Probfuscation: An Obfuscation Approach using Probabilistic Control Flows

Andre Pawlowski, Moritz Contag, and Thorsten Holz

Horst Görtz Institute (HGI), Ruhr-University Bochum, Germany

Abstract. Sensitive parts of a program, such as proprietary algorithms or licensing information, are often protected with the help of code obfuscation techniques. Many obfuscation schemes transform the control flow of the protected program. Typically, the control flow of obfuscated programs is deterministic, *i. e.*, recorded execution traces do not differ for multiple executions using the same input values. An adversary can take advantage of this behavior and create multiple traces to perform analyses on the target program in order to deobfuscate it.
In this paper, we introduce an obfuscation approach which yields *probabilistic control flow* within a given method. That is, for the same input values, multiple execution traces differ, whilst preserving semantics. This effectively renders analyses relying on multiple traces impractical. We have implemented a prototype and applied it to multiple different programs. Our experimental results show that our approach can be used effectively to ensure divergent traces for the same input values and it can significantly improve the resilience against dynamic analysis.

1 Introduction

Obfuscation (lat. *obfuscare* = darken) is the art of disguising a given system such that the analysis becomes harder. In the area of software engineering, obfuscation can be used on either the source code or binary level to obscure the code or data flow. Generally speaking, the goal is to hamper reverse engineering. Code obfuscation plays an important role in practice and such techniques are widely used. On the one hand, obfuscation techniques can be used to protect programs from reverse engineering or to at least increase the costs for such an analysis. Examples include protection systems for sensitive parts or proprietary algorithms of a given program, or digital rights management systems that contain licensing information. On the other hand, obfuscation is widely used by attackers to impede analysis of malicious software such that antivirus companies have a harder time to analyze new samples. As a result, many different kinds of obfuscation techniques were proposed in the last years (*e. g.*, [6, 10, 13, 15]). Note that all obfuscation techniques have one constraint in common: the transformations used to obfuscate the program must ensure that the semantic meaning of the program is not changed.

Current state-of-the-art obfuscation techniques translate the target program's code into custom bytecode [17, 22]. This bytecode is generated specifically for

the obfuscated program and an interpreter is embedded which handles execution of said bytecode. When analyzed statically, the translation to an unknown instruction set forces an analyst to examine the bytecode interpreter first, before actually reverse engineering the original algorithm. Because obfuscation schemes are often difficult to analyze statically, most deobfuscation approaches make use of dynamic analysis [7, 21, 25]. A drawback of current obfuscation techniques is the fact that the control flow does not differ for multiple program executions when using the same input values. Thus, it is easier for an analyst to monitor control flow, which exposes parts of the semantic of the target program. Note that state-of-the-art deobfuscation tools utilize a dynamic trace of the program to reconstruct an unobfuscated version of the program.

In this paper, we propose a novel obfuscation approach that tackles the aforementioned problem. Our obfuscation scheme is constructed in such a way that multiple traces of the same function with the same input values lead to different *observed* control flows, whilst preserving semantics. Our approach is inspired by the idea of Collberg et al. [5], which uses opaque predicates manufactured using a specifically crafted graph data structure. However, their technique is based on a problem that is only difficult to tackle when the attacker is limited to static analysis. Hence, if an analyst employs dynamic analyses, she can easily determine the value of an opaque predicate which has been executed in the recorded trace.

In an empirical evaluation, we show that our proposed obfuscation approach successfully introduces probabilism to the control flow of the target program. Thus, it thwarts dynamic analysis operating on multiple executions of the protected program significantly and does not focus solely on static analysis like other state-of-the-art obfuscation approaches [6, 13, 17, 22].

In summary, we make the following contributions:

- We present a novel obfuscation scheme that introduces probabilistic control flow, but still ensures that the code's semantics are preserved. Due the probabilistic nature of our scheme, it can withstand proposed deobfuscation approaches that rely on a trace-based analysis of several execution runs.
- We implemented a proof-of-concept obfuscation tool for the managed code programming language C#. The tool is freely available at `https://github.com/RUB-SysSec/Probfuscator`.
- We evaluate the prototype and demonstrate that probabilistic obfuscation is a viable obfuscation technique to protect sensitive parts of a given program.

2 Technical Background

The transformations applied by the obfuscation process aim to hide the program's semantics. If successful, the analysis and deobfuscation effort is considerably higher than feasible for an analyst. In the following, we refer to an analyst as *adversary*.

The main class of obfuscation schemes, as well as ours, target the control flow of the target program since it contains vital information about the general structure of a program and exposes high-level constructs such as loops or if-clauses. Doing so, these obfuscation schemes thwart attempts to statically analyze the target program. One building block used by said schemes is the construct of *opaque predicates* [5]. An opaque predicate is a boolean expression whose value is known at obfuscation time. However, its value is difficult to infer by an (automated) attacker. Collberg et al. introduce three types of opaque predicates which we will refer to as *true opaque predicates*, *false opaque predicates*, and *random opaque predicates*, whose expressions evaluate to the boolean values *true*, *false* or evaluate randomly to either, respectively [5]. In the following, we will denote by *(always) taken branch* the branch of an opaque predicate which is known to be always taken.

In case of a true opaque predicate, its taken branch will always be taken, as it corresponds to the predicate evaluating to true. Its other branch also has to point to meaningful code, though, and points to a block of *dead code*. Preferably, from the obfuscator's point of view, it should be difficult to distinguish dead from live code. False opaque predicates operate analogously. Random opaque predicates differ in that their expression yields a random value and both branches may be taken. Consequently, the code blocks the branches point to have to be semantically equivalent for the obfuscation to be semantics-preserving. A resilient random opaque predicate aims to hide this fact by employing several transformations on the blocks to make comparison of their semantics infeasible.

Attacks against opaque predicates make use of *data flow analysis* and try to prove that the expression the predicate checks in fact constant. More resilient opaque predicates hence build expressions involving *pointer aliases* by making use of the hardness of the *intraprocedural may-alias analysis problem* [20]. This problem states that it is generally undecidable if two given pointers into a complex data structure alias each other, *i. e.*, point to the same location in the structure. While algorithms that tackle the problem do exist, many of them are incapable of handling special cases like recursive or cyclic data structures [5].

3 Adversary Model

The goal of the adversary is to analyze and understand a protected algorithm inside the obfuscated method (*e. g.*, a serial key check algorithm or a proprietary algorithm embedded in the method). To this end, the adversary has to understand the effect of the input values on the program's observable behavior, among others. We assume an adversary that bases her deobfuscation attempts solely on dynamic analysis techniques, a common attacker model found in recent literature on attacks against obfuscation schemes [7, 21, 25].

The adversary is able to record multiple traces of the obfuscated method for any inputs as well as set breakpoints on specific points in the control flow. Note that deobfuscation with the help of static analysis is already tackled by obfuscation techniques [1, 5, 20, 23] proposed previously, which are orthogonal to

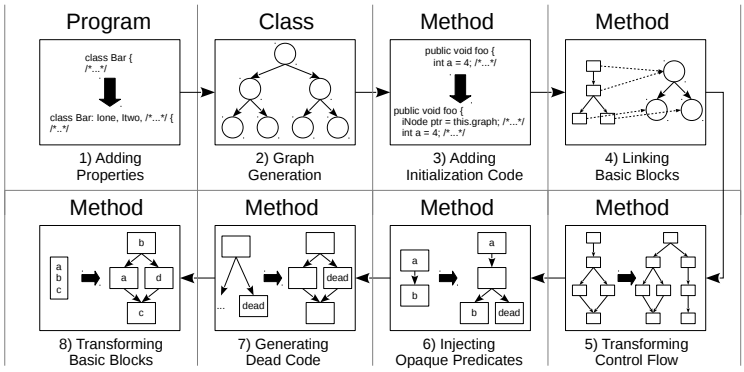

Fig. 1. The eight steps of the obfuscation process. On the top, it is noted which entity is targeted by the current obfuscation step.

our approach. However, the adversary is subject to time constraints in her analysis. Since modern programs change their protection implementations with the release of new versions (*e. g.*, anti-cheat systems, [14]) and recent deobfuscation approaches work solely on execution traces [7,21,25], we deem these assumptions reasonable.

4 Approach

Our approach makes use of an artificial graph, called *obfuscation graph*, whose nodes consist of objects of classes provided by the target program. Each protected method in the target program holds a pointer to the graph, linking both together. Each basic block of the protected method is linked to one or multiple nodes in the obfuscation graph. During the execution of the protected method, the pointer to the obfuscation graph is moved from node to node. The obfuscation only forwards the pointer to nodes linked to the basic blocks which are to be executed next. With the help of opaque predicates, the scheme ensures that tampering with the link most likely results in a crash of the program.

The obfuscation scheme consists of eight steps which are illustrated in Figure 1 and shortly described in the following.

1. *Adding properties.* The scheme uses properties of the nodes in the obfuscation graph for opaque predicates. In order to increase the number of possible opaque predicates, additional properties are added to the nodes.
2. *Generating the obfuscation graph.* The obfuscator then builds the obfuscation graph with the help of the properties. It is then added to the class that contains the method that should be protected.
3. *Adding initialization code.* This step adds additional logic to initialize the obfuscation scheme for all methods that are to be protected.

4. *Linking basic blocks.* The basic blocks of the control flow graph (CFG) are linked to the nodes of the obfuscation graph. This connection is needed to ensure correct evaluation of the boolean expressions of the opaque predicates.
5. *Transforming control flow.* The CFG of the method is transformed with the help of the linked obfuscation graph in such a way that multiple paths through the CFG yield the same output.
6. *Injecting opaque predicates.* Opaque predicates are injected that only evaluate correctly if the pointer to the obfuscation graph points to the correct location during the execution.
7. *Generating dead code.* Dead basic blocks added during the insertion of opaque predicates are filled with artificially created code.
8. *Transforming basic blocks.* The basic blocks themselves are transformed to obfuscate the method's original code.

In the following, the eight steps are described in detail.

Adding Properties. In order to provide a diverse range of opaque predicates for the same node, the nodes should either have a large number of properties or a property which allows a wide range of different states. Note that all nodes in the obfuscation graph have to implement the *same* properties, which may be uncommon for a set of entities in non-obfuscated applications. Therefore, the obfuscator adds a set of random properties to all possible nodes of the obfuscation graph (*i. e.*, to all classes, as a node is an object of a class). However, the random properties use *different* states.

For our obfuscation approach, a *property* can be anything that can be added to all nodes of the obfuscation graph and can hold different *states*, so that boolean expressions for opaque predicates can be built. For example, common attributes or metadata of a class, like implemented interfaces, can be used. The state of an interface would be a boolean variable indicating whether the class implements the interface.

Generating the Obfuscation Graph. The obfuscation graph is embedded into the class that contains the method(s) that should be protected. If multiple methods of the same class should be protected, the same obfuscation graph can be used multiple times. The nodes of the graph consist of objects of different classes of the target program. Hence, every node is related to a specific class of the program and therefore has different states for the added properties. The graph is a tree-like graph structure where the leaf nodes have back-edges to the root of the "tree" (semi-cyclic structure).

The structure of the obfuscation graph allows traversal on multiple paths. The obfuscator chooses random paths through the obfuscation graph and declares them to be *vpaths* (as in *valid paths*). The number of vpaths is given by the user. An example for an obfuscation graph is shown in Figure 2. Classes are randomly assigned to the nodes of the graph. The property states of the nodes on the vpaths are later used to build opaque predicates.

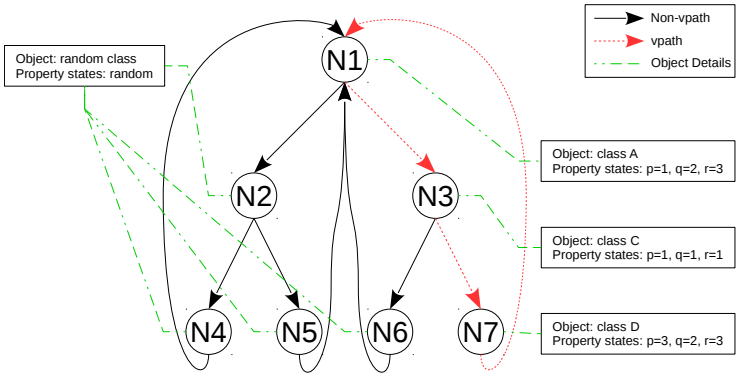

Fig. 2. An example obfuscation graph with one vpath highlighted in red. All classes for the nodes are picked randomly by the obfuscator. The classes and properties that are used for the nodes on the vpath are used to build opaque predicates.

The obfuscation graph is parametrized by its *depth* and *dimension*. The depth specifies the maximum length of a path whereas the dimension specifies the number of children of each node. These parameters can be chosen arbitrarily and determine the obfuscation graph's layout. A detailed evaluation about the effect of chosen parameters is given in Section 6.1.

Adding Initialization Code. Because the opaque predicates use properties of the nodes on the vpaths, each method to protect needs a pointer into the obfuscation graph. In order to be consistent between executions, the pointer has to point to the same starting point each time. Therefore, in the beginning of the method, the pointer is reset to the root node of the graph. This pointer realizes the link between executed basic blocks and the nodes in the obfuscation graph.

Obviously, a single vpath can be easily monitored by an adversary using dynamic analysis. Thus, at least *two* distinct vpaths have to exist in the graph. Probabilistic control flow can then be ensured by letting the obfuscated method determine randomly at runtime which vpath is used. Therefore, a *vpath state* is added to each method which determines the vpath used in current transition. It is initialized randomly in the beginning of the method at runtime.

Linking Basic Blocks. The nodes on the vpaths are linked to basic blocks in the CFG. Detailed information about the links are used later in the obfuscation process to transform the control flow of the method and to build opaque predicates (*e. g.*, the properties used to construct the opaque predicates). This information is only needed during the obfuscation process. During execution of the method, only the states of the properties are used with the help of opaque predicates to position the pointer into the obfuscation graph. The detailed information is merely kept at obfuscation time.

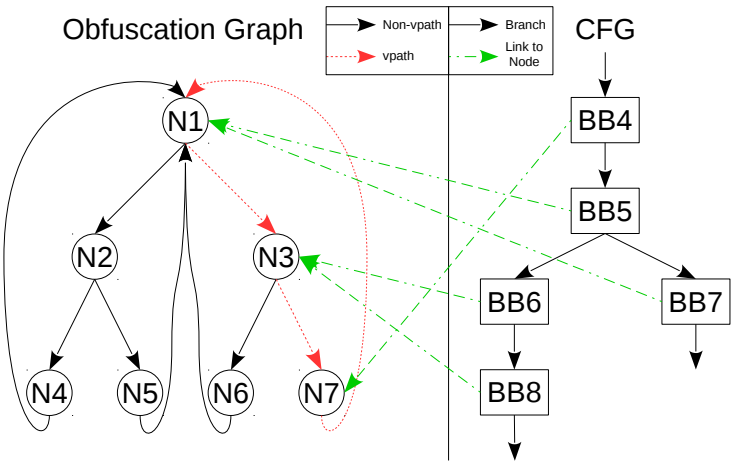

Fig. 3. An example relation between the *obfuscation graph* and the method's control flow. On the right side, a part of the control flow graph is shown. On the left side, the obfuscation graph is shown, where the vpath is highlighted in red. The relation between the nodes of the vpath and the basic blocks is highlighted in green.

An example relation of the obfuscation graph and the CFG of the method to protect is shown in Figure 3. The obfuscator links the first basic block of the CFG to the root node of the obfuscation graph (where the *first block* is the one executed first once the method is called). This is the initial position of the pointer into the graph, which is set by the initialization code added previously. The algorithm then iterates over all remaining basic blocks of the CFG and links each basic block to a node on the vpath of the obfuscation graph. During this process, the obfuscator checks for each basic block which node the preceding block is linked to. It then decides randomly to link the current processed basic block to the same node or to the next node on the vpath. This is done for each vpath the obfuscation graph possesses. Hence, each basic block has a link to one node of each vpath. The algorithm terminates when all basic blocks are linked to a node of the obfuscation graph.

Transforming Control Flow. The outgoing branches of each basic block are processed exactly once. In the following, we describe the control flow transformation process on the basis of the example shown in Figure 4:

1. Each basic block has a link to one node in every vpath. The *vpath state* (introduced to the protected method while adding the initialization code) determines which of the vpaths is currently active during execution. In order to divert the control flow depending on the currently used vpath, logic must be added that switches the control flow accordingly. Hence, the obfuscator replaces the branch of basic block A to B with one branch for every

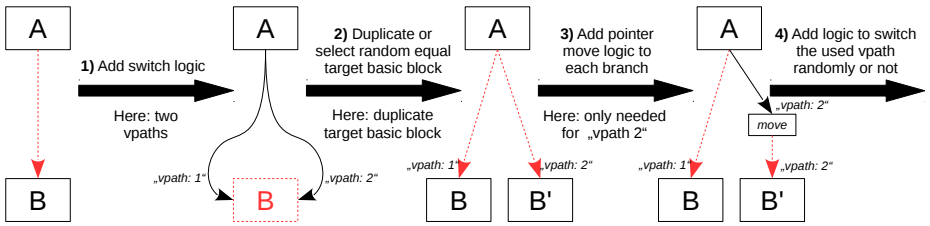

Fig. 4. The control flow transformation process operating on two consecutive basic blocks A and B. The target of the transformation is highlighted in red. The caption "vpath: X" denotes the control flow path corresponding to the respective vpath in the obfuscation graph.

existing vpath (in this example there are two vpaths). At runtime, the branch corresponding to the vpath state is taken.

2. In order to avoid all of these new branches having the same target basic block, the obfuscator either duplicates the target basic block or randomly chooses a semantically equivalent basic block. The list of semantically equivalent basic blocks consists of the target basic block itself and all duplicates of this basic block. In this example, the basic block *B* is duplicated and the new basic block *B'* is executed when *vpath 2* is currently active.
3. The source basic block of a branch and the target basic block may be linked to different nodes on the vpath. Hence, the pointer into the obfuscation graph has to be moved from the node the source basic block is linked to to the node the target basic block is linked to (compare Figure 3). As depicted in our example, basic block *B* is linked to the same node on *vpath 1* as basic block *A*, but basic block *B'* is not linked to the same node on *vpath 2* as *A*. Thus, a *move* block has to be inserted in between *A* and *B'*. Said block moves the pointer into the obfuscation graph to point to the node *B'* is linked to.
4. The current approach would not yield probabilistic control flow at all, as the *vpath state* is only set once in the initialization code of a method. Hence, for each outgoing branch of a basic block, logic may be added (determined during the obfuscation process) that may *switch* the vpath the method currently follows. The switching decision is made at runtime and at random. If switching occurs, the pointer into the graph has to be moved according to the chosen vpath.

Injecting Opaque Predicates. In this step of the obfuscation process, the obfuscator adds opaque predicates to the method that should be protected. For each basic block, the obfuscator randomly decides whether to inject an opaque predicate into the incoming branch. If an opaque predicate is injected, the obfuscator randomly decides to either create a true, false, or random opaque predicate. For the true and false opaque predicates, the never taken branch points to a newly created basic block that is marked as *dead*.

During the execution, the method's pointer into the obfuscation graph has to point to the exact node in the active vpath that is linked to the currently executed basic block. For each opaque predicate, the properties that are given by this node are used for its boolean expression. For example, with the obfuscation graph in Figure 2, the obfuscator can build a true opaque predicate for a basic block that is linked to node *N1* with the boolean expression $q == 2$. Note that this boolean expression is not unique to this node in the obfuscation graph, since it is also fulfilled by node *N7* (and probably by other nodes that do not reside on the vpath). This design decision was made to ensure that an attacker is not able to distinctively connect the opaque predicate to a node in the obfuscation graph. Even if the focus of our approach lies on dynamic analysis, the obfuscation scheme should withstand a shallow static analysis.

Furthermore, true and false opaque predicates are deterministic and do not contribute to the probabilism of the control flow. But since the attacker is allowed to conduct a manual dynamic analysis and change the program state during the execution, it adds a tamper proofing mechanism: if the attacker changes the pointer to the obfuscation graph or the obfuscation graph itself in order to affect execution, one of the following opaque predicates would divert the control flow and with a high probability crash the program. This is an advantage over a solely use of random opaque predicates to create probabilistic control flow.

Generating Dead Code. Basic blocks marked as *dead* are filled with artificially generated code. During this process the obfuscator randomly chooses the terminating instruction (called *exit*) of the dead basic block. If the chosen exit is a branch, the target can either be an arbitrary (existing) basic block in the CFG or a new dead basic block. If the target is a new dead basic block, the process is repeated. Otherwise, if the target is an existing basic block, the interconnectivity of the method's CFG is increased.

Transforming Basic Blocks. The transformation of basic blocks is necessary because the algorithm duplicated basic blocks during the control flow transformation step. If no transformation was applied, a pattern matching of basic blocks could be sufficient to detect the always taken branch of an opaque predicate.

In order to make semantically equivalent blocks harder to detect, the obfuscator employs standard obfuscation techniques [4]. We focus on those affecting control flow (like splitting blocks or outsourcing the last instructions to a common block for a subset of blocks), but other techniques can be applied as well. This includes instruction re-ordering, replacement of instruction sequences with equal ones, or usage of opaque expressions.

5 Implementation

Our prototype obfuscator is written in C# and targets .NET programs. It uses the *CCI Metadata* libraries [11] in order to transform the target program. For now,

the prototype of our obfuscation scheme operates on the bytecode of individual methods a user wishes to protect. In general, however, the approach is not limited to bytecode or methods only (or managed code programming languages). As mentioned in Section 4, the user chooses the method(s) he wants to protect. Note that typically only a very small number of methods in a given software project contain sensitive and valuable information that need to be protected.

All random numbers that are required during the obfuscation process are fetched from the same pseudo random number generator (PRNG). Hence, the seed of the PRNG can be used as a key for the obfuscation. This means the same seed used for the same target method results in the same obfuscated output. In the following, we describe specifics of our implementation.

5.1 Probabilistic Control Flow

The vpath through the obfuscation graph that is used for the current run is randomly determined during execution of the protected method. This randomness is used to implement non-deterministic control flow. We stress that these random numbers are created during the execution of the obfuscated method and not during the obfuscation process.

In our prototype implementation, the random number generator of the .NET *System* namespace is used. This implementation is sufficient for our proof-of-concept tool, but not for a real-world application. An attacker can potentially determine the points in the control flow which generates random numbers and replace them with fixed values. A detailed discussion about the random number generation during the execution of the obfuscated method is given in Section 7.

6 Evaluation

In this section, the prototype of our proposed obfuscation technique is evaluated. Since it is hard to evaluate obfuscation techniques in general, we do our best to evaluate it as thoroughly as possible using the four aspects proposed by Collberg et al. [5]: *cost*, *resilience*, *potency*, and *stealth*. *Cost* gives a measurement of the time and space overhead that is induced by the obfuscation technique. *Resilience* measures how well the protected program resists deobfuscation attempts. *Potency* measures how complex the program has become after the obfuscation process. *Stealth* measures how well the obfuscation blends into the original program.

Since our obfuscation is parametrized, we evaluate the effect of the parameters on the obfuscation first. Afterwards, cost, resilience, potency, and stealth are measured.

6.1 Obfuscator Parameters

The obfuscation graph is the only component of the obfuscation scheme that is memory dependent. Its size is mainly characterized by its *depth* and *dimension*.

Table 1. Relation between the number of vpaths and the size of the obfuscated method.

vpaths	Basic Blocks	Growth Factor	Branches	Growth Factor
4	2,520	504	3,059	611.8
5	5,963	1192.6	7,272	1454.4
6	15,418	3083.6	18,804	3760.8
7	26,215	5243	31,848	6369.6

Table 2. Size of the obfuscation graph and its dependency to the graph's depth and dimension.

Depth	Dim.	|Nodes|	Depth	Dim.	|Nodes|	Depth	Dim.	|Nodes|
6	4	1,365	7	4	5,461	8	4	21,845
6	5	3,906	7	5	19,531	8	5	97,656
6	6	9,331	7	6	55,987	8	6	335,923

Each node of the graph is represented by an object of a class in the target program and incurs an overhead dependent on the classes that are instantiated. Table 2 shows the size of the obfuscation graph for a range of parameters.

The length of the *vpath* is determined by the depth of the obfuscation graph. The number of vpaths affects the number of possible control flows of the method for the same input and thus influences the method's size as well. The effect of multiple possible control flows is further evaluated in Section 6.3. Table 1 shows the outcome of the obfuscation process for different numbers of vpaths for the same example method. The original method's CFG consists of five basic blocks and five edges. As evident from the table, the growth of the method's size proceeds exponentially.

While larger values for the parameters yield better protection levels, one has to weigh up the desired protection level with penalties in terms of size and speed. These penalties are evaluated in detail in Section 6.2.

6.2 Measuring Cost

In order to evaluate the cost of the obfuscation scheme on the program, we measure its performance, file size, and memory consumption during execution. These values are compared to the execution of the original, unobfuscated program. The tests were run on an Intel Core i7 870 CPU with 2.93 GHz using Windows 8.1 as operating system (OS). We set the number of *vpaths* through the obfuscation graph to six, the *depth* of the obfuscation graph to seven, and the *dimension* of the obfuscation graph to five. The chosen numbers provide a balance between the penalty introduced by the obfuscation scheme and the protection level that is provided, as described in Section 6.1. Since obfuscation introduces a performance overhead and is therefore usually only used to protect important parts of the program, we evaluate our approach only on the implementation of certain algorithms (representative of any intellectual property one wishes to protect). Because of its nested loop structure and variable input length,

we deem the SHA-256 hash computation as best suited to represent a worst case for our obfuscation scheme in terms of performance penalties. The nested loop structure increases the effect of the probabilistic control flow and therefore slows down the computation. In the following, we describe this test case in detail. The evaluation of additional test cases can be found in our technical report [18].

Size. To quantify the impact of our obfuscation scheme on the file size, we measure the file size in bytes. In our setting, the size of the original binary is 12,288 bytes and the obfuscated binary has a file size of 7,666,688 bytes. This implies that the obfuscated binary is about 624 times larger than the original binary. This result is similar to the other test cases in the corresponding technical report [18]. Note that, as discussed in Section 6.1, the size of the obfuscated binary highly depends on the parameters chosen for the obfuscator. In order to ensure a variety of possible control flows, the obfuscator has to clone the basic blocks of the target method multiple times. Therefore, our obfuscation scheme also increases the size of the target method multiple times. We stress that the growth of the size is dependent on the target method and not on the entire program. A large program has the same growth as a small program if they implement the same method that is the target of the obfuscation. Nevertheless, due to the resources available to modern devices, we see this growth as acceptable.

Performance. The performance is measured by calculating the SHA-256 hash of a 10 MB file. In order to compensate for outliers, we repeat the calculation 1000 times and calculate the average time. We take two different timings. First, the time needed for the creation of an object of the obfuscated class, and second the time needed for the actual computation of the hash is measured. During the creation of the object itself, the obfuscation graph is built by the constructor of the class. The creation of the obfuscation graph impacts the overall performance depending on the parameters specified by the user. Therefore, we also have to take timings for the creation and not only for the actual computation. Timings are measured with a resolution of 1 ms.

The original binary takes less than 1 ms for object creation. The obfuscated binary takes 3925 ms to create the object (and therefore to build the obfuscation graph). The calculation of the hash is performed in 785 ms by the original binary, whereas 5658 ms are needed by the obfuscated binary. While the obfuscated SHA-256 algorithm takes around 7 times longer to perform the same calculation, we stress that this case constitutes a worst case scenario for our obfuscation scheme in terms of performance. The other tested algorithms in our technical report [18] need roughly the same time to create the object, but only need around 1.6 times longer to perform the same calculation. Again, these values are dependent on the parameters of the obfuscation graph. While parameters exists for which obfuscation graph creation consumes less time, the protection level for the obfuscated method is lowered as well. Additionally, algorithms that are usually protected with obfuscation in real-world applications are sparsely performed during the execution of a program. Therefore, we regard the introduced performance penalty as acceptable.

Memory. The only memory dependent component of the proposed obfuscation technique is the obfuscation graph. Therefore, the memory consumption of the graph is measured after the object of the protected class is created in the program. The parameters yield an obfuscation graph with 19,531 nodes. The original program consumes 1,480 kB of memory after the object is created. The protected program needs 28,852 kB after the target object is allocated. Therefore, the obfuscation graph needs about 27,372 kB for the used parameters. This is similar to the memory consumption of the other test cases in our technical report [18]. Note that the memory required for one obfuscation graph is constant. Larger applications embedding the same obfuscation graph will face the same memory requirements. Having the resources available on today's devices in mind, we believe the impact on memory consumption to be tolerable.

6.3 Measuring Resilience

Resilience measures the resistance of the obfuscation scheme against deobfuscation attempts. Since we focus on thwarting dynamic analyses, we measure the resilience of our obfuscation scheme by quantifying the probabilistic control flow. Therefore, we trace the execution of an obfuscated method with the same input values and compare the similarity of these traces. To this end, we generate a graph from the traced basic blocks in the obfuscated method and compute the *graph-edit distance* between two execution traces using the algorithm proposed by Hu et al. [12]. The graph-edit distance yields the number of edits needed to transform one graph into another graph. Edits are node insertions/deletions and edge insertions/deletions.

We follow the proposal of Chan et al. [2] and normalize the graph-edit distance such that it computes a similarity score using the following formula:

$$\text{similarity}(G_1, G_2) = 1 - \left(\frac{\text{graph-edit distance}}{|G_1| + |G_2|} \right),$$

where the size of the graph G_i is given by the total number of nodes and edges and is denoted by $|G_i|$. The output of the similarity function is a value between 0.0 and 1.0. A result of 1.0 means that the two graphs are identical, whereas a result of 0.0 means they are completely different.

Results. As test case we use our running example, the SHA-256 hash computation. We generated 100 traces by executing the program 100 times in a row with the same input. Since the graph-edit distance calculation is NP-hard in general [26], we have to choose an input size that creates traces with graph dimensions that are still comparable. To this end, we used 100 bytes of random data. Since the SHA-256 hash computation operates on blocks of 512 bits, the algorithm runs through multiple iterations until it terminates. As obfuscation parameters we use the settings evaluated in Section 6.2.

In total, we calculated 4,950 graph comparisons (as graph comparison is commutative). The greatest similarity of two traces was 88.45%. The smallest

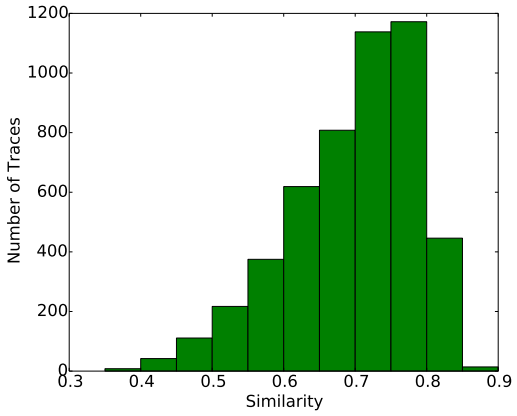

Fig. 5. The 4,950 similarity values of the traces displayed as a histogram. The bin size amounts to 0.05. The smallest similarity was 0.35 and the greatest 0.88. The majority of the values have a similarity of under 0.75.

similarity was 35.29%, while the average of all similarities is 69.65%. An overview of the similarity between the traces is given in Figure 5 as histogram. As can be seen, most of the similarity values are near the calculated average value in the range of 60% to 75%.

The smallest trace regarding the number of unique basic blocks visited 359 unique basic blocks and took 367 unique branches. The largest trace reached 1,183 unique basic blocks and took 1,255 unique branches. On average, 753 unique basic blocks were visited and 793 unique branches were taken by the traces. The number of all visited unique basic blocks and taken unique branches is given in Figure 6. As evident from the figure, the number of visited unique basic blocks and taken unique branches correlate. If more unique basic blocks were executed, more unique branches were used. But still, the number of basic blocks and branches vary greatly between single executions. The size of the traces of our other test cases is provided in the corresponding technical report [18].

These results show that multiple executions for the same input values do not even once have the same execution path. This effectively hinders deobfuscation approaches working on multiple traces, such as state-of-the-art deobfuscation methods like [25]. Also, a manual analysis using breakpoints is rendered unreliable in presence of the probabilistic control flow, as we explain in Section 7.

6.4 Measuring Potency

Potency measures how complex and confusing the program becomes after obfuscation. In order to evaluate the potency of our obfuscation scheme regarding dynamic analysis, we measure the differences between the original and an obfuscated control flow. Therefore, we recorded an execution trace for the original

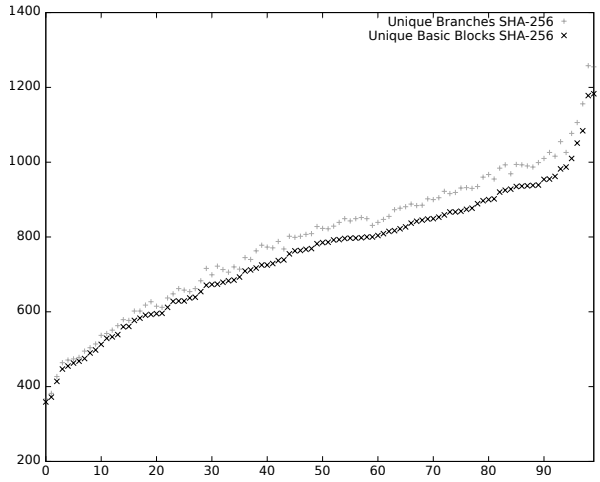

Fig. 6. The number of unique basic blocks and branches each trace used ordered by the number of reached basic blocks. The gray + dots depict the used unique branches and the black x dots show the visited unique basic blocks. On the x-axis the trace number is given. On the y-axis the number of unique basic blocks/unique branches are given.

and obfuscated program with the same input. During the obfuscation process, all semantically equivalent basic blocks were labeled in order to recognize them in the obfuscated CFG. Note that this information is not available for an adversary trying to analyze the obfuscated method.

In order to quantify the *utilization* of the different semantically equivalent basic blocks we visited with respect to all available semantically equivalent basic blocks and the number of executions, we make the following case distinction:

$$\text{utilization} = \begin{cases} \frac{|diff|}{|exec|}, & \text{if } |exec| < |avail| \\ \frac{|diff|}{|avail|}, & otherwise \end{cases},$$

where $|exec|$ gives the number of times one of the semantically equivalent basic blocks were visited, $|avail|$ gives the number of available semantically equivalent basic blocks, and $|diff|$ gives the number of visited different semantically equivalent basic blocks. This way we can differentiate between cases where the total number of visited semantically equivalent basic blocks is lower than the available number of semantically equivalent basic blocks and vice versa. Consider for example a case where only one of the available semantically equivalent basic blocks is executed. If this is the case during multiple iterations of a loop, its utilization of the available semantically equivalent basic blocks is obviously not optimal because control flow visits only this available basic block multiple times. On the other hand, utilization is good if the code contains no loop and control flow visits only one of the semantically equivalent basic blocks during the execution only one single time. Therefore, we have to differentiate.

Table 3. The results of the comparison of the obfuscated method trace with the trace of the original method for the same input (ID = ID for semantically equivalent basic blocks, $|avail|$ = number of available semantically equivalent basic blocks, $|exec|$ = total number of times one of the semantically equivalent basic blocks were visited, $|diff|$ = number of different semantically equivalent basic blocks executed, `Util` = utilization of the reached different semantically equivalent basic blocks with respect to available semantically equivalent basic blocks and the total number of executions in percent).

ID	**0**	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**	**12**	**13**	**14**	**15**	**16**	**Total**
$\|avail\|$	9	43	40	30	35	24	22	20	29	18	22	31	25	22	43	23	33	469
$\|exec\|$	1	20	1	19	3	1	2	1	34	2	32	98	2	96	130	2	128	572
$\|diff\|$	1	10	1	8	3	1	2	1	15	1	10	4	1	4	24	2	20	108
`Util`	100	50	100	42.1	100	100	100	100	51.7	50	45.5	12.9	50	18.2	55.8	100	60.6	71

Results. As input data we used 100 bytes of random data and as obfuscation parameters we use the settings evaluated in Section 6.2. We recorded a trace by executing the obfuscated and original program with the same input. The resulting traces were compared with respect to their executed basic blocks.

The obfuscator cloned the basic blocks of the original method multiple times during the obfuscation process. Remember that the decision to clone a basic block is made randomly during the obfuscation process. The minimum number of semantically equivalent basic blocks in the obfuscated method amounts to 9 and the maximum number to 43. On average, the control flow has 27 different possibilities per basic block to exhibit the same behavior.

During the execution of the obfuscated method, the control flow has visited 572 relevant basic blocks that contribute to the calculation of the result. These basic blocks consist of the basic blocks of the original method and transformed copies of these original basic blocks. The utilization of the available semantically equivalent basic blocks ranges from 12.9% to 100%. In total, 71% of the available semantically equivalent basic blocks were utilized during the execution of the obfuscated method. The results for our test case are shown in Table 3. All test cases in our technical report [18] have similar results.

The results show that an execution of the obfuscated method uses a variety of different but semantically equivalent basic blocks to compute its result. Hence, the number of basic blocks that are actually involved in the computation has been increased by our approach and with it the complexity of the control flow.

6.5 Measuring Stealth

Stealth measures the difficulty for an adversary to determine if the given method is obfuscated, *i. e.*, how well the obfuscated entity fits in legitimate code. Although stealth is not an objective of our approach, we evaluate it for the sake of completeness. Recently published obfuscation papers measure this aspect based on the distribution of instructions [3, 19, 24]. However, as Collberg et al. [5] describe it, stealth is a *context-sensitive* metric. Hence, instead of pursuing a

static approach for evaluating stealth, we consider the *dynamic* behavior of the obfuscated program. This fits our general focus on dynamic analysis.

However, since our approach is, by design, supposed to yield different execution traces for the same input, *stealth* is inherently hard. An adversary only has to execute the program two times with the same input and compare the recorded execution traces. If they differ, the adversary can conclude that the program is most likely protected by our obfuscation approach.

7 Discussion

In the following, we discuss potential limitations of our approach.

Dynamic Analysis. Our approach aims to transform methods such that multiple traces of the same function using the same inputs differ, which implies that dynamic deobfuscation approaches are hampered [7, 21]. Furthermore, this is done to thwart dynamic analyses operating on multiple executions (like [25]). For example, *manual* dynamic analysis of the obfuscated method is hindered by probabilistic control flow: an adversary observing the control flow at some fixed point during execution of the method cannot depend on the program reaching the exactly same point during a following run. Hence, pausing execution using breakpoints is rendered unreliable in presence of our obfuscation approach.

Single Trace Analysis. If an adversary knows that our obfuscation scheme is used, the best way to attack it is by resorting to work on a single execution trace. Since the goal of probabilistic control flow is to make dynamic analyses based on multiple traces harder, deobfuscation methods operating on only *one* trace are only affected if at least one loop is present. In this case, our scheme increases the size of the recorded trace because the obfuscator clones basic blocks in order to have multiple possible control flows to choose from. As shown in Section 6.4, the execution of multiple iterations of a loop results in different semantically equivalent basic blocks that are reached. Algorithms processing the recorded trace *dismiss* basic blocks that do not affect the outcome of the method [7, 21, 25]. Since the visited semantically equivalent basic blocks of the probabilistic control flow affect the outcome of the method, they can not be dismissed. As a result, subsequent analysis of the recorded trace is more complicated due to our obfuscation scheme. As future work, we propose to integrate the use of the obfuscation graph into the calculations of the protected method. This way it gets harder to dismiss instructions based on their usage of the obfuscation graph.

Furthermore, deobfuscation methods operating on only one trace do not perform as good in terms of *code coverage* compared to those using multiple execution paths. This poses a problem for an adversary who wants to analyze multiple execution paths in an algorithmic manner in order to understand the obfuscated program better. Often, *multi-path exploration* techniques are considered when tackling this problem [21, 25]. This is where our approach proves useful: It introduces a variety of valid, but distinct control flows and adds probabilism. For

the adversary, it is hard to distinguish whether a branch was taken due to probabilistic control flow or because the function was run with different input. In order to improve this aspect, we currently work on extending our approach by merging the semantics of multiple methods into one method. The semantic that is actually executed when the method is called is then determined with the help of the obfuscation graph and opaque predicates. Therefore, the same method can have multiple semantics and, depending on the vpath that is used, the correct semantic of the method is chosen.

Probabilistic Control Flow. An important component of our proposed approach is the obfuscation graph with its vpaths. The vpaths are used to select the current control flow through the obfuscated method and therefore to introduce probabilistic control flow. Which vpath is to be used is decided by a random value. In our prototype implementation, the used vpath is merely chosen using the PRNG as provided by the .NET *System* namespace. This implementation is obviously vulnerable, as the call to the PRNG could be replaced by the usage of fixed values. As a result, the probabilistic control flow is then merely reduced to a deterministic one.

A straightforward approach to make the random number generation more resilient is not to use any *external* PRNG. Instead, one could build a PRNG into the obfuscated method itself and replace the calls to the external PRNG with code sequences that generate random numbers. This way, the random number generation is harder to pinpoint by an adversary because the code that generates the random number is concealed by the code of the obfuscated method. The obfuscator is not limited to build only one PRNG into the obfuscated method but could inject multiple ones to make it even harder to find the code sequences that generate random numbers. Furthermore, the random number generation can be protected by additional layers of obfuscation like translating the obfuscated method to custom bytecode [1, 17, 22].

However, even this construct suffers from the problem that it needs an initial random seed to create different control flows every time it is executed. If an adversary is able to set this initial random seed to a fixed value, the PRNG in the obfuscated method generates the same sequence of random numbers every time the program is executed. Even if the user input influences the calculation of the random numbers, the program would only have different traces for different inputs (which still hampers analysis of the program with different inputs, but allows debugging of the function with the same input). This circumstance poses the greatest limitation of our current implementation of the proposed obfuscation scheme. However, due to their huge number, it is not easy in practice to detect every single state that is fetched by a program from the OS or to set every internal state of an OS every time to the exact same value in order to fix the seed. One approach to circumvent fixed OS states would be using non-deterministic sources like the intentional use of race-conditions. For future work, we propose to develop methods to conceal the fetching of external states for the random number generation.

8 Related Work

The basic technique our approach is based on is presented in a paper by Collberg et al. [5]. They propose a method to create opaque constructs based on objects and pointer aliases. Also, they suggest a directed graph as concrete data type. However, their approach is mainly concerned with the creation of cheap, stealthy and resilient opaque constructs. We specifically extend this approach and focus on the different paths we can insert into a target using their construct. This stems from the insight that while their technique efficiently makes static analysis harder, the traces obtained using dynamic analyses are very much the same. This, in turn, helps in determining the concrete value of an opaque predicate and might allow to partly reconstruct the control flow of the program.

Wang et al. describe a technique to obfuscate a target program using control flow transformations as well [23]. They transform a method's CFG in such a way that a new basic block in the beginning of the method decides which original basic block is executed next. These control flow decisions are made based on a state variable which gets updated after every basic block. Similar to the approach of Collberg et al., they transform the control flow analysis problem into a data flow analysis problem. However, their approach also merely aims to make static analysis of an obfuscated program harder.

More recent work focuses on deobfuscation of obfuscated programs [7,21,25]. All of them have in common that they are based on dynamic analysis. Traces of the program's execution are recorded and subsequently used to remove the applied obfuscation schemes. Approaches working on multiple traces in order to tackle the code coverage problem [16] of dynamic analysis are challenged by the probabilistic control flow introduced by our technique.

The recent work of Crane et al. also make use of probabilistic control flow [8]. It enables them to thwart cache side-channel attacks. To this end, they clone program fragments and transform the clone in order to avoid making an exact copy. A stub is used to decide randomly if the clone or the original fragment is executed. Because an attacker has no knowledge about which was executed, it hampers cache side-channel attacks. Additionally, Davi et al. [9] use probabilistic control flow in combination with memory randomization in order to prevent conventional return-oriented programming (ROP) and JIT (just-in-time)-ROP attacks. To this end, they clone and diversify the code that is loaded into memory. Whenever a function is called, their system randomly decides if the original or cloned function is executed. Once the executed function returns, the system checks if execution shall continue at the normal or cloned version of the function caller by adding an offset to the return address. Therefore, an attacker is not able to precisely predict where execution will resume and cannot reliably perform an attack.

9 Conclusion

In this paper, we introduce a novel approach to obfuscate software, including, but not limited to, those written in managed code programming languages. The

proposed scheme is based on a construct introduced by Collberg et al. [5]. However, instead of focusing on protecting the program against static analysis, we introduce a scheme achieving probabilistic control flow, aiming to make dynamic analysis harder. This is achieved by embedding an obfuscation graph containing multiple vpaths. Based on these paths, opaque predicates are constructed and added to the target method. Consequently, control flow may take different paths exhibiting the same observable semantics.

We have implemented a prototype obfuscator for .NET applications and evaluated it using multiple programs. The experiments have shown that the obfuscated methods do not exhibit the same execution trace after executing it 100 times in a row with the same input. Inevitably, this comes with a significant performance and memory penalty. Resilience against dynamic analyses thus has to be weighed up with constraints on time and space. We are confident that the overhead is still acceptable to protect sensitive parts or proprietary algorithms of a given program. Since we believe our obfuscation approach provides a new strategy for tackling dynamic analysis and hence a building block for future research, we are making our obfuscation tool available to the research community.

References

1. Anckaert, B., Jakubowski, M., Venkatesan, R.: Proteus: Virtualization for Diversified Tamper-Resistance. In: Proceedings of the ACM workshop on Digital rights management (2006)
2. Chan, P.P., Collberg, C.: A Method to Evaluate CFG Comparison Algorithms. In: International Conference on Quality Software (QSIC) (2014)
3. Chen, H., Yuan, L., Wu, X., Zang, B., Huang, B., Yew, P.c.: Control Flow Obfuscation with Information Flow Tracking. In: Annual IEEE/ACM International Symposium on Microarchitecture (2009)
4. Collberg, C., Thomborson, C., Low, D.: A Taxonomy of Obfuscating Transformations. Tech. rep., Department of Computer Science, The University of Auckland, New Zealand (1997)
5. Collberg, C., Thomborson, C., Low, D.: Manufacturing Cheap, Resilient, and Stealthy Opaque Constructs. In: ACM Symposium on Principles of Programming Languages (POPL) (1998)
6. Collberg, Christian: The Tigress C Diversifier/Obfuscator. `http://tigress.cs.arizona.edu`
7. Coogan, K., Lu, G., Debray, S.: Deobfuscation of Virtualization-obfuscated Software: a Semantics-based Approach. In: ACM Conference on Computer and Communications Security (CCS) (2011)
8. Crane, S., Homescu, A., Brunthaler, S., Larsen, P., Franz, M.: Thwarting Cache Side-Channel Attacks Through Dynamic Software Diversity. In: Symposium on Network and Distributed System Security (NDSS) (2015)
9. Davi, L., Liebchen, C., Sadeghi, A.R., Snow, K.Z., Monrose, F.: Isomeron: Code Randomization Resilient to (Just-In-Time) Return-Oriented Programming. In: Symposium on Network and Distributed System Security (NDSS) (2015)
10. Fang, H., Wu, Y., Wang, S., Huang, Y.: Multi-stage Binary Code Obfuscation using Improved Virtual Machine. In: Information Security (2011)

11. Guy_Smith: Common Compiler Infrastructure: Metadata API. https://ccimetadata.codeplex.com/
12. Hu, X., Chiueh, T.c., Shin, K.G.: Large-scale Malware Indexing Using Function-Call Graphs. In: ACM Conference on Computer and Communications Security (CCS) (2009)
13. Junod, Pascal: Obfuscator-LLVM. https://github.com/obfuscator-llvm/obfuscator/wiki
14. Kushner, David: Steamed: Valve Software Battles Video-game Cheaters. http://spectrum.ieee.org/consumer-electronics/gaming/steamed-valve-software-battles-videogame-cheaters
15. Lee, B., Kim, Y., Kim, J.: binOb+: A Framework for Potent and Stealthy Binary Obfuscation. In: ACM Symposium on Information, Computer and Communications Security (ASIACCS) (2010)
16. Moser, A., Kruegel, C., Kirda, E.: Exploring Multiple Execution Paths for Malware Analysis. In: IEEE Symposium on Security and Privacy (S&P) (2007)
17. Oreans Technologies: Code Virtualizer: Total Obfuscation against Reverse Engineering. http://oreans.com/codevirtualizer.php
18. Pawlowski, A., Contag, M., Holz, T.: Probfuscation: An Obfuscation Approach using Probabilistic Control Flows. In: Technical Report TR-HGI-2016-002, Ruhr University Bochum (2016)
19. Popov, I.V., Debray, S.K., Andrews, G.R.: Binary Obfuscation Using Signals. In: USENIX Security Symposium (2007)
20. Ramalingam, G.: The Undecidability of Aliasing. ACM Transactions on Programming Languages and Systems (TOPLAS) (1994)
21. Sharif, M., Lanzi, A., Giffin, J., Lee, W.: Automatic Reverse Engineering of Malware Emulators. In: IEEE Symposium on Security and Privacy (S&P) (2009)
22. VMProtect Software: VMProtect: Software protection against reversing and cracking. http://vmpsoft.com/
23. Wang, C., Davidson, J., Hill, J., Knight, J.: Protection of Software-Based Survivability Mechanisms. In: International Conference on Dependable Systems and Networks, 2001. DSN 2001. (2001)
24. Wang, P., Wang, S., Ming, J., Jiang, Y., Wu, D.: Translingual Obfuscation. In: IEEE European Symposium on Security and Privacy (Euro S&P) (2016)
25. Yadegari, B., Johannesmeyer, B., Whitely, B., Debray, S.: A Generic Approach to Automatic Deobfuscation of Executable Code. In: IEEE Symposium on Security and Privacy (S&P) (2015)
26. Zeng, Z., Tung, A.K., Wang, J., Feng, J., Zhou, L.: Comparing Stars: On Approximating Graph Edit Distance. In: International Conference on Very Large Data Bases (VLDB) (2009)