2004] conducted on a real-world program were

¹⁵<https://www.hex-rays.com/products/decompiler> (accessed February 25, 2015)

¹⁶<http://boomerang.sourceforge.net> (accessed February 25, 2015)

promising. Despite the fact that only the core algorithm and not the entire program was decompiled, the experiment was able to demonstrate that decompilation with human assistance can be practical in certain use cases. The project, however, does not seem to receive much attention anymore. Recently, a new binary-to-C decompiler named PHOENIX was introduced by Schwartz et al. [2013]. It implements a structural analysis algorithm that uses iterative refinement strategies as well as the property of semantics-preservation in order to archive significantly more accurate results than previous approaches. The reconstruction of high-level control flow structures is a challenging problem in decompiler research. State-of-the-art decompilers rely on structural analysis and advanced pattern matching and fall back to using `goto`-statements for control flow transfers if no patterns can be identified. Recently Yakdan et al. [2015] proposed a decompiler that does not rely on pattern matching and produces more compact and `goto`-free code. It uses semantics-preserving transformations to restore structured control flow graphs.

A major driving force behind the development of human assisted code analysis approaches is the software cracking scene. In early years the SoftICE debugger was a popular tool among reverse engineers, however it is no longer maintained [Willems and Freiling 2012]. Today, the freely available debugger OllyDBG¹⁷ to some extent continues where SoftICE left off. A large number of plugins that are dedicated to cracking and tampering purposes have been made available by the scene (e.g., Ollybone [Stewart 2006] for semi-automatic unpacking).

The concepts and tools in this code analysis category have in common that the automated analysis process can be directly influenced by the human reverse engineer resulting in very different outcomes depending on the human's decisions. For example, the disassembling process of IDA Pro can be influenced by a human reverse engineer to perform recursive traversal at specific locations of the program's code. This code analysis category is naturally hard to grasp formally. In our analysis in the next section we base the capabilities of human reverse engineers on published results and common knowledge about the state of the art. A comprehensive introduction to software reverse engineering was published by Eilam [2005].

5. ROBUSTNESS ANALYSIS

In this section we evaluate the effectiveness of different classes of software obfuscation schemes against the program analysis scenarios introduced in Section 2.4. The effectiveness of a specific type of code obfuscation is evaluated by comparing it to code analysis approaches described in the literature. We are well aware of the arms race between code obfuscation and analysis and the fact that in theory every obfuscation technique can be broken with targeted analysis techniques (see Section 4). In this survey, we focus on the status quo of code obfuscation in real-life application scenarios and evaluate the capabilities of state-of-the-art code analysis tools (also considering possible non-complex modifications to the analysis techniques) in order to target particular obfuscations. The possibility of developing analysis techniques targeted to break a specific obfuscation scheme does not prove it useless in general, as a small modification to the obfuscation technique can again raise the bar for analysis.

As of today's knowledge, a precise formalization of the security of an obfuscation scheme seems to be difficult to achieve. Previous attempts to quantify the hardness of a particular class of code obfuscation are based on software complexity metrics [Collberg et al. 1997; Anckaert et al. 2007]. However, it remains unclear whether such notions are able to capture all security properties of the obfuscating transformation correctly. In this survey, we thus follow a different approach and rate the strength of each obfus-

¹⁷<http://www.ollydbg.de> (accessed February 25, 2015)

cation class. A summary of the results can be found in Table II. **Black** marks obfuscation techniques that break a certain type of program analysis fundamentally with respect to its state-of-the-art techniques. **Gray** denotes scenarios in which a particular code obfuscation class cannot be considered unbreakable, but still makes analysis substantially more expensive. For example, while an analysis technique might work under lab conditions when dealing with toy programs, limits of available resources could be reached while analyzing larger programs. **White** marks obfuscation techniques that only slightly raise the cost of analysis and therefore cannot justify the cost of adopting them, such as increased code size or added runtime overhead. Furthermore, we distinguish between ratings that are supported by results in the literature (marked with a checkmark in Table II) and ones that we based on our own evaluation of the state of the art.

5.1. Pattern Matching

The most basic form of pattern matching is limited to comparing program code on byte level against some predefined pattern. Consequently, this type of pattern matching is weak against all kinds of code modifications, including obfuscating transformations. However, several more sophisticated pattern matching concepts, mostly for identifying malicious behavior, were introduced in the literature. Such advanced pattern matching arguably can cope with naive obfuscation techniques such as equivalent instructions. In the following, we discuss the state of the art in research on pattern matching-based code analysis in the presence of different classes of code obfuscation schemes.

Locating data. Naive pattern matching techniques on data break as soon as the structure of the data changes. Thus, even simple obfuscations are effective against simple forms of pattern matching. This was confirmed by Moser et al. [2007b], who were able to show that simple data obfuscation techniques are sufficient to make pattern matching ineffective for the identification of data. *Anti-disassembly techniques* are not effective against pattern matching-based analysis of the binary, since they do not necessarily require the program to be disassembled.

Locating code. Moser et al. [2007b] also demonstrated the limitations of pattern matching against the static code rewriting technique *control flow flattening*. Their reasoning can be easily extended to other forms of static rewriting and dynamic code rewriting as any modification of the binary clearly destroys the pattern.

In the literature, locating code through pattern matching is often described in a malware context. The primary aim of techniques in this analysis scenario is the generation of generic patterns describing malicious behavior of program fragments to automatically classify software as malicious or benign. While these approaches do not aim at understanding the semantics of the program, they still can identify locations in the code that implement abnormal (malicious) behavior. HANCOCK [Griffin et al. 2009] is a system for the automated generation of signatures for malware. Through normalization of opcodes it is resistant against simple code transformations such as register re-assignment. Dalla Preda et al. [2011] generated signatures of metamorphic malware through abstract interpretation of semantics developed for self-modifying code. Generally, the expressiveness of the language used for creating patterns (e.g., static strings vs. regular expressions) dictates both their resilience against obfuscations and the performance of matching algorithms. Since existing techniques emphasize performance over expressiveness, we marked static code rewriting techniques in gray.

Dynamic code obfuscation techniques, for instance virtualization, remove the structure of the code entirely, thus rendering pattern matching based analysis approaches ineffective. To some extent, however, code analysis based on patterns is still possible as demonstrated in the literature. Tang and Chen [2007] proposed the identification

Table II. Analysis of the strength of code obfuscation classes in different analysis scenarios (PM = Pattern Matching, LD = Locating Data, LC = Locating Code, EC = Extracting Code, UC = Understanding Code).

Name	PM		Autom. Static				Autom. Dynamic				Human Assisted				
	LD	LC	LD	LC	EC	UC	LD	LC	EC	UC	LD	LC	EC	UC	
Data obfuscation															
Reordering data	■	■	✓	■	■	■	■	■	■	■	■	■	■	■	■
Changing encodings	✓	■	■	■	■	■	■	■	■	■	■	■	■	■	■
Converting static data to procedures	✓	■	■	■	■	■	■	■	■	■	■	■	■	■	■
Static code rewriting															
Replacing instructions	■	✓	■	✓	■	■	■	■	■	■	■	■	■	■	■
Opaque predicates	■	■	■	✓	■	■	■	■	■	■	■	■	■	■	■
Inserting dead code	■	■	■	■	■	■	■	■	■	✓	✓	■	■	■	■
Inserting irrelevant code	■	✓	■	■	■	■	■	■	■	■	■	■	■	■	■
Reordering	■	■	■	■	■	■	■	■	■	■	■	■	■	■	■
Loop transformations	■	■	■	■	■	■	■	■	■	■	■	■	■	■	■
Function splitting/recombination	■	■	■	■	■	■	■	■	■	■	■	■	■	■	■
Aliasing	■	■	■	■	■	■	■	■	■	■	■	■	■	■	■
Control flow obfuscation	✓	■	■	■	■	■	■	■	■	■	■	■	■	■	■
Parallelized code	■	■	■	■	■	■	■	■	■	■	■	■	■	■	■
Name scrambling	■	■	■	■	■	■	■	■	■	■	■	■	■	■	■
Removing standard library calls	■	■	■	■	■	■	■	■	■	■	■	■	■	■	■
Breaking relations	■	■	■	■	■	■	■	■	■	■	■	■	■	■	■
Dynamic code rewriting															
Packing/Encryption	■	✓	■	■	■	■	■	■	■	■	■	■	■	■	■
Dynamic code modifications	■	■	■	■	■	■	■	■	■	■	■	■	■	■	■
Environmental requirements	■	■	■	■	■	■	■	■	■	■	■	■	■	■	■
Hardware-assisted code obfuscation	■	■	■	■	■	■	■	■	■	■	■	■	■	■	■
Virtualization	■	■	✓	✓	■	■	■	■	■	■	■	■	■	■	■
Anti-debugging techniques	■	■	■	■	■	■	■	■	■	■	■	■	■	■	■
Legend	■	obfuscation breaks analysis fundamentally													
	■	obfuscation is not unbreakable, but makes analysis more expensive													
	■	obfuscation only results in minor increases of costs for analysis													
	✓	A checkmark indicates that the rating is supported by results in the literature Scenarios without a checkmark were classified based on theoretical evaluation													

of polymorphic malware on a network stream using statistical analysis. In contrast to fixed string pattern matching, polymorphic versions of a program can be identified with this approach. POLYGRAPH [Newsome et al. 2005] is a concept for the automated identification of polymorphic worms which exploits the existence of invariant substrings in all polymorphic variants of a malware. Both approaches exploit the characteristic structure of polymorphic programs. To summarize, while it was shown in literature that pure identification of the obfuscation method *polymorphism* is feasible in a malware context, locating particular functionality inside the binary is still impos-

sible with pattern matching if the implementation details are unknown to the analyst. As in the previous scenario, *anti-disassembly techniques* do not provide additional protection against pattern matching approaches aiming at locating code.

5.2. Static code analysis

Locating data. Similar to pattern matching, automated static analysis of data is limited when analyzed data is not stored in its original representation. In Table II the data obfuscation class *reordering* is marked in gray, because a simple data flow analysis, which is necessary for the reconstruction of reordered data, is arguably possible with automated static analysis tools such as JAKSTAB [Kinder and Veith 2008]. Furthermore, it can be argued that a static code rewriting technique can be effective against locating data if it complicates the reconstruction of the control flow graph of the program: static data flow analysis strongly depends on knowledge of the control flow. For this reason, control flow modifying obfuscation techniques are marked in gray in Table II.

Dynamic code rewriting in general is effective against locating data inside programs as the original representation of the data is destroyed. At first sight, static analysis appears completely helpless against virtualization obfuscation, as only the code of the interpreter can be directly analyzed. In static analysis, this leads to an effect that Kinder [2012] called *domain flattening*: data flow information from different locations in the original program are merged to one location in the interpreter, resulting in a more imprecise analysis. However, Kinder [2012] was able to demonstrate that by lifting the static analysis to a second dimension of location (the virtual program counter), the same analysis precision as on non-obfuscated code is achievable. While the introduced approach was evaluated on a toy example only, the preliminary results still indicate promising directions for static analysis of VM-protected binaries. For that reason, we marked *virtualization* in gray for this analysis scenario. *Anti-disassembly techniques* can, to some extent, limit static analysis and thus are marked in gray as well.

Locating code. Locating a particular program feature in binary code through static analysis can be based on an analysis of the structure of the control flow graph (e.g., Harris and Miller [2005]) of the program. Therefore, code obfuscation schemes that modify or hide the control flow of a program (*opaque predicates*, *loop transformations*, *parallelized code*, etc.) can be considered as candidates for protection against static analysis techniques. *Opaque predicates* are the most important concept for control flow obfuscation against static analysis tools. Dalla Preda et al. [2006] proposed an abstract interpretation-based methodology for removing simple *opaque predicates*. This automated, static concept was shown to be more complete than approaches for dynamic analysis of opaque predicates. However, as the authors state, their analysis concept is limited to simple types of opaque predicates only. Thus, we consider more complex opaque predicates still effective in making static reconstruction of the control flow graph significantly more difficult.

Tools for matching program code based on control flow similarities include BIN-DIFF [Flake 2004], BINHUNT [Gao et al. 2008], BINS LAYER [Bourquin et al. 2013], and the implementation by Nagarajan et al. [2007]. Moreover, Tsai et al. [2009] introduced a framework for analyzing control flow obfuscation by representing it as a composition of atomic operators to evaluate robustness. Still, it remains unclear to what extent such theoretical results can support the evaluation of the strength of an obfuscation in real-life applications. We conclude that while no general statement regarding the strength of obfuscation can be made for this particular analysis scenario, obfuscation schemes that complicate the reconstruction of the control flow graph can still make the

analyst's aim of identifying the entry point into a particular functionality considerably harder.

Code obfuscation based on dynamic code rewriting makes static analysis considerably more difficult as the analyzed code in the binary does not correspond to the code that is actually executed. However, approaches to circumvent obfuscation were introduced in the literature on malware samples. *Malware packers* were analyzed with heuristics-based static analysis techniques (e.g., Treadwell and Zhou [2009]) and comparison with previously-seen malware samples was performed [Jacob et al. 2012; Karnik et al. 2007]. Similar to pattern matching based approaches, in automated static analysis, the maliciousness of code is evaluated by identification of abnormal structures of the program code. The idea of using model checking for detecting malicious code was proposed by Kinder et al. [2005]. Furthermore, static analysis of *polymorphism* as well as *metamorphism* was discussed in recent literature. Bruschi et al. [2006a] compared a normalized version of the control flow graph of a binary against control flow graphs of known malware in order to detect malicious behavior. In a similar concept, Walenstein et al. [2006] normalized program code (e.g., through simple instruction substitution patterns) to be able to compare it to known malware samples. Another code normalization technique was presented by Bruschi et al. [2006b]. Coogan et al. [2009] used a combination of two static code analysis techniques (slicing and alias analysis) to identify malware packers. The idea of using the frequency of opcodes as a predictor for malware was proposed by Bilar [2007]. For example, massive use of mathematical operations might indicate malware that tries to obfuscate its malicious behavior using a packer based approach. Furthermore, it was shown that a high frequency of a usually rare opcode is an indicator for polymorphic or metamorphic malware. Chouchane and Lakhotia [2006] proposed the detection of metamorphic malware by creating signatures for instruction-substitution engines. However, similar to pattern matching based approaches, none of the introduced concepts for static analysis is able to find the location of specific functionality. Only the characteristic structure of a malware packer can be identified. As a consequence, we marked *packing/encryption* and *replacing instructions* in gray. Like in the previous scenario (locating data), *virtualization obfuscation* was marked gray in Table II due to preliminary results on the challenges of statically analyzing virtualization-obfuscated programs by Kinder [2012].

Extracting code. While the software engineering literature knows of concepts for reusing functionality from legacy binary code (e.g., [Sneed 2000]), the automated static extraction of obfuscated code has not been widely discussed.

Our reasoning regarding the effectiveness of different classes of code obfuscation techniques in this analysis scenario is simple: extracting code from a program through static analysis is at least as difficult as locating code, because the latter is a fundamental requirement for the former. Table II indicates differences between locating and extracting code in the scenario of static code rewriting for four obfuscation techniques. *Aliasing*, *control flow flattening*, *parallelized code*, and *breaking relations* share the common effect of increasing code dependencies by interweaving independent parts of the program. These interweavings are difficult to resolve through static analysis, thus making the extraction of code sections considerably harder than just locating code. Collberg et al. [1997] first described the effect of raising analysis complexity when extending the scope of the obfuscating transformation. Analogously, in the category of dynamic code rewriting, *virtualization* was marked in black because of the interweavings that make it difficult to extract all required code sections. *Anti-debugging* and *Anti-disassembly* was marked in black due to the existence of a wide range of techniques that actively interfere with static disassembly [Branco et al. 2012]. This

includes concepts to prevent a disassembler from recovering a correct higher level representation of the code [Popov et al. 2007].

Understanding Code. For a human analyst to better understand an obfuscated program, an automated analysis tool has to be able to remove at least parts of the applied code obfuscation scheme from the binary.

Several concepts for static de-obfuscation with the aim for making the program code more understandable for a human analyst were proposed in the literature. An early work by Rugaber et al. [1995] has shown that detecting interleaved code (e.g., through *function recombination*) is a time consuming task. Majumdar et al. [2006] evaluated the robustness of the obfuscation technique *aliasing*, where two or more pointers refer to the same memory location. An experimental evaluation showed that resilience expected from the theoretical approach does not hold in real-life scenarios; still, in general it is difficult to evaluate the actual strength of aliasing. Guillot and Gazet [2010] developed techniques for automated static de-obfuscation. The basic concept is to make program code more understandable by automated rewriting based on local semantic analysis, similar to optimization steps made by compilers. Raber and Laspe [2007] introduced a plugin for IDA Pro that is capable of removing basic obfuscation and anti-debugging techniques from a binary. A shared limitation of all introduced concepts is the major gap in success rates between academic examples and real-world scenarios. While these concepts show that theoretical analysis models work under laboratory conditions, practical application is limited. Thus, we marked data obfuscation as well as static code rewriting techniques in gray.

In Table II *name scrambling* was marked in black. Identifier names are often critical to human understanding of a program but cannot be fully restored with the help of automated code analysis techniques. Collberg et al. [1997] described this transformation as one-way, although limited recovery is possible via code analysis (e.g., induction variables of loops can be identified and renamed to “i”, dynamic data structures can be renamed to common identifiers, etc.). *Parallelized code* can implicitly be considered as a strong obfuscation technique in this analysis scenario as code extraction, which is a fundamental requirement for code understanding, is also difficult. Following the same line of reasoning, dynamic code rewriting methods in general are effective against code understanding through automated static analysis.

5.3. Dynamic code analysis

Locating data. Every known data obfuscation technique except for white-box cryptography has a limited practical effect in dynamic analysis scenarios since data is always visible at some point during runtime. For example, although the cryptographic key of a DRM client might be stored in an obfuscated way (e.g., through data encoding) in the binary, during runtime the cryptographic algorithm has to reconstruct the original representation of the key in order to perform decryption tasks. In the literature, several concepts for an automated dynamic extraction of data structures from program binaries were introduced. Shamir and Van Someren [1999] proposed the identification of cryptographic keys in binary code through entropy analysis — a concept which can also be applied to memory.

Zhao et al. [2011] introduced a concept for dynamic extraction of data in malware. Cozzie et al. [2008] extracted data structures from memory dumps using Bayesian unsupervised learning. Lin et al. [2010] introduced a system called REWARDS which reveals data structures through observation of the program execution. It marks each memory location that was accessed at runtime with a timestamp and traces the propagation of data. A similar concept called HOWARD was proposed by Slowinska et al. [2011] in 2011. It aims at extracting data structures from binaries by dynamically trac-

ing (using QEMU-based emulation) how a program accesses the memory. HOWARD is able to reconstruct large parts of the symbol table. It thus simplifies the progress of reverse engineering and improves readability of obfuscated code as well as data.

To sum up, in a dynamic analysis context, most static as well as dynamic code rewriting techniques do not provide significant additional security in the context of data protection. The only exception are obfuscations that require special runtime enablers (additional hardware or environmental conditions) to execute. These techniques can withstand dynamic analysis in situations where the runtime enabler is not present. For this reason, both the techniques *environmental requirements* and *hardware-assisted code obfuscation* were marked in gray in Table II.

Locating code. Locating a particular feature inside binary code through dynamic analysis is based on the observation of the program behavior. Static code rewriting techniques are not effective in dynamic analysis scenarios due to one important factor: most of them are not explicitly targeted at the prevention of dynamic analysis. *Replacing instructions, inserting dead code, opaque predicates, code insertions*, etc. were in the first place developed to obfuscate the static representation of binary code. However, automated dynamic analysis techniques do not depend on the static representation as much as static analysis or even human analysis do; thus they are less affected by the obfuscation. In the literature very diverse concepts for dynamic analysis of obfuscated code were proposed. Li et al. [2009] described a technique that identifies malicious behavior based on the malware's runtime system call sequences. MCVETO [Thakur et al. 2010] is a dynamic test generator and model checker for machine code. While traditional dynamic analysis suffers from the problem of incompleteness as a program behavior can only be analyzed on one input at a time, MCVETO implements a combined static and dynamic approach which aims at reaching more code locations by actively manipulating program inputs. Another strategy for finding a particular feature in program code was discussed by Deprez and Lakhota [2000] and Wilde and Scully [2006]. The basic idea is to execute the program twice with two different inputs whereby one input invokes the feature and the other does not. From calculating the differences between the two traces conclusions on the location of the particular feature can be drawn. Madou et al. [2005] discussed the effectiveness of hybrid (static and dynamic) analysis approaches and demonstrated it in the context of the reconstruction of an obfuscated control flow graph. Zhang and Gupta [2005] proposed the comparison of instrumented executions of different program versions for matching purposes. Brumley et al. [2008] introduced a concept for automated generation of exploits based on static and dynamic comparison of programs and patched versions of it.

While dynamic code analysis is strong against static code obfuscation, dynamic obfuscation techniques are much more robust because executed code differs from the code that is statically visible in the executable. However, analyses targeting dynamic code rewriting obfuscations were presented in the context of malware analysis, in particular for packed programs. Moser et al. [2007a] proposed a solution for the incompleteness problem of dynamic analysis by making the exploration of multiple execution paths possible. Thus, the approach allows the identification of malicious behavior that is executed only if a certain condition is met. REANIMATOR [Comparetti et al. 2010] is a two-step malware identification system. First, known malware is dynamically analyzed and code that is responsible for a particular malicious behavior is modeled. This model can then be used in a second step to identify the same malicious behavior in other code samples through static analysis. However, dynamic code rewriting techniques such as packing can limit the detection rate of this approach significantly. Sharif et al. [2008] introduced the EUREKA framework that automatically extracts the payload of a packed program by running the binary in a virtual machine. RENOVO

by Kang et al. [2007] is another dynamic unpacker that monitors executed instructions and memory writes at runtime to extract a hidden payload and is based on the dynamic code analysis component TEMU [Yin and Song 2010] of the BITBLAZE platform [Song et al. 2008]. Heuristics- and statistics-based strategies are used to determine the exact moment when the unpacking process is finished and the unpacked code is fully stored in memory. Royal et al. [2006] introduced a behavioral based approach for automated unpacking inside a VM. Its combination of static and dynamic analysis identifies unpacking routines through its characteristic behavior. A malware analysis approach by Debray and Patel [2010] focuses on the automated identification of unpacking routines inside binaries. Gröbert et al. [2011] proposed the detection of cryptographic algorithms by analyzing program execution traces, which show unique characteristics depending on the implemented algorithm. In a similar concept named ALIGOT [Calvet et al. 2012], the identification of cryptographic algorithms in execution traces is based on the comparison of input-output relationships with known cryptographic algorithms. With this concept even heavily obfuscated algorithms can be identified because the input-output relationship does not differ from the original version of the algorithm.

Furthermore, several approaches for automated dynamic analysis of programs protected by virtualization were introduced in recent years. Sharif et al. [2009] proposed the use of taint- and data-flow analysis techniques to find the byte code implementing the payload of the virtualized program. The described automated reverse engineering approach is able to reconstruct control flow graphs and was evaluated against the code virtualization tools VMProtect and Code Virtualizer. A different strategy for automated dynamic analysis of virtualized code was proposed by Coogan et al. [2011]. Instructions that contribute to arguments of system calls are collected to understand the functionality of the program. Webster and Malcolm [2009] proposed the use of formal algebraic specifications to detect metamorphic and virtualization-based malware. TRACESURFER [Guizani et al. 2009] uses dynamic binary instrumentation for the detection of self-modifying malware.

To conclude, we marked dynamic code rewriting approaches in gray because practical application of all described approaches is limited to malware identification tools only. In other words, they do not directly aim at locating particular functionality, but malicious behavior in general.

Extracting code. Similar to the scenario of static extraction of code sections, dynamic code extractors have to deal with dependencies between different parts of the program. Several static classes of code obfuscation add bogus dependencies in order to make analysis more difficult. Thus, we can make similar assumptions on the resilience of the code obfuscation techniques. Still, several dynamic concepts for automated extraction of code section were described in the literature. TOP by Zeng et al. [2013] collects instruction traces and translates the executed instructions into a high level program representation that can be reused as a normal C function in new software. The authors claim the concept to be resilient against the obfuscation techniques *packing/encryption, aliasing, control flow obfuscation* (e.g., *flattening*), *inserting dead code*, as well as several popular *anti-debugging techniques* such as “Soft Breakpoint Detection”, “Anti-VMware IN Instruction”, and “IsDebuggerPresent Check”. Following the results of Zeng et al. [2013] we also marked *parallelized code*, which is arguably less effective in dynamic code analysis scenarios, in gray.

Most other concepts introduced in recent literature are focused on malware. Leder et al. [2009] proposed a concept for automated extraction of cryptographic routines through dynamic data analysis. The automated isolation of a single function from a (malicious) binary was proposed by Caballero et al. [2010]. In contrast, INSPECTOR by Kolbitsch et al. [2010] allows the automated extraction of a particular malicious

behavior that does not necessarily have to be limited to one function of the program only. Following the evaluation of this approach, the dynamic obfuscation class of *packing/encryption* has to be considered weak against the described approach.

Other dynamic code rewriting techniques were marked in gray in Table II. In contrast to the locating code scenario we marked *virtualization* in black as the technique by Sharif et al. [2009] for analyzing virtualization-protected code is limited to locating code and does not yield valid executables.

Understanding code. Analogously to previous scenarios, a deeper understanding of the code of a program requires at least basic de-obfuscation in the automated analysis. In literature, several concepts were introduced. Udupa et al. [2005] proposed automated de-obfuscation of *control flow flattening* and *dead code* insertion using a hybrid approach of static and dynamic analysis techniques. The incomplete control flow graph from a dynamic analysis is enriched by adding some control flow edges that could possibly be taken through static analysis. While the results presented in the paper show that the evaluation of isolated analysis problems is possible, it is difficult to reason about the value of the concept for real-life programs.

5.4. Human analysis

The capabilities of a human code analyst are difficult to quantify in general. Tilley et al. [1996] first described a framework for program understanding including cognitive aspects of a human reverse engineer. However, it is still safe to assume that provided with sufficient patience a human analyst can break any class of software-only code obfuscation (i.e., an obfuscation where no trust anchor in hardware exists) and is only limited by scalability constraints.

Locating data. Most data obfuscation techniques have only limited strength in the analysis scenario of a human analyst trying to locate data structures in programs. One important concept for the protection of data (keys) inside a binary is *white-box cryptography*. It was proposed to prevent the extraction of a cryptographic key from the binary by mixing it with the algorithm. A decade ago, the first implementations of white-box algorithms for DES and AES have been proposed [Chow et al. 2003b,a; Link and Neumann 2005; Bringer et al. 2006]. However, all of them have been broken using techniques such as fault injection [Jacob et al. 2003], statistical analysis [Link et al. 2004], condensed implementation [Wyseur and Preneel 2005], differential cryptanalysis [Goubin et al. 2007; Wyseur et al. 2007; Billet et al. 2005; De Mulder et al. 2010] or generic cryptanalysis [Michiels et al. 2009]. Given this mixed history, in recent years, research on white-box cryptography focused on the question how the general idea and its security concepts can be backed by a theoretical foundation. Wyseur [2009] discussed the state of the art of *white-box cryptography* and proposed new block ciphers and design principles for the construction of white-box cryptographic algorithms. Saxena et al. [2009] described a theoretical model of *white-box cryptography* using appropriate security notions and presented both positive and negative results on white-box cryptography. This leads us to the conclusion that in its current state the strength of *white-box cryptography* is unproven, although it can make the extraction of the cryptographic key considerably more complex. Thus, we marked the obfuscation class *converting static data to procedures* in gray. Other obfuscation techniques that can provide at least limited resilience in this analysis scenario are *environmental requirements* and *hardware-assisted code obfuscation* because of their dependencies on external factors such as the presence of a particular hardware token. The strength of dongles for software protection was evaluated by Piazzalunga et al. [2007]. The authors developed a model for forecasting the amount of time an analyst would need to break dongle-based software protection schemes and concluded that today's available dongle solu-

tions provide only minimal protection. Still, we marked the class of *hardware-assisted code obfuscation* techniques in gray as dongles are only one simple concept in this class of obfuscations. Several approaches that are more resilient to dynamic analysis were introduced in the literature [Fu et al. 2007; Zhuang et al. 2004; Bitansky et al. 2011].

Locating Code. Some obfuscation techniques break abstractions in the code that aid human understanding. *Name scrambling*, *removing standard library calls*, and *breaking relations* do not prevent an automated tool from analyzing a program. However, they can make manual analysis by a human more difficult.

Other classes of code obfuscation have limited robustness against a human analyst trying to locate code. For example, Rolles [2009] introduced a semi-automated de-obfuscation approach against virtualization-obfuscation, which is based on reverse engineering the virtual machine, extracting the byte code, and then turning it into native code. Madou et al. [2006b,c] developed an interactive de-obfuscation tool named LOCO that allows an analyst to navigate through the control flow graph of a program to undo static obfuscating transformations such as *control flow flattening*. Quist and Liebrock [2009] demonstrated how a sophisticated visual representation of the control flow of a program can speed up the analysis process. In particular, unpacking routines of malware can be identified efficiently using a visualization approach.

Extracting Code. Finding the location of code is a prerequisite for code extraction. Thus, similar to the scenario of dynamic analysis, the extraction of code can be considered at least as difficult as locating code for a human analyst. One of the most popular approaches for human-assisted code extraction is *program slicing*, which is based on the idea of reducing the program code to a minimum *slice* that still produces a particular behavior or affects the value of a particular variable. In the literature, a multitude of program slicing approaches have been introduced over the past two decades (e.g., [Lanubile and Visaggio 1997] and [Ning et al. 1993]). A major limitation of program slicing is that the human analyst needs a deep understanding of the program internals to be able to specify a slicing criterion such as relevant variables and behavior. Several attempts have been made to raise the level of abstraction in slicing and thus make it less depended on manual analysis, such as conditioned slicing [Canfora et al. 1994, 1998; Cimitile et al. 1996; Danicic et al. 2004] and constraint slicing [Fox et al. 2004; Danicic et al. 2005; Field et al. 1995]. We marked *opaque predicates* and *aliasing* in gray as these obfuscation techniques can make the identification of the minimal subset that still implements a particular functionality more difficult.

Understanding Code. Despite the fact that getting a comprehensive understanding of the program structure and functionality can be considered the most ambitious aim of a human analyst, today's state of the art in code obfuscation provides only limited protection in this analysis scenario. This assumption is backed by a plethora of reports about successfully removed copy protection schemes for digital media such as CSS (DVD copy protection), Windows Media DRM [Myska 2009], and HDCP (high-bandwidth digital content protection) [Lomb and Guneysu 2011]. Another piece of evidence for this assumption is the successful reverse engineering of the VoIP (Voice over IP) software Skype. While Skype is known for its extensive use of code obfuscation, Biondi and Desclaux [2006] still were able to reveal the internal structure of the software. Furthermore, the client software of the cloud service provider Dropbox was successfully reverse engineered despite being heavily obfuscated [Kholia and Wegrzyn 2013]. A decisive factor for these recent success stories in code analysis are today's sophisticated reverse engineering tools that have become better and better in dealing with obfuscated code [Eagle 2008; Ferguson and Kaminsky 2008; Eilam 2005]. For instance, the F.L.I.R.T. library of IDA Pro enables the recognition of standard li-

brary functions generated by a variety of different compilers. It can be concluded that almost all code obfuscation techniques have to be considered ineffective against a human analyst that puts enough time and effort into manual de-obfuscation. We marked *name scrambling*, *removing standard library calls*, *breaking relations*, and *virtualization* in gray as these obfuscations can make manual analysis by a human more difficult. Furthermore, the obfuscation classes *environmental requirements* and *hardware-assisted code obfuscation* can be considered as strong against human analysis as long as the external requirement cannot be accessed by the analyst. *Anti-debugging* and *Anti-disassembly* was also marked in gray as all human-assisted analysis approaches described in the literature are still based on automated static and dynamic analysis tools.

6. CONCLUSIONS

In this survey we addressed the question to what extent software obfuscation is able to provide reasonable protection of programs against state-of-the-art code analysis techniques and tools. Despite more than two decades of research on obfuscation theory, reliable concepts for the evaluation of the resilience of an obfuscation technique have still not been found, and there are similar limitations for the evaluation of code analysis and de-obfuscation techniques. With these constraints in mind, we conducted a literature review of code analysis techniques applied against different classes of obfuscations in specific attack scenarios. Where we could not find direct evidence of the resilience of an obfuscation in a particular analysis scenario in the literature, we made and justified our own inferences. Hence, this survey should not be seen as a formal analysis but as a snapshot of the current state in the arms race between obfuscation and code analysis. We considered the effects of obfuscations applied in isolation of each other, but we would like to point out that the analysis of combinations of multiple types of obfuscations constitutes an interesting area for future research.

Our results indicate that the strength of obfuscations heavily depends on the goals of the analyst and the available resources. Most of the more heavyweight analysis approaches introduced in academia have only been demonstrated to work on relatively small and specific examples, whereas large real-world programs can be significantly harder to analyze. A major limiting factor for code analysis is that the high complexity of the analysis problems often exceeds the resources available to the analyst. Therefore, simple obfuscations can still be quite effective where the deployed techniques have to be fast and lightweight, like many pattern matching or static analysis algorithms. This explains the unbroken popularity of software obfuscation among malware writers. Where more resources are available to run an expensive dynamic analysis or even perform manual reverse engineering, obfuscations are much less effective. As a result, intellectual property protection against a human adversary remains challenging.

Another observation is that much of the current academic research on code analysis in the presence of obfuscations focuses on malware. Frequently, the relevant literature describes methods for classifying programs as malicious based on identifying obfuscation techniques that are common in malware. The analysis of the actual functionality of an obfuscated program is typically left out of scope. Substantially less academic research has been done on reverse engineering obfuscated general (non-malicious) programs, and tool support is worse.

The arms race between software obfuscation and analysis is still ongoing and the fundamental challenge of devising software protection mechanisms that are resistant against a human analyst remains. With today's software obfuscation techniques one has to assume that a dedicated analyst who is willing to spend enough time and effort will always be able to successfully analyze a program. Nevertheless, as this survey

shows, specific classes of software obfuscation can be effective in more restricted analysis scenarios.

Acknowledgments

This work has been carried out partially within the scope of TARGET, the Josef Ressel Center for Unified Threat Intelligence on Targeted Attacks. The financial support by the Austrian Federal Ministry of Science, Research and Economy and the National Foundation for Research, Technology and Development is gratefully acknowledged. The research was partially funded under Grant 826461 (FIT-IT) by the FFG – Austrian Research Promotion Agency.

REFERENCES

- ANATI, I., GUERON, S., JOHNSON, S., and SCARLATA, V. 2013. Innovative technology for cpu based attestation and sealing. In *Proceedings of the 2nd International Workshop on Hardware and Architectural Support for Security and Privacy*.
- ANCKAERT, B., DE SUTTER, B., and DE BOSSCHERE, K. 2004. Software Piracy Prevention through Diversity. In *Proceedings of the 4th Acm Workshop on Digital Rights Management*. ACM, 63–71.
- ANCKAERT, B., JAKUBOWSKI, M., and VENKATESAN, R. 2006. Proteus: Virtualization for Diversified Tamper-Resistance. In *Proceedings of the ACM Workshop on Digital Rights Management*. ACM, 47–58.
- ANCKAERT, B., JAKUBOWSKI, M. H., VENKATESAN, R., and SAW, C. W. 2009. Runtime Protection Via Dataflow Flattening. In *3rd International Conference on Emerging Security Information, Systems and Technologies (SECURWARE '09)*. IEEE, 242–248.
- ANCKAERT, B., MADOU, M., DE SUTTER, B., DE BUS, B., DE BOSSCHERE, K., and PRENEEL, B. 2007. Program Obfuscation: A Quantitative Approach. In *Proceedings of the 2007 ACM Workshop on Quality of Protection*. ACM, 15–20.
- AVOINE, G., JUNOD, P., and OECHSLIN, P. 2007. *Computer System Security: Basic Concepts and Solved Exercises*. EPFL Press.
- BACON, D., GRAHAM, S., and SHARP, O. 1994. Compiler Transformations for High-Performance Computing. *ACM Computing Surveys* 26, 4 (1994), 345–420.
- BALAKRISHNAN, G. and REPS, T. W. 2004. Analyzing Memory Accesses in x86 Executables. In *Compiler Construction*, Evelyn Duesterwald (Ed.). Vol. 2985. Springer, 5–23.
- BARAK, B., GARG, S., KALAI, Y. T., PANETH, O., and SAHAI, A. 2014. Protecting obfuscation against algebraic attacks. In *Advances in Cryptology–EUROCRYPT 2014*. Springer, 221–238.
- BARAK, B., GOLDREICH, O., IMPAGLIAZZO, R., RUDICH, S., SAHAI, A., VADHAN, S., and YANG, K. 2001. On the (Im)possibility of Obfuscating Programs. In *Advances in Cryptology–Crypto 2001*. Springer, 1–18.
- BARDIN, S., HERRMANN, P., and VÉDRINE, F. 2011. Refinement-Based CFG Reconstruction from Unstructured Programs. In *Proceedings of the 12th International Conference on Verification, Model Checking, and Abstract Interpretation (VMCAI '11)*. 54–69.
- BAYER, U., KRUEGEL, C., and KIRDA, E. 2006. TTAalyze: A Tool for Analyzing Malware. In *15th Annual Conference of the European Institute for Computer Antivirus Research (EICAR)*.
- BILAR, D. 2007. Opcodes as Predictor for Malware. *International Journal of Electronic Security and Digital Forensics* 1, 2 (2007), 156–168.
- BILLET, O., GILBERT, H., and ECH-CHATBI, C. 2005. Cryptanalysis of a White Box

- AES Implementation. In *Proceedings of the 11th International Conference on Selected Areas in Cryptography*. Springer, 227–240.
- BIONDI, P. and DESCLAUX, F. 2006. Silver Needle in the Skype. *Black Hat Europe 6* (2006), 25–47.
- BITANSKY, N., CANETTI, R., COHN, H., GOLDWASSER, S., KALAI, Y. T., PANETH, O., and ROSEN, A. 2014. The impossibility of obfuscation with auxiliary input or a universal simulator. In *Advances in Cryptology—CRYPTO 2014*. Springer, 71–89.
- BITANSKY, N., CANETTI, R., GOLDWASSER, S., HALEVI, S., KALAI, Y. T., and ROTHBLUM, G. N. 2011. Program Obfuscation with Leaky Hardware. In *Advances in Cryptology—Asiacrypt 2011*. Vol. 7073. Springer, 722–739.
- BOURQUIN, M., KING, A., and ROBBINS, E. 2013. BinSlayer: accurate comparison of binary executables. In *Proceedings of the 2nd ACM SIGPLAN Program Protection and Reverse Engineering Workshop*. ACM.
- BRAKERSKI, Z. and ROTHBLUM, G. N. 2014. Virtual black-box obfuscation for all circuits via generic graded encoding. In *Theory of Cryptography*. Springer, 1–25.
- BRANCO, R. R., BARBOSA, G. N., and NETO, P. D. 2012. Scientific but Not Academical Overview of Malware Anti-Debugging, Anti-Disassembly and Anti-VM Technologies. In *Blackhat 2012*.
- BRAND, M. 2010. *Analysis Avoidance Techniques of malicious software*. Ph.D. Dissertation. Edith Cowan University.
- BRINGER, J., CHABANNE, H., and DOTTA, E. 2006. White Box Cryptography: Another Attempt. *IACR Cryptology Eprint Archive 2006* (2006).
- BROSCH, T. and MORGENSTERN, M. 2006. Runtime Packers: The Hidden Problem. *Black Hat USA* (2006).
- BRUMLEY, D., JAGER, I., AVGERINOS, T., and SCHWARTZ, E. J. 2011. BAP: A Binary Analysis Platform. In *23th International Conference on Computer Aided Verification (CAV '11)*. 463–469.
- BRUMLEY, D., POOSANKAM, P., SONG, D., and ZHENG, J. 2008. Automatic patch-based exploit generation is possible: Techniques and implications. In *Security and Privacy, 2008. SP 2008. IEEE Symposium on*. IEEE, 143–157.
- BRUSCHI, D., MARTIGNONI, L., and MONGA, M. 2006a. Detecting Self-Mutating Malware Using Control-Flow Graph Matching. *Detection of Intrusions and Malware & Vulnerability Assessment* (2006), 129–143.
- BRUSCHI, D., MARTIGNONI, L., and MONGA, M. 2006b. Using code normalization for fighting self-mutating malware. In *Proceedings of the International Symposium on Secure Software Engineering*. 37–44.
- CABALLERO, J., JOHNSON, N. M., MCCAMANT, S., and SONG, D. 2010. Binary Code Extraction and Interface Identification for Security Applications. In *Proceedings of Network and Distributed System Security Symposium (NDSS '09)*.
- CADAR, C., GANESH, V., PAWLOWSKI, P. M., DILL, D. L., and ENGLER, D. R. 2006. EXE: Automatically Generating Inputs of Death. In *Proceedings of the 13th ACM Conference on Computer and Communications Security*. 322–335.
- CALVET, J., FERNANDEZ, J. M., and MARION, J.-Y. 2012. Aligot: Cryptographic Function Identification in Obfuscated Binary Programs. In *Proceedings of the 19th ACM Conference on Computer and Communications Security*. ACM, 169–182.
- CANETTI, R. and DAKDOUK, R. 2008. Obfuscating Point Functions with Multibit Output. *Advances in Cryptology—Eurocrypt 2008* (2008), 489–508.
- CANFORA, G., CIMITILE, A., and DE LUCIA, A. 1998. Conditioned Program Slicing. *Information and Software Technology* 40, 11 (1998), 595–607.
- CANFORA, G., CIMITILE, A., DE LUCIA, A., and DI LUCCA, G. A. 1994. Software Salvaging Based on Conditions. In *Proceedings of the International Conference on Software Maintenance (ICSM '94)*. IEEE, 424–433.

- CAPPAERT, J., KISSERLI, N., SCHELLEKENS, D., and PRENEEL, B. 2006. Self-Encrypting Code to Protect against Analysis and Tampering. In *1st Benelux Workshop on Information and System Security*.
- CAPPAERT, J. and PRENEEL, B. 2010. A General Model for Hiding Control Flow. In *Proceedings of the 10th Annual ACM Workshop on Digital Rights Management*. ACM, 35–42.
- CHANG, H. and ATALLAH, M. J. 2002. Protecting Software Code by Guards. In *Revised Papers from the ACM CCS-8 Workshop on Security and Privacy in Digital Rights Management*. Springer, 160–175.
- CHOUCHANE, M. R. and LAKHOTIA, A. 2006. Using Engine Signature to Detect Metamorphic Malware. In *Proceedings of the 4th ACM Workshop on Recurring Malcode*. ACM, 73–78.
- CHOW, S., EISEN, P., JOHNSON, H., and VAN OORSCHOT, P. 2003a. White-Box Cryptography and an AES Implementation. In *Revised Papers from the 9th Annual International Workshop on Selected Areas in Cryptography*. Springer, 250–270.
- CHOW, S., EISEN, P., JOHNSON, H., and VAN OORSCHOT, P. C. 2003b. A White-Box DES Implementation for DRM Applications. In *Digital Rights Management*. Vol. 2696. Springer, 1–15.
- CHOW, S., GU, Y., JOHNSON, H., and ZAKHAROV, V. A. 2001. An Approach to the Obfuscation of Control-Flow of Sequential Computer Programs. In *Information Security*. Springer, 144–155.
- CHRISTODORESCU, M., JHA, S., KINDER, J., KATZENBEISSER, S., and VEITH, H. 2007. Software Transformations to Improve Malware Detection. *Journal in Computer Virology* 3, 4 (2007), 253–265.
- CHRISTODORESCU, M., JHA, S., SESHIA, S., SONG, D., and BRYANT, R. 2005. Semantics-Aware Malware Detection. In *26th IEEE Symposium on Security and Privacy*. IEEE, 32–46.
- CIFUENTES, C. and GOUGH, K. J. 1995. Decompilation of Binary Programs. *Software: Practice and Experience* 25, 7 (1995), 811–829.
- CIMITILE, A., DE LUCIA, A., and MUNRO, M. 1996. A Specification Driven Slicing Process for Identifying Reusable Functions. *Journal of Software Maintenance: Research and Practice* 8, 3 (1996), 145–178.
- COHEN, F. 1993. Operating System Protection through Program Evolution. *Computers & Security* 12, 6 (1993), 565–584.
- COLLBERG, C. and NAGRA, J. 2009. *Surreptitious Software: Obfuscation, Watermarking, and Tamperproofing for Software Protection*. Addison-Wesley Professional.
- COLLBERG, C., THOMBORSON, C., and LOW, D. 1997. *A Taxonomy of Obfuscating Transformations*. Technical Report. Department of Computer Science, The University of Auckland, New Zealand.
- COLLBERG, C., THOMBORSON, C., and LOW, D. 1998a. Breaking Abstractions and Unstructuring Data Structures. In *Proceedings of the 1998 International Conference on Computer Languages*. IEEE, 28–38.
- COLLBERG, C., THOMBORSON, C., and LOW, D. 1998b. Manufacturing Cheap, Resilient, and Stealthy Opaque Constructs. In *Proceedings of the 25th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. ACM, 184–196.
- COMPARETTI, P. M., SALVANESCHI, G., KIRDA, E., KOLBITSCH, C., KRUEGEL, C., and ZANERO, S. 2010. Identifying Dormant Functionality in Malware Programs. In *30th IEEE Symposium on Security and Privacy*. IEEE, 61–76.
- COOGAN, K., DEBRAY, S., KAOCHAR, T., and TOWNSEND, G. 2009. Automatic Static Unpacking of Malware Binaries. In *16th Working Conference on Reverse Engineering (WCRE '09)*. IEEE, 167–176.
- COOGAN, K., LU, G., and DEBRAY, S. 2011. Deobfuscation of Virtualization-

- Obfuscated Software: A Semantics-Based Approach. In *Proceedings of the 18th ACM Conference on Computer and Communications Security*. ACM, 275–284.
- COPPENS, B., DE SUTTER, B., and MAEBE, J. 2013. Feedback-driven binary code diversification. *ACM Transactions on Architecture and Code Optimization (TACO)* 9, 4 (2013).
- COZZIE, A., STRATTON, F., XUE, H., and KING, S. T. 2008. Digging for Data Structures. In *Symposium on Operating Systems Design and Implementation (OSDI '08)*.
- CRANDALL, J. R., WASSERMANN, G., DE OLIVEIRA, D. A., SU, Z., WU, S. F., and CHONG, F. T. 2006. Temporal Search: Detecting Hidden Malware Timebombs with Virtual Machines. *ACM SIGPLAN Notices* 41, 11 (2006), 25–36.
- DALLA PREDI, M. and GIACOBAZZI, R. 2005. Semantic-Based Code Obfuscation by Abstract Interpretation. In *Automata, Languages and Programming*. Springer, 1325–1336.
- DALLA PREDI, M., GIACOBAZZI, R., DEBRAY, S., COOGAN, K., and TOWNSEND, G. 2011. Modelling Metamorphism by Abstract Interpretation. *Static Analysis* (2011), 218–235.
- DALLA PREDI, M., MADOU, M., DE BOSSCHERE, K., and GIACOBAZZI, R. 2006. Opaque Predicates Detection by Abstract Interpretation. *Algebraic Methodology and Software Technology* (2006), 81–95.
- DANICIC, S., DAOUDI, M., FOX, C., HARMAN, M., HIERONS, R. M., HOWROYD, J. R., OURABYA, L., and WARD, M. 2005. Consus: A Light-Weight Program Conditioner. *Journal of Systems and Software* 77, 3 (2005), 241–262.
- DANICIC, S., DE LUCIA, A., and HARMAN, M. 2004. Building Executable Union Slices Using Conditioned Slicing. In *Proceedings of the 12th IEEE International Workshop on Program Comprehension*. IEEE, 89–97.
- DAS, M., LERNER, S., and SEIGLE, M. 2002. ESP: Path-Sensitive Program Verification in Polynomial Time. In *Proceedings of the ACM SIGPLAN 2002 Conference on Programming Language Design and Implementation*. 57–68.
- DAVI, L., DMITRIENKO, A., NÜRNBERGER, S., and SADEGHI, A.-R. 2012. XIFER: A Software Diversity Tool against Code-Reuse Attacks. In *4th ACM International Workshop on Wireless of the Students, by the Students, for the Students (S3 2012)*.
- DE MULDER, Y., WYSEUR, B., and PRENEEL, B. 2010. Cryptanalysis of a Perturbated White-Box AES Implementation. In *Progress in Cryptology - INDOCRYPT 2010*. Springer, 292–310.
- DE SUTTER, B., ANCKAERT, B., GEIREGAT, J., CHANET, D., and DE BOSSCHERE, K. 2009. Instruction Set Limitation in Support of Software Diversity. *Information Security and Cryptology* (2009), 152–165.
- DEBRAY, S. and PATEL, J. 2010. Reverse Engineering Self-Modifying Code: Unpacker Extraction. In *17th Working Conference on Reverse Engineering (WCRE '10)*. IEEE, 131–140.
- DEDIĆ, N., JAKUBOWSKI, M., and VENKATESAN, R. 2007. A Graph Game Model for Software Tamper Protection. In *Proceedings of the 9th International Conference on Information Hiding*. Springer-Verlag, 80–95.
- DEPREZ, J. and LAKHOTIA, A. 2000. A Formalism to Automate Mapping from Program Features to Code. In *Proceedings of the 8th International Workshop on Program Comprehension*. IEEE, 69–78.
- DYER, K. P., COULL, S. E., RISTENPART, T., and SHRIMPTON, T. 2013. Protocol misidentification made easy with format-transforming encryption. In *Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security*. ACM, 61–72.
- EAGLE, C. 2008. *The IDA Pro Book: The Unofficial Guide to the World's Most Popular Disassembler*. No Starch Press.

- EGELE, M., SCHOLTE, T., KIRDA, E., and KRUEGEL, C. 2012. A Survey on Automated Dynamic Malware-Analysis Techniques and Tools. *ACM Computing Surveys* 44, 2 (2012).
- EILAM, E. 2005. *Reversing: Secrets of Reverse Engineering*. Wiley.
- EMMERIK, M. and WADDINGTON, T. 2004. Using a Decompiler for Real-World Source Recovery. In *In Proceedings of the 11th Working Conference on Reverse Engineering*. IEEE, 27–36.
- FERGUSON, J. and KAMINSKY, D. 2008. *Reverse Engineering Code with IDA Pro*. Syngress.
- FIELD, J., RAMALINGAM, G., and TIP, F. 1995. Parametric Program Slicing. In *Proceedings of the 22nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. ACM, 379–392.
- FLAKE, H. 2004. Structural Comparison of Executable Objects. In *Detection of Intrusions and Malware & Vulnerability Assessment, GI SIG SIDAR Workshop, DIMVA 2004, Dortmund, Germany, July 6-7, 2004, Proceedings*. 161–173.
- FOKET, C., DE SUTTER, B., COPPENS, B., and DE BOSSCHERE, K. 2013. A novel obfuscation: class hierarchy flattening. In *Foundations and Practice of Security*. Springer, 194–210.
- FOKET, C., DE SUTTER, B., and DE BOSSCHERE, K. 2014. Pushing Java Type Obfuscation to the Limit. *IEEE Transactions on Dependable and Secure Computing* 6 (2014), 553–567.
- FORREST, S., SOMAYAJI, A., and ACKLEY, D. H. 1997. Building diverse computer systems. In *Operating Systems, 1997., The Sixth Workshop on Hot Topics in*. IEEE, 67–72.
- FOX, C., DANICIC, S., HARMAN, M., and HIERONS, R. M. 2004. ConSIT: A Fully Automated Conditioned Program Slicer. *Software: Practice and Experience* 34, 1 (2004), 15–46.
- FRANZ, M. 2010. E Unibus Pluram: Massive-Scale Software Diversity as a Defense Mechanism. In *Proceedings of the 2010 Workshop on New Security Paradigms*. ACM, 7–16.
- FU, B., ARAVALLI, S., and ABRAHAM, J. 2007. Software Protection by Hardware and Obfuscation. In *Proceedings of the 2007 International Conference on Security & Management (SAM 2007)*. 367–373.
- GAO, D., REITER, M. K., and SONG, D. 2008. Binhunt: Automatically finding semantic differences in binary programs. In *Information and Communications Security*. Springer, 238–255.
- GARG, S., GENTRY, C., HALEVI, S., RAYKOVA, M., SAHAI, A., and WATERS, B. 2013. Candidate indistinguishability obfuscation and functional encryption for all circuits. In *Foundations of Computer Science (FOCS), 2013 IEEE 54th Annual Symposium on*. IEEE, 40–49.
- GHOSH, S., HISER, J. D., and DAVIDSON, J. W. 2010. A secure and robust approach to software tamper resistance. In *Information Hiding*. Springer, 33–47.
- GIACOBazzi, R. 2008. Hiding Information in Completeness Holes: New Perspectives in Code Obfuscation and Watermarking. In *6th IEEE International Conference on Software Engineering and Formal Methods (SEFM '08)*. IEEE, 7–18.
- GIACOBazzi, R. and MASTROENI, I. 2012. Making Abstract Interpretation Incomplete: Modeling the Potency of Obfuscation. In *Proceedings of the 19th International Symposium Static Analysis (SAS 2012)*. Springer, 129–145.
- GODEFROID, P., KLARLUND, N., and SEN, K. 2005. DART: Directed Automated Random Testing. In *Proceedings of the 2005 ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI '05)*. 213–223.
- GODEFROID, P., LEVIN, M. Y., and MOLNAR, D. A. 2008. Automated Whitebox Fuzz

- Testing. In *Proceedings of Network and Distributed System Security Symposium (NDSS '08)*.
- GOLDWASSER, S. and ROTHBLUM, G. N. 2007. On Best-Possible Obfuscation. In *Theory of Cryptography*. Vol. 4392. Springer, 194–213.
- GOUBIN, L., MASEREEL, J., and QUISQUATER, M. 2007. Cryptanalysis of White Box DES Implementations. In *Selected Areas in Cryptography*. Vol. 4876. Springer, 278–295.
- GRIFFIN, K., SCHNEIDER, S., HU, X., and CHIUEH, T. 2009. Automatic Generation of String Signatures for Malware Detection. In *Recent Advances in Intrusion Detection*. Lecture Notes in Computer Science, Vol. 5758. Springer, 101–120.
- GRÖBERT, F., WILLEMS, C., and HOLZ, T. 2011. Automated Identification of Cryptographic Primitives in Binary Programs. In *Recent Advances in Intrusion Detection*. Lecture Notes in Computer Science, Vol. 6961. Springer, 41–60.
- GROVER, D. 1992. *Protection of Computer Software: Its Technology and Application*. Cambridge University Press.
- GUILLOT, Y. and GAZET, A. 2010. Automatic Binary Deobfuscation. *Journal in Computer Virology* 6, 3 (2010), 261–276.
- GUIZANI, W., MARION, J.-Y., and REYNAUD-PLANTEY, D. 2009. Server-side dynamic code analysis. In *Malicious and unwanted software (MALWARE), 2009 4th international conference on*. IEEE, 55–62.
- HARRIS, L. and MILLER, B. 2005. Practical Analysis of Stripped Binary Code. *ACM SIGARCH Computer Architecture News* 33, 5 (2005), 63–68.
- HORNE, B., MATHESON, L., SHEEHAN, C., and TARJAN, R. E. 2002. Dynamic Self-Checking Techniques for Improved Tamper Resistance. In *Revised Papers from the ACM CCS-8 Workshop on Security and Privacy in Digital Rights Management*. Springer, 141–159.
- HORWITZ, S. 1997. Precise Flow-Insensitive May-Alias Analysis is NP-Hard. *ACM Transactions on Programming Languages and Systems* 19, 1 (1997), 1–6.
- JACOB, G., COMPARETTI, P. M., NEUGSCHWANDTNER, M., KRUEGEL, C., and VIGNA, G. 2012. A Static, Packer-Agnostic Filter to Detect Similar Malware Samples. In *Proceedings of the 9th International Conference on Detection of Intrusions and Malware, and Vulnerability Assessment*. Springer-Verlag, 102–122.
- JACOB, M., BONEH, D., and FELTEN, E. 2003. Attacking an Obfuscated Cipher by Injecting Faults. *Digital Rights Management (2003)*, 16–31.
- JACOB, M., JAKUBOWSKI, M. H., and VENKATESAN, R. 2007. Towards Integral Binary Execution: Implementing Oblivious Hashing Using Overlapped Instruction Encodings. In *Proceedings of the 9th Workshop on Multimedia & Security*. ACM, 129–140.
- JAKUBOWSKI, M., NALDURG, P., PATANKAR, V., and VENKATESAN, R. 2007. Software Integrity Checking Expressions (ICEs) for Robust Tamper Detection. In *Information Hiding*. Vol. 4567. Springer, 96–111.
- KANG, M. G., POOSANKAM, P., and YIN, H. 2007. Renovo: A Hidden Code Extractor for Packed Executables. In *Proceedings of the 2007 ACM Workshop on Recurring Malcode*. ACM, 46–53.
- KANZAKI, Y., MONDEN, A., NAKAMURA, M., and MATSUMOTO, K.-I. 2003. Exploiting Self-Modification Mechanism for Program Protection. In *Proceedings of the 27th Annual International Conference on Computer Software and Applications*. IEEE, 170–179.
- KARNIK, A., GOSWAMI, S., and GUHA, R. 2007. Detecting Obfuscated Viruses Using Cosine Similarity Analysis. In *1st Asia International Conference on Modelling & Simulation (AMS '07)*. IEEE, 165–170.
- KHOLIA, D. and WEGRZYN, P. 2013. Looking inside the (Drop)Box. In *7th Usenix*

- Workshop on Offensive Technologies (Woot '13)*.
- KINDER, J. 2012. Towards Static Analysis of Virtualization-Obfuscated Binaries. In *Proceedings of the 19th Working Conf. Reverse Engineering (WCRE 2012)*. IEEE, 61–70.
- KINDER, J., KATZENBEISSER, S., SCHALLHART, C., and VEITH, H. 2005. Detecting Malicious Code by Model Checking. In *Detection of Intrusions and Malware, and Vulnerability Assessment*. Vol. 3548. Springer, 174–187.
- KINDER, J. and VEITH, H. 2008. Jakstab: A Static Analysis Platform for Binaries. In *20th International Conference on Computer Aided Verification (CAV '08)*. Springer, 423–427.
- KINDER, J., ZULEGER, F., and VEITH, H. 2009. An Abstract Interpretation-Based Framework for Control Flow Reconstruction from Binaries. In *Proceedings of the 10th International Conference on Verification, Model Checking, and Abstract Interpretation (VMCAI '09)*. Springer, 214–228.
- KING, J. C. 1976. Symbolic execution and program testing. *Commun. ACM* 19, 7 (1976), 385–394.
- KING, S. and CHEN, P. 2006. SubVirt: Implementing Malware with Virtual Machines. In *27th IEEE Symposium on Security and Privacy*. IEEE.
- KOLBITSCH, C., HOLZ, T., KRUEGEL, C., and KIRDA, E. 2010. Inspector Gadget: Automated Extraction of Proprietary Gadgets from Malware Binaries. In *30th IEEE Symposium on Security and Privacy*. IEEE, 29–44.
- KOLBITSCH, C., KIRDA, E., and KRUEGEL, C. 2011. The Power of Procrastination: Detection and Mitigation of Execution-Stalling Malicious Code. In *Proceedings of the 18th ACM Conference on Computer and Communications Security*. ACM, 285–296.
- KRÜGEL, C., ROBERTSON, W. K., VALEUR, F., and VIGNA, G. 2004. Static Disassembly of Obfuscated Binaries. In *Usenix Security Symposium*. 255–270.
- LAKHOTIA, A., BOCCARDO, D. R., SINGH, A., and MANACERO JR, A. 2010. Context-sensitive analysis without calling-context. *Higher-Order and Symbolic Computation* 23, 3 (2010), 275–313.
- LANUBILE, F. and VISAGGIO, G. 1997. Extracting Reusable Functions by Flow Graph Based Program Slicing. *IEEE Transactions on Software Engineering* 23, 4 (1997), 246–259.
- LÁSZLÓ, T. and KISS, Á. 2009. Obfuscating C++ Programs Via Control Flow Flattening. *Annales Universitatis Scientiarum Budapestinensis De Rolando Eötvös Nominatae, Sectio Computatorica* 30 (2009), 3–19.
- LEDER, F., MARTINI, P., and WICHMANN, A. 2009. Finding and Extracting Crypto Routines from Malware. In *IEEE 28th International Performance Computing and Communications Conference (IPCCC '09)*. IEEE, 394–401.
- LI, J., XU, M., ZHENG, N., and XU, J. 2009. Malware Obfuscation Detection Via Maximal Patterns. In *3rd International Symposium on Intelligent Information Technology Application (LITA '09)*, Vol. 2. IEEE, 324–328.
- LIN, Z., ZHANG, X., and XU, D. 2010. Automatic Reverse Engineering of Data Structures from Binary Execution. In *17th Network and Distributed System Security Symposium*.
- LINK, H. E. and NEUMANN, W. D. 2005. Clarifying Obfuscation: Improving the Security of White-Box DES. In *International Conference on Information Technology: Coding and Computing (ITCC '05)*, Vol. 1. IEEE, 679–684.
- LINK, H. E., SCHROEPEL, R. C., NEUMANN, W. D., CAMPBELL, P. L., BEAVER, C. L., PIERSON, L. G., and ANDERSON, W. E. 2004. *Securing Mobile Code*. Technical Report. Sandia National Laboratories.
- LINN, C. and DEBRAY, S. 2003. Obfuscation of Executable Code to Improve Resistance to Static Disassembly. In *Proceedings of the 10th ACM Conference on Computer and*

- Communications Security*. ACM, 290–299.
- LOMB, B. and GUNEYSU, T. 2011. Decrypting HDCP-Protected Video Streams Using Reconfigurable Hardware. In *International Conference on Reconfigurable Computing and FPGAs (ReConFig '11)*. IEEE, 249–254.
- LYNN, B., PRABHAKARAN, M., and SAHAI, A. 2004. Positive Results and Techniques for Obfuscation. In *Advances in Cryptology–Eurocrypt 2004*. Springer, 20–39.
- MADOU, M., ANCKAERT, B., BUS, B. D., BOSSCHERE, K. D., CAPPAERT, J., and PRENEEL, B. 2006. On the Effectiveness of Source Code Transformations for Binary Obfuscation. In *Proceedings of the International Conference on Software Engineering Research and Practice (SERP06)*. 527–533.
- MADOU, M., ANCKAERT, B., DE SUTTER, B., and DE BOSSCHERE, K. 2005. Hybrid Static-Dynamic Attacks against Software Protection Mechanisms. In *Proceedings of the 5th Acm Workshop on Digital Rights Management*. ACM, 75–82.
- MADOU, M., ANCKAERT, B., MOSELEY, P., DEBRAY, S., DE SUTTER, B., and DE BOSSCHERE, K. 2006a. Software Protection through Dynamic Code Mutation. In *Information Security Applications*. Springer, 194–206.
- MADOU, M., VAN PUT, L., and DE BOSSCHERE, K. 2006b. LOCO: An Interactive Code (De) Obfuscation Tool. In *Proceedings of the 2006 ACM SIGPLAN Symposium on Partial Evaluation and Semantics-Based Program Manipulation*. ACM, 140–144.
- MADOU, M., VAN PUT, L., and DE BOSSCHERE, K. 2006c. Understanding Obfuscated Code. In *14th IEEE International Conference on Program Comprehension (ICPC '06)*. IEEE, 268–274.
- MAJUMDAR, A., MONSIFROT, A., and THOMBORSON, C. 2006. On Evaluating Obfuscatory Strength of Alias-Based Transforms Using Static Analysis. In *International Conference on Advanced Computing and Communications (ADCOM 2006)*. IEEE, 605–610.
- MAJUMDAR, A. and THOMBORSON, C. 2006. Manufacturing opaque predicates in distributed systems for code obfuscation. In *Proceedings of the 29th Australasian Computer Science Conference-Volume 48*. Australian Computer Society, Inc., 187–196.
- MASON, J., SMALL, S., MONROSE, F., and MACMANUS, G. 2009. English shellcode. In *Proceedings of the 16th ACM conference on Computer and communications security*. ACM, 524–533.
- MATROSOV, A., RODIONOV, E., HARLEY, D., and MALCHO, J. 2010. Stuxnet under the microscope. *ESET LLC (September 2010)* (2010).
- MAVROGIANNOPOULOS, N., KISSERLI, N., and PRENEEL, B. 2011. A Taxonomy of Self-Modifying Code for Obfuscation. *Computers & Security* 30, 8 (2011), 679–691.
- MICHIELS, W., GORISSEN, P., and HOLLMANN, H. D. 2009. Cryptanalysis of a Generic Class of White-Box Implementations. In *Selected Areas in Cryptography*. Vol. 5381. Springer, 414–428.
- MILES, C., LAKHOTIA, A., and WALLENSTEIN, A. 2012. In situ reuse of logically extracted functional components. *Journal in Computer Virology* 8, 3 (2012), 73–84.
- MONDEN, A., MONSIFROT, A., and THOMBORSON, C. 2004. A framework for obfuscated interpretation. In *Proceedings of the second workshop on Australasian information security, Data Mining and Web Intelligence, and Software Internationalisation-Volume 32*. Australian Computer Society, Inc., 7–16.
- MOSER, A., KRUEGEL, C., and KIRDA, E. 2007a. Exploring Multiple Execution Paths for Malware Analysis. In *28th IEEE Symposium on Security and Privacy*. IEEE, 231–245.
- MOSER, A., KRUEGEL, C., and KIRDA, E. 2007b. Limits of Static Analysis for Malware Detection. In *23rd Annual Computer Security Applications Conference (ACSAC '07)*. IEEE, 421–430.

- MYSKA, M. 2009. The True Story of DRM. *Masaryk Ujl & Tech.* 3 (2009).
- NACHENBERG, C. 1997. Computer Virus-Coevolution. *Communications of the ACM* 50, 1 (1997), 46–51.
- NAGARAJAN, V., GUPTA, R., ZHANG, X., MADOU, M., and DE SUTTER, B. 2007. Matching control flow of program versions. In *Software Maintenance, 2007. ICSM 2007. IEEE International Conference on.* IEEE, 84–93.
- NEWSOME, J., KARP, B., and SONG, D. 2005. Polygraph: Automatically Generating Signatures for Polymorphic Worms. In *26th IEEE Symposium on Security and Privacy.* IEEE, 226–241.
- NIELSON, F., NIELSON, H. R., and HANKIN, C. 1999. *Principles of Program Analysis.* Springer.
- NING, J. Q., ENGBERTS, A., and KOZACZYNSKI, W. 1993. Recovering Reusable Components from Legacy Systems by Program Segmentation. In *Proceedings of the Working Conference on Reverse Engineering.* IEEE, 64–72.
- PALSBERG, J., KRISHNASWAMY, S., KWON, M., MA, D., SHAO, Q., and ZHANG, Y. 2000. Experience with software watermarking. In *Computer Security Applications, 2000. ACSAC'00. 16th Annual Conference.* IEEE, 308–316.
- PIAZZALUNGA, U., SALVANESCHI, P., BALDUCCI, F., JACOMUZZI, P., and MORONCELLI, C. 2007. Security Strength Measurement for Dongle-Protected Software. *Security & Privacy, IEEE* 5, 6 (2007), 32–40.
- POPOV, I. V., DEBRAY, S. K., and ANDREWS, G. R. 2007. Binary Obfuscation Using Signals. In *Usenix Security Symposium.* 275–290.
- QUIST, D. A. and LIEBROCK, L. M. 2009. Visualizing Compiled Executables for Malware Analysis. In *6th International Workshop on Visualization for Cyber Security, 2009 (VizSec '09).* IEEE, 27–32.
- RABER, J. and LASPE, E. 2007. Deobfuscator: An Automated Approach to the Identification and Removal of Code Obfuscation. In *14th Working Conference on Reverse Engineering (WCRE '07).* IEEE, 275–276.
- RAMALINGAM, G. 1994. The Undecidability of Aliasing. *ACM Transactions on Programming Languages and Systems* 16, 5 (1994), 1467–1471.
- RIORDAN, J. and SCHNEIER, B. 1998. Environmental Key Generation Towards Clueless Agents. *Mobile Agents and Security* (1998), 15–24.
- ROLLES, R. 2009. Unpacking Virtualization Obfuscators. In *3rd Usenix Workshop on Offensive Technologies (Woot '09).*
- ROUNDY, K. A. and MILLER, B. P. 2013. Binary-Code Obfuscations in Prevalent Packer Tools. *ACM Computing Surveys* 46, 1 (2013).
- ROYAL, P., HALPIN, M., DAGON, D., EDMONDS, R., and LEE, W. 2006. Polyunpack: Automating the Hidden-Code Extraction of Unpack-Executing Malware. In *22nd Annual Computer Security Applications Conference (ACSAC '06).* IEEE, 289–300.
- RUGABER, S., STIREWALT, K., and WILLS, L. 1995. The Interleaving Problem in Program Understanding. In *Proceedings of the 2nd Working Conference on Reverse Engineering.* IEEE, 166–175.
- SAKABE, Y., SOSHI, M., and MIYAJI, A. 2005. Java obfuscation approaches to construct tamper-resistant object-oriented programs. *IPSJ Digital Courier* 1 (2005), 349–361.
- SAXENA, A., WYSEUR, B., and PRENEEL, B. 2009. Towards Security Notions for White-Box Cryptography. In *Information Security.* Springer, 49–58.
- SCHRITTWIESER, S. and KATZENBEISSER, S. 2011. Code Obfuscation against Static and Dynamic Reverse Engineering. In *Proceedings of the 13th International Conference on Information Hiding (IH '11).* Springer, 270–284.
- SCHRITTWIESER, S., KATZENBEISSER, S., KIESEBERG, P., HUBER, M., LEITHNER, M., MULAZZANI, M., and WEIPPL, E. 2013. Covert Computation: Hiding Code in

- Code for Obfuscation Purposes. In *Proceedings of the 8th ACM SIGSAC Symposium on Information, Computer and Communications Security*. ACM, 529–534.
- SCHWARTZ, E. J., AVGERINOS, T., and BRUMLEY, D. 2010. All You Ever Wanted to Know about Dynamic Taint Analysis and Forward Symbolic Execution (but Might Have Been Afraid to Ask). In *31st IEEE Symposium on Security and Privacy, S&P 2010*. 317–331.
- SCHWARTZ, E. J., LEE, J., WOO, M., and BRUMLEY, D. 2013. Native x86 Decompilation Using Semantics-Preserving Structural Analysis and Iterative Control-Flow Structuring. In *Usenix Security Symposium*.
- SCHWARZ, B., DEBRAY, S., and ANDREWS, G. 2002. Disassembly of Executable Code Revisited. In *9th Working Conference on Reverse Engineering*. IEEE, 45–54.
- SEN, K., MARINOV, D., and AGHA, G. 2005. CUTE: a concolic unit testing engine for C. In *Proceedings of the 10th European Software Engineering Conference held jointly with 13th ACM SIGSOFT International Symposium on Foundations of Software Engineering, 2005, Lisbon, Portugal, September 5-9, 2005*. 263–272.
- SHACHAM, H. 2007. The geometry of innocent flesh on the bone: Return-into-libc without function calls (on the x86). In *Proceedings of the 14th ACM conference on Computer and communications security*. ACM, 552–561.
- SHAMIR, A. and VAN SOMEREN, N. 1999. Playing 'Hide and Seek' with Stored Keys. In *Financial Cryptography*. Vol. 1648. Springer, 118–124.
- SHARIF, M., LANZI, A., GIFFIN, J., and LEE, W. 2009. Automatic Reverse Engineering of Malware Emulators. In *30th IEEE Symposium on Security and Privacy*. IEEE, 94–109.
- SHARIF, M., YEGNESWARAN, V., SAIDI, H., PORRAS, P., and LEE, W. 2008. Eureka: A Framework for Enabling Static Malware Analysis. *Computer Security-Esorics 2008* (2008), 481–500.
- SHARIF, M. I., LANZI, A., GIFFIN, J. T., and LEE, W. 2008. Impeding Malware Analysis Using Conditional Code Obfuscation.. In *Proceedings of Network and Distributed System Security Symposium (NDSS '08)*.
- SLOWINSKA, A., STANCESCU, T., and BOS, H. 2011. Howard: A Dynamic Excavator for Reverse Engineering Data Structures. In *Proceedings of Network and Distributed System Security Symposium (NDSS '11)*.
- SNEED, H. M. 2000. Encapsulation of Legacy Software: A Technique for Reusing Legacy Software Components. *Annals of Software Engineering* 9, 1-2 (2000), 293–313.
- SONG, D., BRUMLEY, D., YIN, H., CABALLERO, J., JAGER, I., KANG, M. G., LIANG, Z., NEWSOME, J., POOSANKAM, P., and SAXENA, P. 2008. BitBlaze: A New Approach to Computer Security Via Binary Analysis. In *Proceedings of the 4th International Conference on Information Systems Security. Keynote Invited Paper*.
- SONG, Y., LOCASTO, M. E., STAVROU, A., KEROMYTIS, A. D., and STOLFO, S. J. 2010. On the infeasibility of modeling polymorphic shellcode. *Machine learning* 81, 2 (2010), 179–205.
- SOSONKIN, M., NAUMOVICH, G., and MEMON, N. 2003. Obfuscation of design intent in object-oriented applications. In *Proceedings of the 3rd ACM workshop on Digital rights management*. ACM, 142–153.
- STEWART, J. 2006. Ollybone: Semi-Automatic Unpacking on IA-32. In *Proceedings of the 14th Def Con Hacking Conference*.
- TANG, Y. and CHEN, S. 2007. An Automated Signature-Based Approach against Polymorphic Internet Worms. *IEEE Transactions on Parallel and Distributed Systems* 18, 7 (2007).
- THAKUR, A., LIM, J., LAL, A., BURTON, A., DRISCOLL, E., ELDER, M., ANDERSEN, T., and REPS, T. 2010. Directed Proof Generation for Machine Code. In *22th Inter-*

- national Conference on Computer Aided Verification (CAV '10)*. Springer, 288–305.
- TILLEY, S., PAUL, S., and SMITH, D. 1996. Towards a Framework for Program Understanding. In *4th Workshop on Program Comprehension*. IEEE, 19–28.
- TREADWELL, S. and ZHOU, M. 2009. A Heuristic Approach for Detection of Obfuscated Malware. In *IEEE International Conference on Intelligence and Security Informatics (ISI '09)*. IEEE, 291–299.
- TSAI, H., HUANG, Y., and WAGNER, D. 2009. A Graph Approach to Quantitative Analysis of Control-Flow Obfuscating Transformations. *IEEE Transactions on Information Forensics and Security* 4, 2 (2009), 257–267.
- UDUPA, S., DEBRAY, S., and MADOU, M. 2005. Deobfuscation: Reverse Engineering Obfuscated Code. In *12th Working Conference on Reverse Engineering*. IEEE.
- VRBA, Z., HALVORSEN, P., and GRIWODZ, C. 2010. Program Obfuscation by Strong Cryptography. In *International Conference on Availability, Reliability, and Security (ARES '10)*. IEEE, 242–247.
- WALENSTEIN, A., MATHUR, R., CHOUCHANE, M., and LAKHOTIA, A. 2006. Normalizing Metamorphic Malware Using Term Rewriting. In *6th IEEE International Workshop on Source Code Analysis and Manipulation (SCAM '06)*. IEEE, 75–84.
- WANG, C., DAVIDSON, J., HILL, J., and KNIGHT, J. 2001. Protection of Software-Based Survivability Mechanisms. In *Proceedings of the 2001 International Conference on Dependable Systems and Networks (Formerly: FTCS)*. IEEE, 193–202.
- WANG, C., HILL, J., KNIGHT, J., and DAVIDSON, J. 2000. *Software Tamper Resistance: Obstructing Static Analysis of Programs*. Technical Report. CS-2000-12, University of Virginia.
- WEBSTER, M. and MALCOLM, G. 2009. Detection of Metamorphic and Virtualization-Based Malware Using Algebraic Specification. *Journal in Computer Virology* 5, 3 (2009), 221–245.
- WEE, H. 2005. On Obfuscating Point Functions. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing*. ACM, 523–532.
- WILDE, N. and SCULLY, M. 2006. Software Reconnaissance: Mapping Program Features to Code. *Journal of Software Maintenance: Research and Practice* 7, 1 (2006), 49–62.
- WILLEMS, C. and FREILING, F. C. 2012. Reverse Code Engineering-State of the Art and Countermeasures. *IT-Information Technology* 54, 2 (2012), 53–63.
- WOLFE, M., SHANKLIN, C., and ORTEGA, L. 1995. *High Performance Compilers for Parallel Computing*. Addison-Wesley Longman Publishing Co., Inc.
- WU, Z., GIANVECCHIO, S., XIE, M., and WANG, H. 2010. Mimimorphism: A New Approach to Binary Code Obfuscation. In *Proceedings of the 17th ACM Conference on Computer and Communications Security*. ACM, 536–546.
- WYSEUR, B. 2009. *White-Box Cryptography*. Ph.D. Dissertation. KU Leuven.
- WYSEUR, B., MICHIELS, W., GORISSEN, P., and PRENEEL, B. 2007. Cryptanalysis of White-Box DES Implementations with Arbitrary External Encodings. In *Proceedings of the 14th International Conference on Selected Areas in Cryptography*. Springer, 264–277.
- WYSEUR, B. and PRENEEL, B. 2005. Condensed White-Box Implementations. In *Proceedings of the 26th Symposium on Information Theory in the Benelux*. 296–301.
- YAKDAN, K., ESCHWEILER, S., GERHARDS-PADILLA, E., and SMITH, M. 2015. No More Gotos: Decompilation Using Pattern-Independent Control-Flow Structuring and Semantics-Preserving Transformations. *Proceedings of the 22nd Network and Distributed Systems Security Symposium (NDSS)* (2015).
- YIN, H. and SONG, D. 2010. *TEMU: Binary Code Analysis Via Whole-System Layered Annotative Execution*. Technical Report UCB/EECS-2010-3. EECS Department, University of California, Berkeley.

- ZENG, J., FU, Y., MILLER, K. A., LIN, Z., ZHANG, X., and XU, D. 2013. Obfuscation Resilient Binary Code Reuse through Trace-Oriented Programming. In *Proceedings of the 20th ACM Conference on Computer and Communications Security*.
- ZHANG, X. and GUPTA, R. 2005. Matching execution histories of program versions. In *ACM SIGSOFT Software Engineering Notes*, Vol. 30. ACM, 197–206.
- ZHAO, Z., AHN, G., and HU, H. 2011. Automatic Extraction of Secrets from Malware. In *18th Working Conference on Reverse Engineering (WCRE '11)*. IEEE, 159–168.
- ZHOU, Y., MAIN, A., GU, Y. X., and JOHNSON, H. 2007. Information hiding in software with mixed boolean-arithmetic transforms. In *Information Security Applications*. Springer, 61–75.
- ZHUANG, X., ZHANG, T., LEE, H., and PANDE, S. 2004. Hardware Assisted Control Flow Obfuscation for Embedded Processors. In *Proceedings of the 2004 International Conference on Compilers, Architecture, and Synthesis for Embedded Systems*. 292–302.