

Memoization Attacks and Copy Protection in Partitioned Applications

Charles W. O'Donnell¹, G. Edward Suh², Marten van Dijk¹, and Srinivas Devadas^{1†}

¹Massachusetts Institute of Technology, Cambridge, MA 02139
{cwo,marten,devadas}@mit.edu

²Cornell University, Ithaca, NY 14853
suh@csl.cornell.edu

Abstract— Application source code protection is a major concern for software architects today. Secure platforms have been proposed that protect the secrecy of application algorithms and enforce copy protection assurances. Unfortunately, these capabilities incur a sizeable performance overhead. Partitioning an application into secure and insecure regions can help diminish these overheads but invalidates guarantees of code secrecy and copy protection.

This work examines one of the problems of partitioning an application into public and private regions, the ability of an adversary to recreate those private regions. To our knowledge, it is the first to analyze this problem when considering application operation as a whole. Looking at the fundamentals of the issue, we analyze one of the simplest attacks possible, a “Memoization Attack.” We implement an efficient Memoization Attack and discuss necessary techniques that limit storage and computation consumption. Experimentation reveals that certain classes of real-world applications are vulnerable to Memoization Attacks. To protect against such an attack, we propose a set of indicator tests that enable an application designer to identify susceptible application code regions.

I. INTRODUCTION

Proprietary software architects have long been concerned with Intellectual Property (IP) protection to guard trade-secret algorithms from theft, enforce licensing agreements, and to prevent application vandalism caused by viruses and Trojan horses. However, the scale with which the Internet facilitates piracy and application “cracking” has now made IP protection a first order concern. Consequently, a number of software protection techniques have arisen, primarily based on hiding application functionality, but without strict security guarantees.

In this work we investigate state-of-the-art security systems that protect applications by partitioning code into public and private regions of execution [27][49]. Specifically, we analyze one of the most basic methods for an adversary to determine the functionality of hidden application code, a “Memoization Attack.” We have implemented a Memoization Attack, run it against a number of applications, and developed methods of identifying when an arbitrary application might be vulnerable to such an attack.

Naively, a secure computing system can perfectly protect application IP by executing all software on an impenetrable Trusted Computing Base (TCB), where only the final application results are observable. To practically achieve this one could modify a standard processor

with hardware safeguards and encryption mechanisms, encrypt the application, and only allow that processor to decrypt and re-encrypt software instructions and data. Unfortunately, whole-application encryption can inhibit the use of shared libraries, complicate upgrades and patches, and most importantly, require the use of cryptographic resources throughout all of software execution — incurring a sizable performance and power usage penalty. Since typically only a small portion of an application is considered sensitive IP, all of these problems can be mitigated by partitioning the application and only requiring the TCB for execution of the sensitive IP.

The idea of application partitioning has been proposed often as a performance-friendly mechanism for secure processors [27][49], secure co-processors [59], dongles [36], and remote servers [13] to defend against attack [2][24][35]. However, while it seems self-evident that running an entire piece of software on a TCB guarantees its IP privacy and licensing assurances, *it is not clear if these guarantees hold true for individual partitions of a partitioned application*. For example, if an adversary is able to duplicate the functionality of IP sensitive partitions by simply observing the execution of the remainder of the code, the IP protection is wholly invalidated. Therefore, this paper’s analysis of this kind of software vulnerability is of great importance to guarantee the security of partitioned applications.

In Section II, we put forth a simple adversarial model. Section III defines what a Memoization Attack is, and shows that it is the “best” possible attack an adversary can mount given our model. Section IV describes one practical and efficient implementation of this attack, and Section V describes when the attack can be effective. Using these insights, Section VI proposes heuristic metrics that can be used to identify whether a partitioned region of application code is susceptible to an Memoization Attack. Section VII discusses other work in this area and Section VIII concludes.

II. ATTACK MODEL

In this work, we restrict our focus to one of the simplest types of adversary imaginable. As will be described in Section II-D, our adversary can only *observe* the execution of a partitioned application and then attempt to reconstruct the hidden regions that are run on a TCB using that observation. More sophisticated adversaries are easy to envision, however, we feel it prudent to explore a very basic model

[†]This work is generously supported by ITRI, Taiwan.

to its fullest extent. Further, the adversarial powers we describe here can be considered necessary for a number of more complex types of adversaries.

A. TCB and Partitioned Application Model

For the sake of clarity we will only focus on one type of TCB model so that we can describe more concretely what actions an adversary can and cannot take, and what constitutes a partitioned applications. To this end, we look at physically secure processors and co-processors [27][49][59] since these represent some of the most secure methods that exist for TCB code execution. Specifically, we choose the AEGIS secure architecture [49] because of its fairly straightforward protocol for the execution of partitioned applications. The remainder of this paper and all of our experiments assume this model for a TCB.

In the AEGIS secure architecture a partitioned application is merely a combination of *private* encrypted regions and *public* unencrypted regions of code that switch back and forth during execution using two distinct processor modes. Application memory is also separated into encrypted and unencrypted regions, conceptually forming private and public divisions of data and code. The encrypted portions of code can only run in a secure mode that decrypts instructions, executes them, and protects the secrecy and integrity of any private data these instructions operate on. While executing in this mode an adversary can only observe accesses to public data, and cannot observe or modify private data or program execution. Unencrypted portions of a partitioned application run in an insecure mode, with no protection of the data the instructions operate on.

For simplicity, we assume that *procedures* and the data structures they “own” are the fundamental units of public and private division. We also assume that procedures do not maintain state across calls within encrypted regions of memory. Disallowing a procedure to maintain encrypted state between calls allows for a more clean analysis and is fairly realistic for a large number of procedures within applications

B. What an Adversary Can Observe

Fig. 1 depicts a fragment of a partitioned application while it is run on an AEGIS secure architecture. Progressing downward is an execution trace of an application as it switches from a public region of code to a private region and back. To reduce clutter, we only show reads and writes to main memory and do not show any other machine operations (such as add, etc.).

Beginning in box I a public region of code executes and performs some arbitrary procedure. Note that, this being a public procedure, all accesses to memory can only touch regions of memory that are also public. Since this is unencrypted code executing on a conventional processor, an adversary can inspect everything involved with the procedure. The procedure itself can be read to determine its

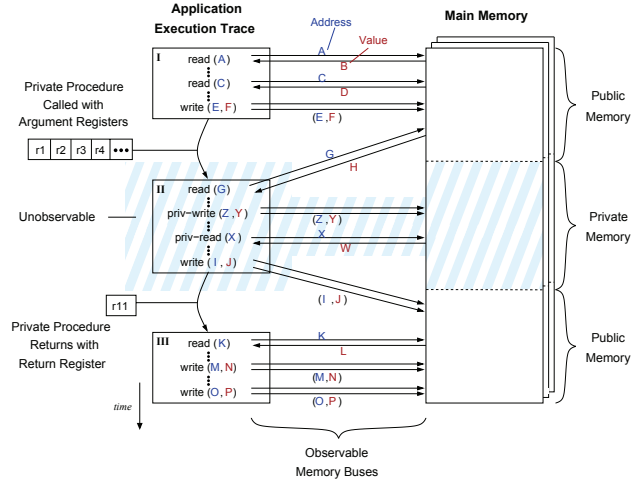


Fig. 1. Partitioned application input/output visibility.

control flow, the processor state can be examined cycle-by-cycle, and all memory requests and responses can be sniffed.

At the end of box I the procedure calls a private region of code (box II) and transfers control to the TCB to execute that private procedure. This call requires procedure arguments to be passed to the TCB, as shown by the registers “r1,” etc. in Fig. 1 (as defined by the application binary interface and including the stack and frame pointers). Similarly, once the procedure completes, a return value is also passed back from the TCB to the conventional processor. Since the private procedure was encrypted, an adversary cannot inspect the code directly to determine its control flow, nor can it examine the processor state cycle-by-cycle since it’s a TCB. Further, this TCB model *hides* any accesses to its private memory stack that the private procedure makes.

Therefore, the only information an adversary can observe relating to the private code is the arguments passed into the procedure, the return value passed back from the procedure, and any accesses to *public* memory that the private procedure makes (since public memory requests cannot be hidden by the TCB and the values within public memory are unencrypted). All three of these can be described as a collection of Address/Value (AV) pairs, where the “Address” indicates a memory address or argument register identifier, and the “Value” is the actual data being accessed. Once the private procedure returns to execution of public code (box III), the adversary can again observe everything.

C. Adversary Goals

Principally, IP secrecy and copy protection depends on preventing an adversary from discovering the contents of a partitioned application’s private code. However, it is critical to note that an adversary does not need to *exactly* determine the contents of a private region of code, *but must only reproduce a private procedure’s effect on the system*

sufficiently well as to allow the entire partitioned application to continue to function as designed. Therefore, an adversary’s most simple goal is to replace “authentic” private procedures with indistinguishable “counterfeit” procedures that can reproduce the adversary’s desired “functionality” and “utility” of the partitioned application as a whole.

Ultimately, the only functionality that matters to an adversary is the set of application outputs that result from some set of inputs he is interested in. If the set of inputs are time-dependent, then the adversary may further only be interested in reproducing functionality for a limited amount of time. To this extent, an adversary need not *understand* each private procedure, but must only duplicate its external effects. For example, assume a fragment of an application performing the power function $f(x, p) = x^p$ is made secret. If an adversary only ever *cares* about executions when $p = 3$ then his only interest is in duplicating code that performs $f(x, 3) = x^3$.

Consequently, this limited sense of duplication of functionality is exactly what we should be concerned with when analyzing possible attacks on a partitioned application. This can be formally defined as *Temporal Application Operation Equivalence* (T-AOE).

Definition 1: T-AOE($APP', APP'', \langle \Lambda \rangle, t_s, \omega$): Assume two applications APP' and APP'' begin execution at time 0 and finish execution at the same time H . During each unit of time between 0 and H , both applications are given the same, new vector of inputs Λ_t chosen from some set of many input vectors, the total available input set $\langle \Lambda \rangle$. These applications are T-AOE at some time t_s for the length of time ω if, during the period $[t_s, t_s + \omega]$, the responses or results of both applications Ψ'_t and Ψ''_t , are exactly equivalent (assuming $0 \leq t_s \leq H$ and $(t_s + \omega) \leq H$).

Given this definition, the adversary we are concerned with aspires to create a counterfeit private region of code for a specific partitioned application that maximizes T-AOE time ω (ideally, $\omega \rightarrow (H - t_s)$). This ω can be thought of as the adversary’s “time-till-failure.”

D. Adversarial Powers

Unfortunately, any realistic adversary we try to model involves a human who has some *innate prior knowledge* about the application under attack that can make the process of recreating a hidden region of code trivial. For all we know, the adversary may even be the author of the original source code for a private procedures.

Given the inability to formally capture such knowledge, we will simply treat a private procedure as a mathematical function with inputs and outputs the adversary is capable of observing. Our adversary has no understanding of the purpose of a private procedure and can only obtain knowledge of the code’s functionality by observing the procedure arguments, the public memory accesses, and the procedure output of an authentic application run on a TCB. Note, although it seems probable for a real-world attack, our ad-

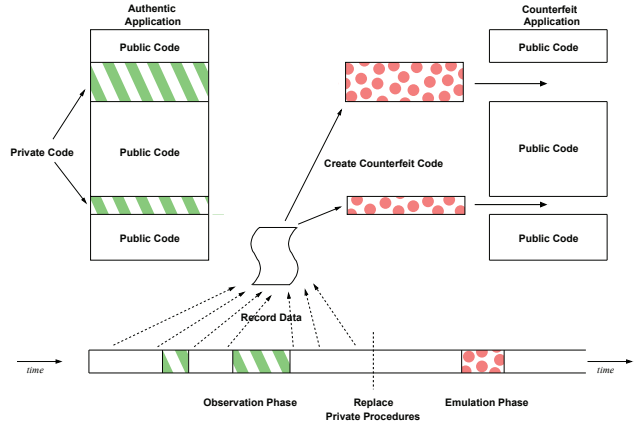


Fig. 2. Basic technique a Memoization Attack.

versary does not analyze the available public code at all to infer any “meaning” to the application. This would again prove quite troublesome to model.

To formally specify these powers, we say that for each call to a private procedure, an adversary is aware of an *input set* λ of memory reads and the procedure’s arguments, and an *output set* ψ of memory writes and the procedure’s return value. These can be thought of as a vector of data values indexed by addresses or register numbers. Combining these vectors forms a single *input/output relationship pair* (λ, ψ) . An adversary observing multiple calls to a private procedure can collect a multiple number of pairs.

III. MEMOIZATION ATTACKS

A “*Memoization Attack*” is the name we give to an attack that uses an authentic private procedure’s input/output relationship pairs to create an alternate counterfeit version. Fig. 2 shows one way a Memoization Attack can be performed. The adversary begins by running an authentic application using a TCB for some amount of time. During this time all input/output relationship pairs (λ, ψ) of the private procedures are observed and stored into a single “*Interaction Table*.” At some point the adversary stops executing the authentic application on the TCB and constructs replacement private procedures using the interaction table that was captured. He can then continue to execute the application using these counterfeit private procedures. Whenever a counterfeit procedure is called, the set of inputs λ are read, and the interaction table is searched for a match. If a match is found the counterfeit procedure returns the corresponding output set ψ , emulating the procedure, and continues execution of the application. Otherwise the application fails and terminates. The application continues running as long as calls to the counterfeit procedure are completed correctly, agreeing with our previous definition of failure under T-AOE.

To determine just how powerful this attack can be, let us first assume there exists an adversary with infinite memory and computational power, but who must also abide by the restrictions on adversarial powers discussed in Section

II. Note, while this adversary may have infinite general purpose computational power, we assume that he cannot decrypt private procedures and is restricted to the use of a real TCB to run authentic applications. Therefore what an adversary can observe from the execution of an authentic application remains the same as in a Memoization Attack since this is essentially defined by our model (although with infinite memory this adversary can save everything). Our question is then: can this omnipotent adversary mount a different type of attack that can outperform a Memoization Attack (that is, have a longer ω value for T-AOE).

Now let us assume that this adversary observes L calls to an authentic private procedure, somehow creates and inserts his own counterfeit procedure, and continues running the application. No matter how the counterfeit procedure is constructed, when the counterfeit procedure is called during emulation only one of two things can happen. If the exact set of inputs λ had been seen during the observation phase, then the adversary can simply return the set of corresponding outputs ψ that it had saved. However, if the set of inputs to the counterfeit procedure contain any new elements, then the adversary must rely on some other knowledge to decide what to output. Unfortunately, the only knowledge the adversary has is the set of input/output relationship pairs already seen. In the absence of any other knowledge, there's no reason to believe there's any correlation between prior input/output relationship pairs and the current input. Such a correlation requires a "prior" or an abstract learning model. Therefore this adversary's best option is simply to *uniformly guess* the values within the output set.

If the maximum number of outputs the procedure returns is σ , and the number of possible different outputs is κ , then the probability of the adversary guessing every output correctly is $(1/\kappa)^\sigma$. Assuming each set of inputs is uniformly selected from the set of all possible inputs, Γ , the probability of an adversary correctly emulating a new call to a counterfeit procedure is

$$P_{call} = \left(\frac{L}{|\Gamma|} \right) + \left(1 - \frac{L}{|\Gamma|} \right) \left(\frac{1}{\kappa} \right)^\sigma.$$

Assuming an extremely large set Γ , the probability of a guess being correct is practically zero, which makes this attack look quite the same as a Memoization Attack. Further, the probability of an adversary successfully emulating ω calls is simply a sequence of Bernoulli trials, so $P_\omega = (P_{call})^\omega$. Therefore, even an adversary with access to unlimited memory and unlimited computational power cannot do better than simply memoizing every input/output relationship pair he observes. Given no innate knowledge of a private procedure, nor any idea of its output distribution, an attack as simple as a Memoization Attack is also the best attack possible.

Note, we do not consider adversaries with knowledge of a non-uniform output distribution of a private procedure. These adversaries may be able to increase their probability

of success when guessing, but algorithmic effort to increase this probability likely exhibits diminishing returns. That is, it is our intuition that the computation time required to approximate the output distribution grows exponentially as a function of this approximation's accuracy.

IV. IMPLEMENTING MEMOIZATION

To investigate the feasibility of a Memoization Attack, we implemented a tool that is capable of observing the execution of a partitioned application, constructing an interaction table, replacing all private procedures with counterfeit procedures, and re-running the partitioned application on alternate inputs. Although building such a tool may sound easy, a naive implementation (as will be first described) would run quite slowly, and require orders of magnitude more storage than typically found on computer workstations. Therefore, we present here some of the tactics used to streamline the implementation so that it can be executed efficiently.

A. System Setup

To create a tool that performs a Memoization Attack, we created a special functional simulator for our chosen TCB (an AEGIS processor). Binary applications are run on the simulator, using some input data set, while attack-specific tasks are performed in the background whenever a transition is encountered from a public to a private procedure. In an attack, the simulator is first started in an observation mode, which saves an interaction table to disk as the application is run. The simulator is then restarted, using a different input data set, and uses the saved interaction table to emulate all private procedures. An assembly rewriting tool was constructed to automate the separation of public and private procedures in compiled applications.

B. Creating an Interaction Table

At first glance, creating an interaction table sounds quite simple: observe an application execute and create a flat mapping of each private procedure's inputs and corresponding outputs. However, emulating a private procedure with such a flat lookup does not work on real systems. The problem is: at the moment when a private procedure is first called an adversary cannot know all of the input values to that private procedure. This is because the procedure itself determines what memory addresses are to be read as inputs *during* execution. This is what we call "*Input Self-Determination*," and is demonstrated in Fig. 3. As we can see in the figure, an adversary cannot know whether




Fig. 3. At call-time, inputs unknown due to self-determination.




Fig. 4. Basic private procedure interaction table.

address B is actually part of the input set until we know the value at address A .

Therefore, our interaction table must contain more information than just the input/output relationship pairs; the table must keep information about the *temporal ordering* of those pairs as they occurred during the execution of the authentic application. One way to visualize such a table is shown in Fig. 4, where each “column” represents a call to the procedure which holds an ordered list of A/V pairs.

C. Emulation using an Interaction Table

Once an adversary has created an interaction table of the type discussed in Section IV-B, he can quite easily emulate any private procedure that is run on inputs he has already seen. We show how this emulation can be done in Fig. 5. When a private procedure is called, the input arguments (registers $r1$, etc.) are matched against the previously seen arguments found at the beginning of each column (independent procedure call) in the table. The set of columns with matching arguments then constitutes a set of candidate previously observed procedure calls that the current procedure call might be an exact copy of. Notice, the next row of all the candidate columns is the same and dictates whether a memory read or write happened in the previously observed calls. If the next row is a write, the write is performed and the following row is inspected (which will still be the same for all columns). However, if the next row is a read, there exists a new input value and the set of candidate rows can possibly be reduced. This continues until a row contains the procedure’s return value, in which case the emulation succeeds, or the set of candidate columns is reduced to zero, in which case the emulation fails.

D. Compressing an Interaction Table

The method for creating an interaction table described in Section IV-B is sound and will work on procedures that have very few inputs or are called only a few times. However, because of the way it maintains order information, this table’s size will grow unmanageably when dealing with procedures that have numerous inputs (many memory reads) or procedures that are called often with many different values for their inputs.

To solve this issue, we instead imagine the structure keeping track of the ordering of inputs and outputs as a tree. Instead of each column representing a unique call to the procedure, the root of the tree represents the beginning of any call to the private procedure and each branch




Fig. 5. Emulation steps using a basic interaction table.

leaving the root is one possible execution path. Such a tree can reduce the amount of *redundant* data found in our interaction table (as might have been noticed in Fig. 5).

An example of what this tree might look like is shown in Fig. 6. Notice that since only memory *reads* can change an execution path, each “tree node” contains the memory address of the next read that should be performed in that one execution path, as well as any writes that must be made before that next read occurs. A private procedure can be emulated using this tree in much the same way as an interaction table is used as described in Section IV-C.

Although this tree drastically reduces the amount of data we must save to perform a Memoization Attack, it still contains some redundancies. The actual data structure used in the implementation of our attack tool is significantly more complicated, and is basically a full, possibly cyclic graph that intelligently keeps track of unique paths from a single start node to many possible end nodes.

The motivation for this is based on the following observation: when a tree keeps track of inputs and outputs to a private procedure, loops within a single call can create

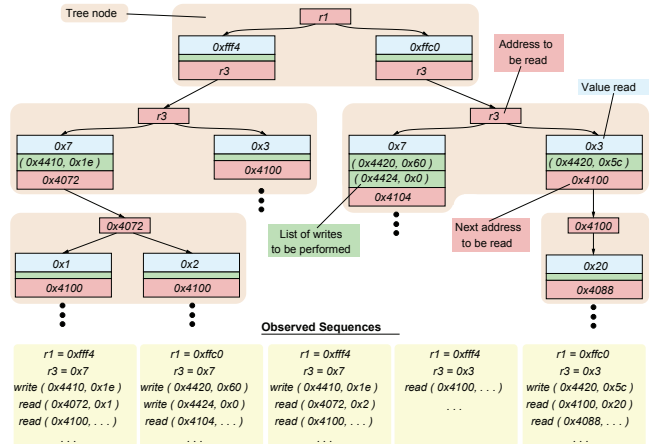


Fig. 6. Interaction tree representing compressed Fig. 5 data.

an extremely deep and repetitive tree, even when the same memory addresses and values are being read over and over. Further, often multiple calls to a private procedure that differ in their initial arguments can later exhibit identical input and output traces for long periods of time. For example, a procedure that takes two input arguments, computes many values using the first argument, but only uses the second argument at the very end of the procedure. When using a tree data structure this would create two separate, but *nearly* identical branches from the root. To construct the graph data structure useful for emulation, we use a unique-numbering method that pinpoints divergences and convergences of execution traces. More details are given in a related thesis [40].

V. EFFECTIVENESS OF MEMOIZATION

After running our implementation of a Memoization Attack on a number of applications in the SPEC CPU2000 [20] suite, we found that two particular types of partitioned applications are susceptible to this attack. By this we mean that an adversary has a good chance of successfully emulating the private procedures of a partitioned application, given an arbitrary or naive partitioning of that application into public and private regions. The two classes of applications can be identified by the types of inputs they require to accomplish their task, applications with *Partially Repeated Input Sets* and applications with *Composite Input Sets*.

A. Applications with Partially Repeated Input Sets

The first class of partitioned applications we found to be susceptible to a Memoization Attack are those with private procedures that are called when the exact same inputs are given to a single execution of an application over and over (such as a repeated common function like “save”), or those that have private procedures that are only ever called when the application reads identical inputs on every execution (such as fixed runtime flags or a common data set).

To illustrate this class, we examined the “*Parser*” application found in the SPEC CPU2000 suite, and assumed a partitioning scheme that only treats individual procedures as private or public. *Parser* executes in two stages, first it processes a fixed dictionary file, and second it analyzes sentences for their grammatic structure using that dictionary. During the analysis, it can also accept special directives from a user that perform standard, repeatable operations. These two traits (two stage execution and special directives) are indicative of an application with partially repeated input sets.

To test whether a Memoization Attack would succeed, we designated the `special_command()` procedure in `main.c` to be private. We then observed the application while sending the `!echo` directive (that sets whether to display output to the screen) which uses the `special_command()` procedure. *Parser* was then run on new input data and we were able to emulate the call the `!echo` without any problems.

Next, the `is_equal()` procedure in `read-dict.c` was made private and observed the application when run on the standard dictionary file and the input file “`smred.in`” taken from MinneSPEC [41]. This procedure is only called while *Parser* reads the dictionary file. In our attack, we were able to correctly emulate this procedure when executing the entire *Parser* application on much larger input files `mdred.in` and `lgred.in`. Both attacks proved successful. Further, Table I shows that the storage requirements for the interaction table (or actually graph, cf. Section IV-D) are not large at all.

TABLE I
SIZE OF MEMOIZED PRIVATE PROCEDURES

Metric	<i>Parser</i>	<i>Parser</i>
	<code>special_command()</code>	<code>is_equal()</code>
Total number of nodes in graph	283	5
Size on disk (in Bytes)	26,972	3,042,968
Maximum number of inputs per call	743	5

B. Applications with Composite Input Sets

The second class of partitioned applications we found susceptible to Memoization Attacks are those that contain private deeply “inner” procedures (such as libraries) that are only fed a finite number of unique inputs (due to the control flow of the calling procedures), no matter what external inputs are given to the application as a whole. In this case a Memoization Attack might succeed by simply observing an authentic application run on *any* large set of inputs, hoping to “cover” or “saturate” the set of inputs to the inner procedure. Because these “saturating” procedures are often not immediately apparent (unlike, perhaps, those mentioned in Section V-A) this class of applications represents a significant problem for a software architecture who would like to prevent Memoization Attacks.

To test whether a Memoization Attack would succeed on this class of applications we attempted to emulate a few procedures from the *Gzip* and *Parser* applications in the SPEC CPU2000 suite, assuming a partitioning scheme that only treats individual procedures as private or public. Table II summarizes the results.

In our attack of *Gzip*, we attempt to emulate a number of procedures using the input file `ref.log` after observing the execution of *Gzip* on just the `ref.random` input file, both the `ref.random` and `ref.graphic` input files, and so on. Even though there is virtually no overlap between these inputs, we found that the `bi_reverse()` procedure can be emulated almost entirely correctly. Of the 1,797 calls made to `bi_reverse()` while processing `ref.log`, 1,741 of the calls contained the *exact* same procedure inputs as had been observed when running *Gzip* on the first four input files.

Similarly, our attack on *Parser* attempted to emulate a number of procedures using the `mdred.in` and `smred.in` input files after observing the execution of the

TABLE II

SUCCESS OF A MEMOIZATION ATTACK ON APPLICATIONS WITH COMPOSITE INPUT SETS

<i>Gzip</i> procedure (Lines of assembly)	Percentage of correct procedure calls while emulating <code>ref.log</code> after observing input set(s) <code>ref.*</code>			
	{ <code>random</code> }	{ <code>random.graphic</code> }	{ <code>random.graphic.program</code> }	{ <code>random.graphic.program.source</code> }
<code>bi_reverse</code> (11)	38% (681/1797)	76% (1362/1797)	84% (1518/1797)	97% (1741/1797)
<code>huft_build</code> (438)	0% (0/27)	0% (0/27)	0% (0/27)	0% (0/27)

<i>Parser</i> procedure (Lines of assembly)	Percentage of correct procedure calls after observing input set <code>lgred.in</code>	
	emulating <code>mdred.in</code>	emulating <code>smred.in</code>
<code>contains_one</code> (123)	33% (1136/3485)	0% (0/71)

application using the `lgred.in` input file. Although none of the procedures could be fully emulated after memoizing input/output relationships pairs from `lgred.in`, it is clear that there are still many duplicated procedure calls between the two unrelated input files. It might even be possible for an adversary to fully emulate the `contains_one()` procedure if he simply observes a large enough set of application inputs from an input file.

From this experimentation we see that a Memoization Attack may be able to succeed even when application inputs seen during emulation are completely unrelated to application inputs recorded during observations.

VI. IDENTIFYING VULNERABLE APPLICATIONS

In Section V we have shown that a Memoization Attack can succeed on certain classes of applications. This may be useful information for an attacker, however, we would rather help software architects *avoid* such attacks through their choice of what procedures to make private and what procedures to make public.

Ideally, we would like to have some *test* that tells us whether a particular private procedure can be easily emulated via a Memoization Attack. The simplest test could be to just run our a Memoization Attack on that procedure. However, to run this attack on all procedures in an application would be computationally infeasible. Instead, information theoretic analyses could be applied, but these might also prove ineffective for practical applications because of their assumptions on entropy, complexity, input space, and “learnability” may be too general. [12][53][54].

Thus, we propose the use of two heuristics, or “*indicators of insecurity*,” that *speculate* upon the likelihood that a private procedure can be emulated in a partitioned application. Importantly, these indicators focus on the interaction of the procedure with the application, rather than the procedure itself. While these tests are not absolute, a procedure that passes them can be given a high confidence that it is immune to a Memoization Attack. Such methods of identifying negative results are used often in problems that do not have a clear positive indicator, for example, tests determining the “randomness” of a random number generator [23][31][33].

A. Indicator I: Input Saturation

Our first test, *Input Saturation* tracks whether a private procedure is only ever fed a finite number of distinct inputs

(AV pairs) by the rest of its application. A simple way to detect this is to run an application using successively more inputs, and observe whether the number of distinct inputs fed to a procedure is linearly related to the number of inputs fed to the application, or if the number of distinct inputs fed to that procedure *saturates* at some level. Procedures that are input saturating are likely easy to emulate through a Memoization Attack (assuming a correlation between the number of unique input AV pairs and the total number of ordered sets of input AV pairs — in reality this correlates well, but not 100% of the time).

Many techniques exist that can estimate the number of unique input AV pairs given to a procedure, however, we simply created a tool that counts and efficiently stores this number while an application is run on a specific input set. Given a large enough input set, this method quickly separates input saturating procedures from those that are not.

As a case example we used this tool to identify input saturating procedures in the SPEC CPU2000 application *Gzip*. Fig. 7 plots the number of unique AV pairs that are fed to the *Gzip* procedure `ct_tally()` during *Gzip*’s execution on five large, orthogonal input sets. For normalization purposes, the x-axis represents the number of calls made to the procedure instead of time. We call this a “*cumulative input density*” plot, and use it as a helpful visualization of when a procedure might be input saturating.

In Fig. 7, we see that the rate of increase in the number of unique AV pairs decreases as more input sets are applied. In fact, the input set `ref.log` did not cause *any* new AV

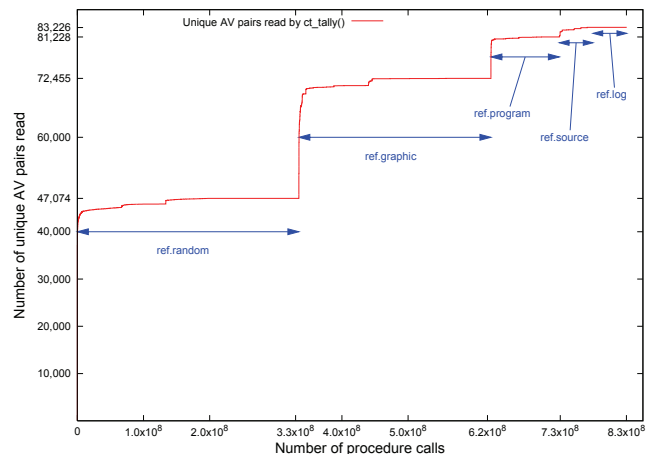


Fig. 7. Cumulative input density plot of unique AV pairs for *Gzip*’s `ct_tally`, when run on a large input set.




Fig. 8. Cumulative input density plots from *Gzip*.

pairs to be fed to the procedure, implying that an adversary might be able to emulate `ct_tally()` on `ref.log` given the observation of the prior four input sets.

To numerically quantify the information in the cumulative input density plot we can use two specific metrics. First, the “average input delta” ($Avg.I\Delta\%$) can tell us the percentage increase in the number of unique AV pairs input to a procedure from call to call. This gives an estimate of how many procedure calls are expected before a new input is seen, and correlates exactly to ω in the formulation of T-AOE in Section II-C. Second, the “saturation weight” (SW) is a single number that gives an idea of the shape of the cumulative input density plot. If the function $w(c)$ represents the number of total unique inputs after the c^{th} call out of N calls, then SW is the normalized sum

$$SW = \frac{1}{Nw(N)} \sum_{c=0}^N w(c).$$

Looking closer at our example, we’ve run *Gzip* on a smaller version of the same inputs and highlighted five of its procedures to demonstrate different levels of input saturation typical in applications. Fig. 8 shows a cumulative input density plot for these procedures executing on the five inputs (normalizing total calls and inputs between procedures to make comparison easier), and Table III gives their average input delta and saturation weight values.

Inspecting the plots in Fig. 8, we see that the two `ct_tally()` and `bi_reverse()` procedures are probably in-

put saturating, and would be susceptible to a Memoization Attack, while the `build_tree()` and `longest_match()` procedures are probably not. It is less clear if `huft_build()` is input saturating since it does not appear to plateau overall, but does plateau for each workload. We attempted a full Memoization Attack on `huft_build()` (Table II) and found that we could not successfully emulate the procedure. This may suggest that, conservatively, only “strongly” saturating procedures might be vulnerable.

The SW values for these procedures also agree with our conjectures, giving values close to 1.0 for easily emulated procedures. However, the $Avg.I\Delta\%$ values do not show a correlation. This is important to note: $Avg.I\Delta\%$ only estimates ω for one specific procedure and has little meaning when used for comparison.

Finally, Table III also shows the average input delta and saturation weight values of the *Gzip* `ct_tally()` procedure using the larger version of the input set. Over the smaller input set, this larger input set causes a drastic decrease in SW from 0.87 to 0.77, lowering it to levels near `huft_build()` (0.72). However, `huft_build()` shows very little susceptibility to a Memoization Attack and `ct_tally()` shows very high susceptibility (as mentioned earlier, no new inputs are seen when the input set `ref.log` is applied). This underscores the need for multiple metrics and a conservative interpretation when making final security decisions.

B. Indicator II: Output Weighting

Our second test, *Output Weighting*, tracks the values that a private procedure outputs and those values’ usefulness to the application as a whole. In essence, this determines how important a private procedure is to the entire application. Output Weighting is possibly a better test than input saturation since it indicates how important a private procedure is to the entire application.

Since an adversary only cares about whole-application functionality, partitioning an application by making less important procedures private may lead to a more successful Memoization Attack. For example, assume during a memoization attack that an adversary cannot return the correct outputs for a private procedure call but continues running. If the previous values in memory still produce the correct behavior (because of range checks, etc.) then the

TABLE III
RATE OF INPUT SATURATION FOR FIVE *Gzip* PROCEDURES

Procedure	Total unique inputs seen after execution on the input set(s) <code>ref.*</code>					$Avg. I\Delta\%$	SW
	{random}	{random,graphic}	{random,graphic, program}	{random,graphic, program,source}	{random,graphic, program,source,log}		
<code>ct_tally</code> (large input)	47,074	72,455	81,228	83,226	83,226	9.7×10^{-9}	0.77
<code>ct_tally</code> (small input)	2,304	2,550	2,768	2,836	2,837	6.9×10^{-7}	0.87
<code>bi_reverse</code>	569	580	580	580	581	6.3×10^{-5}	0.99
<code>huft_build</code>	0	2,500	3,170	3,510	3,586	7.4×10^{-3}	0.72
<code>build_tree</code>	11,873	23,611	29,945	32,103	32,672	5.9×10^{-3}	0.51
<code>longest_match</code>	4.78 M	8.33 M	10.13 M	11.19 M	11.61 M	2.7×10^{-6}	0.51

TABLE IV
OUTPUT WEIGHTS FOR SIX *Gzip* PROCEDURES

Procedure	Tot. uniq. reads	Tot. uniq. writes	Public readers	$\Phi(\cdot)$ weight
<code>inflate_codes</code>	4,240,569	5,151,281	9	390,657
<code>ct_tally</code>	2,837	4,214,758	4	1,343,144
<code>bi_reverse</code>	581	259	2	93
<code>huft_build</code>	3,586	59,224	4	96
<code>build_tree</code>	32,672	21,000	4	2
<code>longest_match</code>	11,610,835	515	1	13,010

adversary will still be content. This “low importance” of the outputs of the private procedure has allowed a Memoization Attack to succeed. Another good example arises when a private procedure’s outputs are only ever used by a single, simple public procedure that is itself easily emulated. In this case, the inputs and outputs of the private procedure do not matter, and a Memoization Attack can succeed by simply emulating the simple public procedure wrapping the private procedure. The output weighting of a private procedure should be able to identify both cases as susceptible to a Memoization Attack.

Tracing the entire flow of data throughout an application and deriving “usefulness” information is again a hard task. Therefore we suggest a simple test that estimates how much “usefulness” public procedures derive from the outputs of a private procedure. This test recognizes that a private procedure can only impact the outputs of the entire application if its own outputs are passed along and or used by other public procedures. This test is called the the “*output weight*” $\Phi(\cdot)$ of a procedure, and is defined by

$$\Phi(\eta) = \sum_{\forall(\iota, \kappa_i) \in \eta} \frac{\kappa_i}{\iota_i}$$

Here η is a set of pairs (ι, κ) , where ι is the number of unique outputs written by a private procedure and read by a public procedure, and κ is of the total number of unique outputs from that public procedure. For example, if five public procedures use the outputs of one private procedure as their input, then $|\eta| = 5$. Here the fraction κ_i/ι_i indicates the impact of a private procedure’s outputs on the outputs of public procedures. In other words, this indicates how the utility of a private procedure’s outputs is “*amplified*” as the outputs are used by the rest of the application. A very large value of κ_i/ι_i implies that a private procedure’s outputs are important.

As with input saturation, the output weight of a procedure can be estimated using many techniques. However, for simplicity we made a tool that efficiently tabulated the number of unique outputs that are transferred between procedures while an application is run on some input set. From these tabulations we can compute the output weight, as shown in Fig. 9 where the number of unique outputs out of the `inflate_codes()` procedure in *Gzip* are used to determine an output weight of $\Phi(\cdot) = 390,657$.

Looking again at our *Gzip* example, Table IV gives the computed output weight of six select procedures from an execution of *Gzip* on the small version of the same five input

sets. We also show the number of unique reads and write a private procedure performs, and the number of public procedures that read what that private procedure writes.

We see from Table IV that procedures `inflate_codes()` and `ct_tally()` produce many unique outputs that are then fed to other procedures that in turn produce many unique outputs. If this trend continues, it is highly likely that the outputs of the application as a whole will depend significantly on the outputs of `inflate_codes()` and `ct_tally()`. Alternately, the outputs of the `bi_reverse()`, `huft_build()`, `build_tree()`, and `longest_match()` procedures only produce a limited number of unique output AV pairs, and these outputs are passed to procedures that do not produce many more unique outputs. Therefore, it might be easy for an adversary to perform a Memoization function on these latter four procedures

C. Interpreting Indicators

The two indicators presented are not absolute and are not the only possible metrics of whether an application is susceptible to a Memoization Attack. In practice a software architect should apply as many tests as possible, including more complicated tests, before confidently labeling a private procedure safe from attack. Further, the indicator results should be viewed in tandem. As seen in the examples, the set of “safe” procedures determined by the input saturation test does not perfectly overlap with the set of “safe” procedures determined by the output weighting test. Ultimately, no test can rule out the possibility of a Memoization Attack since this attack depends directly on the input set applied to the application. Thus, the amount of testing performed is yet another design choice when partitioning an application into private and public regions.

VII. RELATED WORK

Only a few studies [6][7][8][17][28] have specifically examined software secrecy and modification of application code to prevent an adversary from determining its contents, sometimes suggesting techniques with which to decipher these contents. To counter such techniques *obfuscation* transforms have been proposed that make an application incomprehensible, but still functionally correct [9][10]. Unfortunately, it has been proven that cryptographically secure obfuscation is generally impossible for a large family of functions [4] (although a few specific families have been shown to be obfuscatable [29][56]).

A more popular way of concealing application instructions is through encryption. Homomorphic encryption schemes [45][46] allow meaningful computations to be performed on encrypted data, but are not general enough for practical use. Instead, many have suggested using a small trusted computing base to decrypt ciphertext applications and to execute instructions confidentially [3][22][32][38]. This idea of using specialized security hardware and secure coprocessors has seen many manifestations

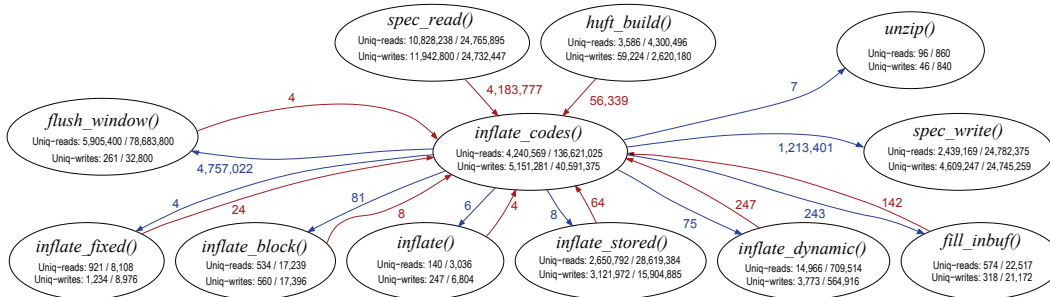


Fig. 9. The unique outputs of the `inflate_codes()` procedure in *Gzip* used to compute its output weighting of $\Phi(\cdot) = 390,657$.

[14][25][58][59][61].

Recently, physically secure architectures have been proposed that reduce this trusted computing base to a single chip while also supporting application partitioning [27][49][50][51]. These allow applications to be encrypted and executed on a processor without revealing any information to even the device owner. Even though these architectures encrypt application instructions, additional methods must still be employed to defend against side-channel attacks [1][15][16][62].

Application encryption can be used for copy protection (by binding software execution to a specific key), but concepts of watermarking [10][55], renewability [30], online-verification [11][13], and hardware assisted authentication [18][26][37][43] have also all been suggested as means to enforce basic software licensing. Unfortunately, many of these methods suffer from the same fundamental problem: they add *extra* code to the application. While it may be extremely difficult, a motivated adversary can almost always remove this extra code.

Architectural support for application partitioning is not a new concept [48][59], however we believe that this paper is the first analysis of the problem of code secrecy when considering application operation as a whole. Program slicing has been proposed [60] as a means to prevent piracy, however it does not address the possibility that program secrecy may not be guaranteed in a partitioned application. Other compiler and language support for secure partitioning has been proposed [5][44], but focuses on a different problem of application etiquette and information flow.

Finally, the indicators discussed are basically an analysis of code complexity. Many empirical software complexity metrics have been developed over the years [19][21][34][39][42][47][52]. Of these, one [57] does discuss the complexity of de-constructing an application, but does not focus on security. Deeper investigation of many of the problems examined here can be found in a related thesis [40].

VIII. CONCLUSIONS

Application partitioning has been suggested by a number of works as a means to allow an application to run efficiently on a TCB while preserving security guarantees. We have investigated the problem of maintaining IP se-

crecy and copy protection in a partitioned application by looking at how to prevent an adversary from duplicating the *functionality* or *utility* of a private partition. Importantly, this analysis often depends more on the make-up of the entire application than it does on the make-up of a private partition. This is formalized by the adversarial goal of Temporal Application Operation Equivalence.

To tackle this question we have analyzed the simplest form of attack, a Memoization Attack. This attack tabulates input/output relationships seen during legitimate executions of a private partition, and replays what had been saved when it later emulates that private partition. Under certain assumptions this attack is the best an adversary can possibly perform. We implemented a full Memoization Attack, and described some necessary techniques that allow this attack to run with reasonable storage and computation time restrictions. Running this attack on real-world applications, two classes of partitioned applications were found that are susceptible to Memoization Attacks.

To help software architects identify vulnerable classes of partitioned applications during development we proposed two efficient tests that can be used to identify an application’s susceptibility to a Memoization Attack. These tests were implemented efficiently and run on an example application to demonstrate their usefulness.

This has only been an initial step in the investigation of the security hazards inherent in partitioned applications. There are also many complicating issues that can be explored on this topic, such as the ease with which private libraries can be attacked (since multiple input sets are used in a predictable way), or if multiple versions of an application can make it any easier on an adversary.

Ultimately, the question of whether or not an adversary can duplicate an unseen private partition is problematic at best. The exact security of a system can only be guaranteed in terms of the adversarial model it defends against. Further, such concepts like “human intuition” do not easily fit into models, even though this is often the most important factor when performing such attacks. Therefore, our approach in answering this question is to begin with a simple, practical, but universal attack model that can then be built upon by more complicated attack models that address specific domains of human-injected knowledge.

REFERENCES

- [1] J. Agat. Transforming out timing leaks. In *27th ACM Principles of Programming Languages*, January 2000.
- [2] R. Anderson and M. Kuhn. Low cost attacks on tamper resistant devices. In *IWSP: LNCS*, volume 1361, pages 125–136, 1997.
- [3] D. Aucsmith. Tamper Resistant Software: An Implementation. In *Proceeding of the 1st Information Hiding Workshop*, 1996.
- [4] B. Barak, O. Goldreich, R. Impagliazzo, S. Rudich, A. Sahai, S. Vadhan, and K. Yang. On the (im)possibility of obfuscating programs. *CRYPTO'01: LNCS*, 2139:1–18, 2001.
- [5] D. Brumley and D. Song. Privtrans: Automatically partitioning programs for privilege separation. In *13th USENIX*, Aug. 2004.
- [6] E. J. Byrne. Software reverse engineering: a case study. *Software Practice and Experience*, 21(12):1349–1364, 1991.
- [7] E. J. Chikofsky and J. H. C. II. Reverse engineering and design recovery: A taxonomy. *IEEE Software*, 7(1):13–17, 1990.
- [8] S. C. Choi and W. Scacchi. Extracting and restructuring the design of large systems. *IEEE Software*, 7(1):66–71, 1990.
- [9] C. Collberg, C. Thomborson, and D. Low. A taxonomy of obfuscating transformations. Technical Report 148, Department of Computer Science, University of Auckland, July 1997.
- [10] C. S. Collberg and C. Thomborson. Watermarking, tamper-proofing, and obfuscation: Tools for software protection. *IEEE Transactions on Software Engineering*, 28(8):735–746, 2002.
- [11] M. Corporation. Technical Overview of Windows Rights Management Services. Microsoft white paper, Apr. 2005.
- [12] T. M. Cover and J. A. Thomas. *Elements of Information Theory*. John Wiley and Sons, 1991.
- [13] O. Dvir, M. Herlihy, and N. N. Shavit. Virtual Leashing: Internet-Based Software Piracy Protection. In *25th ICDCS*, pages 283–292, 2005.
- [14] T. Gilmont, J.-D. Legat, and J.-J. Quisquater. Enhancing security in the memory management unit. In *EUROMICRO*, 1999.
- [15] O. Goldreich. Towards a theory of software protection and simulation by oblivious rams. In *ACM STOC*, pages 182–194, 1987.
- [16] O. Goldreich and R. Ostrovsky. Software Protection and Simulation on Oblivious RAMs. *J. of the ACM*, 43(3):431–473, 1996.
- [17] J. R. Gosler. Software protection: Myth or reality? In *CRYPTO'85 (LNCS No. 218)*, pages 140–157, 1986.
- [18] T. C. Group. TCG Specification Architecture Overview Revision 1.2. <http://www.trustedcomputinggroup.com/home>, 2004.
- [19] W. A. Harrison and K. I. Magel. A complexity measure based on nesting level. *SIGPLAN Not.*, 16(3):63–74, 1981.
- [20] J. L. Henning. SPEC CPU2000: Measuring CPU performance in the new millennium. *IEEE Computer*, July 2000.
- [21] S. Henry and D. Kafura. Software structure metrics based on information flow. *IEEE Trans. on Soft. Eng.*, 7(5):510–518, 1981.
- [22] S. T. Kent. *Protecting Externally Supplied Software in Small Computers*. PhD thesis, Massachusetts Institute of Tech., 1980.
- [23] S. Kim, K. Umeno, and A. Hasegawa. On the NIST Statistical Test Suite for Randomness. In *IEICE Technical Report, Vol. 103, No. 449, pp. 21-27*, 2003.
- [24] P. Kocher, J. Jaffe, and B. Jun. Differential Power Analysis. *Lecture Notes in Computer Science*, 1666:388–397, 1999.
- [25] M. Kuhn. The TrustNo1 Cryptoprocessor Concept. Technical Report, Purdue University, April 1997, 1997.
- [26] R. B. Lee, P. C. S. Kwan, J. P. McGregor, J. Dwoskin, and Z. Wang. Architecture for protecting critical secrets in microprocessors. In *ISCA*, pages 2–13, 2005.
- [27] D. Lie. *Architectural Support for Copy and Tamper-Resistant Software*. PhD thesis, Stanford University, 2003.
- [28] R. Lutz. Recovering high-level structure of software systems using a minimum description length principle. In *AICS*, 2002.
- [29] B. Lynn, M. Prabhakaran, and A. Sahai. Positive results and techniques for obfuscation. In *EUROCRYPT*, pages 20–39, 2004.
- [30] M. Jakobsson, M.K. Reiter. Discouraging software piracy using software aging. In *DRM (CCS-8 Workshop)*, pages 1–12, 2002.
- [31] G. Marsaglia and W. W. Tsang. Some difficult-to-pass tests of randomness. *Journal of Statistical Software*, 7(3):1–8, 2002.
- [32] T. Maude and D. Maude. Hardware protection against software piracy. *Communications of the ACM*, 27(9):950–959, 1984.
- [33] NIST Special Publication 800-22. A statistical test suite for random and pseudorandom number generators for cryptographic applications. IT Laboratory of NIST, May 2000.
- [34] McCabe. A complexity measure. *IEEE Trans. on Soft. Eng.*, 2(4):308–320, Dec. 1976.
- [35] T. S. Messerges, E. A. Dabbish, and R. H. Sloan. Examining smart-card security under the threat of power analysis attacks. *IEEE Transactions on Computers*, 51(5):541–552, 2002.
- [36] Microcosm. DinkeyDongle. www.microcosm.co.uk, 2007.
- [37] Microsoft. Next-Generation Secure Computing Base. www.microsoft.com/resources/ngscb/default.aspx, 2007.
- [38] C. Morgan. How Can We Stop Software Piracy. *BYTE*, 6(5):6–10, May 1981.
- [39] J. C. Munson and T. M. Khoshgoftaar. Measurement of data structure complexity. *Journal of System Software*, 20(3):217–225, 1993.
- [40] C. W. O'Donnell. Secure Application Partitioning for Intellectual Property Protection. Master's thesis, Massachusetts Institute of Technology, Aug. 2005.
- [41] A. J. K. Osowski and D. J. Lilja. MinneSPEC: A New SPEC Benchmark Workload for Simulation-Based Computer Architecture Research. *Computer Architecture Letters*, 1, 2002.
- [42] E. I. Oviedo. Control Flow, Data Flow, and Program Complexity. In *Proceedings of COMPSAC*, pages 145–152, 1980.
- [43] P. England, B. Lampson, J. Manferdelli, M. Peinado, B. Willman. A trusted open platform. *Computer*, 36(7):55–62, 2003.
- [44] S. Zdancewic, L. Zheng, N. Nystrom, A. Myers. Secure program partitioning. *ACM Trans. Comp. Sys.*, 20(3):283–328, 2002.
- [45] T. Sander and C. F. Tschudin. On Software Protection via Function Hiding. *Lecture Notes in Comp. Sci.*, 1525:111–123, 1998.
- [46] T. Sander and C. F. Tschudin. Towards mobile cryptography. In *Proc. of the IEEE Symposium on Security and Privacy*, 1998.
- [47] S.R. Chidamber, C.F. Kemerer. A metrics suite for object oriented design. *IEEE Trans. on Soft. Eng.*, 20(6):476–493, 1994.
- [48] S.R. White, L. Comerford. Abyss: An architecture for software protection. *IEEE Trans. on Soft. Eng.*, 16(6):619–629, 1990.
- [49] G. E. Suh. *AEGIS: A Single-Chip Secure Processor*. PhD thesis, Massachusetts Institute of Technology, 2005.
- [50] G. E. Suh, D. Clarke, B. Gassend, M. van Dijk, and S. Devadas. AEGIS: Architecture for Tamper-Evident and Tamper-Resistant Processing. In *17th ICS*, June 2003.
- [51] G. E. Suh, C. W. O'Donnell, I. Sachdev, and S. Devadas. Design and Implementation of the AEGIS Single-Chip Secure Processor Using Physical Random Functions. In *32nd ISCA*, June 2005.
- [52] K.-C. Tai. A program complexity metric based on data flow information in control graphs. In *7th ICSE*, pages 239–248, 1984.
- [53] V. N. Vapnik. *Statistical Learning Theory*. Wiley & Sons, 1998.
- [54] V. N. Vapnik. *The Nature of Statistical Learning Theory, Second Edition*. Springer, 1999.
- [55] R. Venkatesan, V. V. Vazirani, and S. Sinha. A graph theoretic approach to software watermarking. In *4th International Workshop on Information Hiding*, pages 157–168, 2001.
- [56] H. Wee. On obfuscating point functions. In *ACM STOC*, pages 523–532, May 2005.
- [57] H. Yang, P. Luker, and W. C. Chu. Measuring abstractness for reverse engineering in a re-engineering tool. In *13th ICSM*, 1997.
- [58] L. G. J. Yang and Y. Zhang. Fast Secure Processor for Inhibiting Software Piracy and Tampering. In *36th MICRO*, Dec. 2003.
- [59] B. S. Yee. *Using Secure Coprocessors*. PhD thesis, CMU, 1994.
- [60] X. Zhang and R. Gupta. Hiding program slices for software security. In *1st CGO*, pages 325–336, 2003.
- [61] Y. Zhang, J. Yang, Y. Lin, and L. Gao. Architectural Support for Protecting user Privacy on Trusted Processors. *SIGARCH Computer Architecture News*, 33(1):118–123, 2005.
- [62] X. Zhuang, T. Zhang, and S. Pande. Hide: an infrastructure for efficiently protecting information leakage on the address bus. *SIGPLAN Not.*, 39(11):72–84, 2004.