



Automating the Development of Chosen Ciphertext Attacks

Gabrielle Beck, Maximilian Zinkus, and Matthew Green,
Johns Hopkins University

<https://www.usenix.org/conference/usenixsecurity20/presentation/beck>

This paper is included in the Proceedings of the
29th USENIX Security Symposium.

August 12-14, 2020

978-1-939133-17-5

Open access to the Proceedings of the
29th USENIX Security Symposium
is sponsored by USENIX.

Automating the Development of Chosen Ciphertext Attacks



Gabrielle Beck*
Johns Hopkins University
becgabri@cs.jhu.edu

Maximilian Zinkus*
Johns Hopkins University
zinkus@cs.jhu.edu

Matthew Green
Johns Hopkins University
mgreen@cs.jhu.edu

Abstract

In this work we investigate the problem of automating the development of adaptive chosen ciphertext attacks on systems that contain vulnerable *format oracles*. Unlike previous attempts, which simply automate the execution of known attacks, we consider a more challenging problem: to programmatically derive a novel attack strategy, given only a machine-readable description of the plaintext verification function and the malleability characteristics of the encryption scheme. We present a new set of algorithms that use SAT and SMT solvers to reason deeply over the design of the system, producing an automated attack strategy that can entirely decrypt protected messages. Developing our algorithms required us to adapt techniques from a diverse range of research fields, as well as to explore and develop new ones. We implement our algorithms using existing theory solvers. The result is a practical tool called *Delphinium* that succeeds against real-world and contrived format oracles. To our knowledge, this is the first work to automatically derive such complex chosen ciphertext attacks.

1 Introduction

The past decades have seen enormous improvement in our understanding of cryptographic protocol design. Despite these advances, vulnerable protocols remain widely deployed. In many cases this is a result of continued support for legacy protocols and ciphersuites, such as TLS’s CBC-mode ciphers [7, 64], export-grade encryption [4, 9, 19], and legacy email encryption [59]. However, support for legacy protocols does not account for the presence of vulnerabilities in more recent protocols and systems [36, 42, 47, 72, 74].

*These authors contributed equally to the work.

In this work we consider a specific class of vulnerability: the continued use of unauthenticated symmetric encryption in many cryptographic systems. While the research community has long noted the threat of adaptive-chosen ciphertext attacks on malleable encryption schemes [17, 18, 56], these concerns gained practical salience with the discovery of *padding oracle* attacks on a number of standard encryption protocols [6, 7, 13, 22, 30, 40, 51, 52, 73]. Despite repeated warnings to industry, variants of these attacks continue to plague modern systems, including TLS 1.2’s CBC-mode ciphersuite [5, 7, 48] and hardware key management tokens [10, 13]. A generalized variant, the *format oracle attack* can be constructed when a decryption oracle leaks the result of applying some (arbitrarily complex) format-checking predicate F to a decrypted plaintext. Format oracles appear even in recent standards such as XML encryption [42, 45], Apple’s iMessage [36] and modern OpenPGP implementations [47, 59]. These attacks likely represent the “tip of the iceberg”: many vulnerable systems may remain undetected, due to the difficulty of exploiting non-standard format oracles.

From a constructive viewpoint, format oracle vulnerabilities seem easy to mitigate: simply mandate that protocols use authenticated encryption. Unfortunately, even this advice may be insufficient: common authenticated encryption schemes can become insecure due to implementation flaws such as nonce re-use [21, 43, 46]. Setting aside implementation failures, the continued deployment of unauthenticated encryption raises an obvious question: *why do these vulnerabilities continue to appear in modern protocols?* The answer highlights a disconnect between the theory and the practice of applied cryptography. In many cases, a vulnerable protocol is not obviously an *exploitable* protocol. This is particularly true for non-standard format oracles which require en-




Figure 1: Output of a format oracle attack that our algorithms developed against a bitwise padding check oracle F_{bitpad} (see §5.2 for a full description). The original ciphertext is a valid 128-bit (random) padded message encrypted using a stream cipher. Each row of the bitmap represents a *malleation string* that was exclusive-ORed with the ciphertext prior to making a decryption query.

tirely new exploit strategies. As a concrete example, the authors of [36] report that Apple did not repair a complex gzip compression format oracle in the iMessage protocol when the lack of authentication was pointed out; but did mitigate the flaw when a concrete exploit was demonstrated. Similar flaws in OpenPGP clients [36, 59] and PDF encryption [55] were addressed only when researchers developed proof-of-concept exploits. The unfortunate aspect of this strategy is that cryptographers’ time is limited, which leads protocol designers to discount the exploitability of real cryptographic flaws.

Removing the human element. In this work we investigate the feasibility of *automating the design and development* of adaptive chosen ciphertext attacks on symmetric encryption schemes. We stress that our goal is not simply to automate the execution of known attacks, as in previous works [45]. Instead, we seek to develop a methodology and a set of tools to (1) evaluate if a system is vulnerable to practical exploitation, and (2) programmatically derive a novel exploit strategy, given only a description of the target. This removes the expensive human element from attack development.

To emphasize the ambitious nature of our problem, we summarize our motivating research question as follows:

Given a machine-readable description of a format checking function F along with a description of the encryption scheme’s malleation

properties, can we programmatically derive a chosen-ciphertext attack that allows us to efficiently decrypt arbitrary ciphertexts?

Our primary requirement is that the software responsible for developing this attack should require no further assistance from human beings. Moreover, the developed attack must be efficient: ideally it should not require substantially more work (as measured by number of oracle queries and wall-clock execution time) than the equivalent attack developed through manual human optimization.

To our knowledge, this work represents the first attempt to automate the discovery of *novel* adaptive chosen ciphertext attacks against symmetric format oracles. While our techniques are designed to be general, in practice they are unlikely to succeed against every possible format checking function. Instead, in this work we initiate a broader investigation by exploring the limits of our approach against various real-world and contrived format checking functions. Beyond presenting our techniques, our practical contribution of this work is a toolset that we name *Delphinium*, which produces highly-efficient attacks across several such functions.

Relationship to previous automated attack work. Previous work [12, 26, 58] has looked at automatic discovery and exploitation of *side channel* attacks. In this setting, a program combines a fixed secret input with many “low” inputs that are (sometimes adaptively) chosen by an attacker, and produces a signal, *e.g.*, modeling a timing result. This setting can be viewed as a special case of our general model (and vice versa). Like our techniques, several of these works employ SAT solvers and model counting techniques. However, beyond these similarities, there are fundamental differences that manifest in our results: (1) in this work we explore a new approach based on approximate model counting, and (2) as a result of this approach, our results operate over much larger secret domains than the cited works. To illustrate the differences, our experimental results succeed on secret (message) domains of several hundred bits in length, with malleation strings (“low inputs”) drawn from similarly-sized domains. By contrast, the cited works operate over smaller secret domains that rarely even reach a size of 2^{24} . Moreover, our format functions are relatively complex. It is an open question to determine whether the experimental results in the cited works can be scaled using our techniques.

Our contributions. In this work we make the following contributions:

- We propose new, and *fully automated* algorithms

for developing format oracle attacks on symmetric encryption (and hybrid encryption) schemes. Our algorithms are designed to work with arbitrary format checking functions, using a machine-readable description of the function and the scheme’s malleation features to develop the attack strategy.

- We design and implement novel attack-development techniques that use approximate model counting techniques to achieve significantly greater efficiency than previous works. These techniques may be of independent interest.
- We show how to implement our technique practically with existing tools such as SAT and SMT solvers; and propose a number of efficiency optimizations designed to improve performance for specific encryption schemes and attack conditions.
- We develop a working implementation of our techniques using “off-the-shelf” SAT/SMT packages, and provide the resulting software package (which we call *Delphinium*), an artifact accompanying this submission, as an open source tool for use and further development by the research community.
- We validate our tool experimentally, deriving several attacks using different format-checking functions. These experiments represent, to our knowledge, the first evidence of a completely functioning end-to-end machine-developed format oracle attack.

1.1 Intuition

Implementing a basic format oracle attack. In a typical format oracle attack, the attacker has obtained some target ciphertext $C^* = \text{Encrypt}_K(M^*)$ where K and M^* are unknown. She has access to a decryption oracle that, on input any chosen ciphertext C , returns $F(\text{Decrypt}_K(C)) \in \{0, 1\}$ for some known predicate F . The attacker may have various goals, including plaintext recovery and forgery of new ciphertexts. Here we will focus on the former goal.

Describing malleability. Our attacks exploit the malleability characteristics of symmetric encryption schemes. Because the encryption schemes themselves can be complex, we do not want our algorithms to reason over the encryption mechanism itself. Instead, for a given encryption scheme Π , we require the user to develop two efficiently-computable functions that define the malleability properties of the scheme. The function $\text{Maul}_{\text{ciph}}^{\Pi}(C, S) \rightarrow C'$ takes as input a valid ciphertext and some opaque *malleation instruction string* S (henceforth “malleation

string”), and produces a new, mauled ciphertext C' . The function $\text{Maul}_{\text{plain}}^{\Pi}(M, S) \rightarrow M'$ computes the equivalent malleation over some plaintext, producing a plaintext (or, in some cases, a set of possible plaintexts¹). The essential property we require from these functions is that the plaintext malleation function should “predict” the effects of encrypting a plaintext M , mauling the resulting ciphertext, then subsequently decrypting the result. For some typical encryption schemes, these functions can be simple: for example, a simple stream cipher can be realized by defining both functions to be bitwise exclusive-OR. However, malleation functions may also implement features such as truncation or more sophisticated editing, which could imply a complex and structured malleation string.

Building block: theory solvers. Our techniques make use of efficient theory solvers, such as SAT and Satisfiability Modulo Theories (SMT) [1, 49]. SAT solvers apply a variety of tactics to identify or rule out a satisfying assignment to a boolean constraint formula, while SMT adds a broader range of theories and tactics such as integer arithmetic and string logic. While in principle our techniques can be extended to work with either system, in practice we will focus our techniques to use quantifier-free operations over bitvectors (a theory that easily reduces to SAT). In later sections, we will show how to realize these techniques efficiently using concrete SAT and SMT packages.

Anatomy of our attack algorithm. The essential idea in our approach is to model each phase of a chosen ciphertext attack as a constraint satisfaction problem. At the highest level, we begin by devising an initial constraint formula that defines the known constraints on (and hence, implicitly, a set of candidates for) the unknown plaintext M^* . At each phase of the attack, we will use our current knowledge of these constraints to derive an *experiment* that, when executed against the real decryption oracle, allows us to “rule out” some non-zero number of plaintext candidates. Given the result of a concrete experiment, we can then update our constraint formula using the new information, and continue the attack procedure until no further candidates can be eliminated.

In the section that follows, we use $\mathcal{M}_0, \mathcal{M}_1$ to represent the partition of messages induced by a malleation string. M_0 and M_1 represent concrete plaintext message assignments chosen by the solver, members of the respective partitions.

¹This captures the fact that, in some encryption schemes (e.g., CBC-mode encryption), malleation produces *key-dependent* effects on the decrypted message. We discuss and formalize this in §2.

The process of deriving the malleation string represents the core of our technical work. It requires our algorithms to reason deeply over both the plaintext malleation function and the format checking function in combination. To realize this, we rely heavily on theory solvers, together with some novel optimization techniques.

Attack intuition. We now explain the full attack in greater detail. To provide a clear exposition, we will begin this discussion by discussing a simplified and *inefficient* precursor algorithm that we will later optimize to produce our main result. Our discussion below will make a significant simplifying assumption that we will later remove: namely, that $\text{Maul}_{\text{plain}}$ will output exactly one plaintext for any given input. This assumption is compatible with common encryption schemes such as stream ciphers, but will not be valid for other schemes where malleation can produce key-dependent effects following decryption.

We now describe the basic steps of our first attack algorithm.

Step 0: Initialization. At the beginning of the attack, our attack algorithm receives as input a target ciphertext C^* , as well as a machine-readable description of the functions F and $\text{Maul}_{\text{plain}}$. We require that these descriptions be provided in the form of a constraint formula that a theory solver can reason over. To initialize the attack procedure, the user may also provide an initial constraint predicate $G_0 : \{0, 1\}^n \rightarrow \{0, 1\}$ that expresses all known constraints over the value of M^* .² (If we have no *a priori* knowledge about the distribution of M^* , we can set this initial formula G_0 to be trivial).

Beginning with $i = 1$, the attack now proceeds to iterate over the following two steps:

Step 1: Identify an experiment. Let G_{i-1} be the current set of known constraints on M^* . In this first step, we employ the solver to identify a malleation instruction string S as well as a pair of distinct plaintexts M_0, M_1 that each satisfy the constraints of G_{i-1} . Our goal is to identify an assignment for (S, M_0, M_1) that induces the following specific properties on $\mathcal{M}_0, \mathcal{M}_1$: namely, that each message in the pair, when mauled using S and then evaluated using the format checking function, results in a *distinct* output from F . Expressed more concretely, we require the solver to identify an assignment that satisfies the following constraint formula:

$$G_{i-1}(M_0) = G_{i-1}(M_1) = 1 \wedge \quad (1) \\ \forall b \in \{0, 1\} : F(\text{Maul}_{\text{plain}}(M_b, S)) = b$$

²Here n represents an upper bound on the length of the plaintext M^* .

If the solver is unable to derive a satisfying assignment to this formula, we conclude the attack and proceed to Step (3). Otherwise we extract a concrete satisfying assignment for S , assign this value to \mathbf{S} , and proceed to the next step.

Step 2: Query the oracle; update the constraints. Given a concrete malleation string \mathbf{S} , we now apply the ciphertext malleation function to compute an experiment ciphertext $C \leftarrow \text{Maul}_{\text{ciph}}(C^*, \mathbf{S})$, and submit C to the decryption oracle. When the oracle produces a concrete result $\mathbf{r} \in \{0, 1\}$, we compute an updated constraint formula G_i such that for each input M , it holds that:

$$G_i(M) \leftarrow (G_{i-1}(M) \wedge F(\text{Maul}_{\text{plain}}(M, \mathbf{S})) = \mathbf{r})$$

If possible, we can now ask the solver to *simplify* the formula G_i by eliminating redundant constraints in the underlying representation. We now set $i \leftarrow i + 1$ and return to Step (1).

Step 3: Attack completion. The attack concludes when the solver is unable to identify a satisfying assignment in Step (1). In the ideal case, this occurs because the constraint system G_{i-1} admits only one possible candidate plaintext, M^* : when this happens, we can employ the solver to directly recover M^* and complete the attack. However, the solver may also fail to find an assignment because no further productive experiment can be generated, or simply because finding a solution proves computationally intractable. When the solver conclusively rules out a solution at iteration $i = 1$ (*i.e.*, prior to issuing any decryption queries) this can be taken as an indication that a viable attack is not practical using our techniques. Indeed, this feature of our work can be used to rule out the exploitability of certain systems, even without access to a decryption oracle. In other cases, the format oracle may admit only partial recovery of M^* . If this occurs, we conclude the attack by applying the solver to the final constraint formula G_{i-1} to extract a human-readable description of the remaining candidate space (*e.g.*, the bits of M^* we are able to uniquely recover).

Remark on efficiency. A key feature of the attack described above is that it is *guaranteed* to make progress at each round in which the solver is able to find a satisfying assignment to Equation (1). This is fundamental to the constraint system we construct: our approach forces the solver to ensure that each malleation string S implicitly partitions the candidate message set into a pair (M_0, M_1) , such that malleation of messages in either subset by S will produce distinct outputs from the format checking function F . As a consequence of this, for any possible result from the real-world decryption oracle, the updated




Figure 2: Left: illustration of a plaintext candidate space defined by G_{i-1} , highlighting the two subsets $\mathcal{M}_0, \mathcal{M}_1$ induced by a specific malleation string \mathbf{S} . Right: the candidate space defined by G_i , in which many candidates have been eliminated following an oracle response $b = 1$.

constraint formula G_i must eliminate at least one plaintext candidate that satisfied the previous constraints G_{i-1} .

While this property ensures progress, it does not imply that the resulting attack will be *efficient*. In some cases, the addition of a new constraint will fortuitously rule out a large number of candidate plaintexts. In other cases, it might only eliminate a single candidate. As a result, there exist worst-case attack scenarios where the algorithm requires *as many queries as there are candidates for M^** , making the approach completely unworkable for practical message sizes. Addressing this efficiency problem requires us to extend our approach.

Improving query profitability. We can define the *profitability* $\psi(G_{i-1}, G_i)$ of an experimental query by the number of plaintext candidates that are “ruled out” once an experiment has been executed and the constraint formula updated. In other words, this value is defined as the number of plaintext candidates that satisfy G_{i-1} but do *not* satisfy G_i . The main limitation of our first attack strategy is that it does not seek to optimize each experiment to maximize query profitability.

To address this concern, let us consider a more general description of our attack strategy, which we illustrate in Figure 2. At the i^{th} iteration, we wish to identify a malleation string \mathbf{S} that defines two disjoint subsets $\mathcal{M}_0, \mathcal{M}_1$ of the current candidate plaintext space, such that for any concrete oracle result $\mathbf{r} \in \{0, 1\}$ and $\forall M \in \mathcal{M}_{\mathbf{r}}$ it holds that $F(\text{Maul}_{\text{plain}}(M, \mathbf{S})) = \mathbf{r}$. In this description, any concrete decryption oracle result must “rule out” (at a minimum) every plaintext contained in the subset $\mathcal{M}_{1-\mathbf{r}}$. This sets $\psi(G_{i-1}, G_i)$ equal to the cardinality of $\mathcal{M}_{1-\mathbf{r}}$.

To increase the profitability of a given query, it is therefore necessary to maximize the size of $\mathcal{M}_{1-\mathbf{r}}$. Of course, since we do not know the value \mathbf{r} prior to issuing a decryption oracle query, the obvious strategy is to find \mathbf{S} such

that *both* $\mathcal{M}_0, \mathcal{M}_1$ are as large as possible. Put slightly differently, we wish to find an experiment \mathbf{S} that maximizes the cardinality of the smaller subset in the pair. The result of this optimization is a greedy algorithm that will seek to eliminate the largest number of candidates with each query.

Technical challenge: model count optimization.

While our new formulation is conceptually simple, actually realizing it involves overcoming serious limitations in current theory solvers. This is due to the fact that, while several production solvers provide optimization capabilities [49], these heuristics optimize for the *value* of specific variables. Our requirement is subtly different: we wish to solve for a candidate S that maximizes the *number of satisfying solutions* for the variables M_0, M_1 in Equation (1).³

Unfortunately, this problem is both theoretically and practically challenging. Indeed, merely *counting* the number of satisfying assignments to a constraint formula is known to be asymptotically harder than SAT [69, 70], and practical counting algorithms solutions [14, 20] tend to perform poorly when the combinatorial space is large and the satisfying assignments are sparsely distributed throughout the space, a condition that is likely in our setting. The specific optimization problem our techniques require proves to be even harder. Indeed, only recently was such a problem formalized, under the name Max#SAT [35].

Approximating Max#SAT. While an exact solution to Max#SAT is NP^{PP} -complete [35, 69], several works have explored *approximate* solutions to this and related counting problems [25, 35, 37, 65]. One powerful class of approximate counting techniques, inspired by the theoretical work of Valiant and Vazirani [71] and Stockmeyer [67], uses a SAT oracle as follows: given a constraint formula F over some bitvector T , add to F a series of s random parity constraints, each computed over the bits of T . For $j = 1$ to s , the j^{th} parity constraint can be viewed as requiring that $H_j(T) = 1$ where $H_j : \{0, 1\}^{|T|} \rightarrow \{0, 1\}$ is a universal hash function. Intuitively, each additional constraint reduces the number of satisfying assignments approximately by half, independently of the underlying distribution of valid solutions. The implication is as follows: if a satisfying assignment to the enhanced formula exists, we should be convinced (probabilistically) that the original formula

³Some experimental SMT implementations provide logic for reasoning about the cardinality of small sets, these strategies scale poorly to the large sets we need to reason about in practical format oracle attacks.

is likely to possess on the order of 2^s satisfying assignments. Subsequently, researchers in the model counting community showed that with some refinement, these approximate counting strategies can be used to approximate Max#SAT [35], although with an efficiency that is substantially below what we require for an efficient attack.

To apply this technique efficiently to our attack, we develop a custom count-optimization procedure, and apply it to the attack strategy given in the previous section. At the start of each iteration, we begin by conjecturing a candidate set size 2^s for some non-negative integer s , and then we query the solver for a solution to (S, M_0, M_1) in which approximately 2^s solutions can be found for *each* of the abstract bitvectors M_0, M_1 . This involves modifying the equation of Step (1) by adding s random parity constraints to *each* of the abstract representations of M_0 and M_1 . We now repeatedly query the solver on variants of this query, with increasing (resp. decreasing) values of s , until we have identified the maximum value of s that results in a satisfying assignment.⁴ For a sufficiently high value of s , this approach effectively eliminates many “unprofitable” malleation string candidates and thus significantly improves the efficiency of the attack.

The main weakness of this approach stems from the probabilistic nature of the approximation algorithm. Even when 2^s satisfying assignments exist for M_0, M_1 , the solver may deem the extended formula unsatisfiable with relatively high probability. In our approach, this false-negative will cause the algorithm to reduce the size of s , potentially resulting in the selection of a less-profitable experiment S . Following Gomes *et al.* [37], we are able to substantially improve our certainty by conducting t trials within each query, accepting iff at least $\lceil (\frac{1}{2} + \delta)t \rceil$ trials are satisfied, where δ is an adjustable tolerance parameter.

Putting it all together. The presentation above is intended to provide the reader with a simplified description of our techniques. However, this discussion does not convey most challenging aspect of our work: namely, the difficulty of implementing our techniques and making them practical, particularly within the limitations of existing theory solvers. Achieving the experimental results we present in this work represents the result of months of software engineering effort and manual algorithm optimization. We discuss these challenges more deeply in §4.

Using our techniques we were able to re-discover both well known and entirely novel chosen ciphertext attacks,

⁴Note that $s = 0$ represents the original constraint formula, and so a failure to find a satisfying assignment at this size triggers the conclusion of the attack.

all at a query efficiency nearly identical to the (optimal in expectation) human-implemented attacks. Our experiments not only validate the techniques we describe in this work, but they also illustrate several possible avenues for further optimization, both in our algorithms and in the underlying SMT/SAT solver packages. Our hope is that these results will inspire further advances in the theory solving community.

2 Preliminaries

2.1 Encryption Schemes and Malleability

Our attacks operate assume that the target system is using a malleable symmetric encryption scheme. We now provide definitions for these terms.

Definition 1 (Symmetric encryption) A symmetric encryption scheme Π is a tuple of algorithms $(\text{KeyGen}, \text{Encrypt}, \text{Decrypt})$ where $\text{KeyGen}(1^\lambda)$ generates a key, the probabilistic algorithm $\text{Encrypt}_K(M)$ encrypts a plaintext M under key K to produce a ciphertext C , and the deterministic algorithm $\text{Decrypt}_K(C)$ decrypts C to produce a plaintext or the distinguished error symbol \perp . We use \mathcal{M} to denote the set of valid plaintexts accepted by a scheme, and \mathcal{C} to denote the set of valid ciphertexts.

2.1.1 Malleation Functions

The description of malleation functions is given in the form of two functions. The first takes as input a ciphertext along with an opaque data structure that we refer to as a *malleation instruction string*, and outputs a mauled ciphertext. The second function performs the analogous function on a plaintext. We require that the following intuitive relationship hold between these functions: given a plaintext M and an instruction string, the plaintext malleation function should “predict” the effect of mauling (and subsequently decrypting) a ciphertext that encrypts M .

Definition 2 (Malleation functions) The *malleation functions* for a symmetric encryption scheme Π comprise a pair of efficiently-computable functions $(\text{Maul}_{\text{ciph}}^\Pi, \text{Maul}_{\text{plain}}^\Pi)$ with the following properties. Let \mathcal{M}, \mathcal{C} be the plaintext (resp. ciphertext) space of Π . The function $\text{Maul}_{\text{ciph}}^\Pi : \mathcal{C} \times \{0, 1\}^* \rightarrow \mathcal{C} \cup \{\perp\}$ takes as input a ciphertext and a *malleation instruction string*. It outputs a ciphertext or the distinguished error symbol \perp . The function $\text{Maul}_{\text{plain}}^\Pi : \mathcal{M} \times \{0, 1\}^* \rightarrow \hat{\mathcal{M}}$, on input a plaintext and a malleation instruction string, outputs a

set $\hat{\mathcal{M}} \subseteq \mathcal{M} \cup \{\perp\}$ of possible plaintexts (augmented with the decryption error symbol \perp). The structure of the malleation string is entirely defined by these functions; since our attack algorithms will reason over the functions themselves, we treat S itself as an opaque value.

We say that $(\text{Maul}_{\text{ciph}}^{\Pi}, \text{Maul}_{\text{plain}}^{\Pi})$ describes the malleability features of Π if malleation of a ciphertext always induces the expected effect on a plaintext following encryption, malleation and decryption. More formally, $\forall K \in \text{KeyGen}(1^\lambda), \forall C \in \mathcal{C}, \forall S \in \{0, 1\}^*$ the following relation must hold whenever $\text{Maul}_{\text{ciph}}^{\Pi}(C, S) \neq \perp$:

$$\text{Decrypt}_K(\text{Maul}_{\text{ciph}}^{\Pi}(C, S)) \in \text{Maul}_{\text{plain}}^{\Pi}(\text{Decrypt}_K(C), S)$$

In §4.2.1 we discuss a collection of encryption schemes and implementing their associated malleation functions.

2.2 Theory Solvers and Model Counting

Solvers take as input a system of constraints over a set of variables, and attempt to derive (or rule out the existence of) a satisfying solution. Modern SAT solvers generally rely on two main families of theorem solver: DPLL [28, 29] and Stochastic Local Search [39]. Satisfiability Modulo Theories (SMT) solvers expand the language of SAT to include predicates in first-order logic, enabling the use of several theory solvers ranging from string logic to integer logic. Our prototype implementation uses a quantifier-free bitvector (QFBV) theory solver. In practice, this is implemented using SMT with a SAT solver as a back-end.⁵ For the purposes of describing our algorithms, we specify a query to the solver by the subroutine $\text{SATsolve}\{(A_1, \dots, A_N) : G\}$ where A_1, \dots, A_N each represent abstract bitvectors of some defined length, and G is a constraint formula over these variables. The response from this call provides one of three possible results: (1) sat, as well as a concrete satisfying solution $(\mathbf{A}_1, \dots, \mathbf{A}_N)$, (2) the distinguished response unsat, or (3) the error unknown.

Model counting and Max#SAT. While SAT determines the existence of a single satisfying assignment, a more general variant of the problem, #SAT, determines the *number* of satisfying assignments. In the literature this problem is known as *model counting* [11, 14, 20, 24, 37, 63, 70, 75].

In this work we make use of a specific optimization variant of the model count problem, which was formulated as Max#SAT by Fremont *et al.* [35]. In a streamlined form, the problem can be defined as follows: given a boolean

⁵In principle our attacks can be extended to other theories, with some additional work that we describe later in this section.

formula $\phi(X, Y)$ over abstract bitvectors X and Y , find a concrete assignment to X that *maximizes* the number of possible satisfying assignments to Y .⁶ We will make use of this abstraction in our attacks, with realizations discussed in §3.2. Specifically, we define our main attack algorithm in terms of a generic Max#SAT oracle that has the following interface:

$$\text{Max\#SAT}(\phi, X, Y) \rightarrow \mathbf{X}$$

2.3 Format Checking Functions

Our attacks assume a decryption oracle that, on input a ciphertext C , computes and returns $F(\text{Decrypt}_K(C))$. We refer to the function $F : \mathcal{M} \cup \{\perp\} \rightarrow \{0, 1\}$ as a *format checking function*. Our techniques place two minimum requirements on this function: (1) the function F must be efficiently-computable, and (2) the user must supply a machine-readable implementation of F , expressed as a constraint formula that a theory solver can reason over.

Function descriptions. Requiring format checking functions to be usable within SAT/SMT solvers raises additional implementation considerations. Refer to the full version of this paper [15] for discussion of these considerations, and to the artifact accompanying this work for implemented examples.

3 Constructions

In this section we present a high-level description of our main contribution: a set of algorithms for programmatically conducting a format oracle attack. First, we provide pseudocode for our main attack algorithm, which uses a generic Max#SAT oracle as its key ingredient. This first algorithm can be realized approximately using techniques such as the MaxCount algorithm of Fremont *et al.* [35], although this realization will come at a significant cost to practical performance. To reduce this cost and make our attacks practical, we next describe a concrete replacement algorithm that can be used in place of a Max#SAT solver. The combination of these algorithms forms the basis for our tool Delphinium.

3.1 Main Algorithm

Algorithm 1 presents our main attack algorithm, which we name *DeriveAttack*. This algorithm is parameterized

⁶The formulation of Fremont *et al.* [35] includes an additional set of boolean variables Z that must also be satisfied, but is not part of the optimization problem. We omit this term because it is not used by our algorithms. Note as well that, unlike Fremont *et al.*, our algorithms are not concerned with the actual *count* of solutions for Y .

by three subroutines: (1) a subroutine for solving the Max#SAT problem, (2) an implementation of the ciphertext malleation function $\text{Maul}_{\text{ciph}}$, and (3) a decryption oracle \mathcal{O}_{dec} . The algorithm takes as input a target ciphertext C^* , constraint formulae for the functions $\text{Maul}_{\text{plain}}$, F , and an (optional) initial constraint system G_0 that defines known constraints on M^* .

This algorithm largely follows the intuition described in §1.1. At each iteration, it derives a concrete malleation string \mathbf{S} using the Max#SAT oracle in order to find an assignment that maximizes the number of solutions to the abstract bitvector $M_0 \parallel M_1$. It then mauls C^* using this malleation string, and queries the decryption oracle \mathcal{O}_{dec} on the result. It terminates by outputting a (possibly incomplete) description of M^* . This final output is determined by a helper subroutine `SolveForPlaintext` that uses the solver to find a unique solution for M^* given a constraint formula, or else to produce a human-readable description of the resulting model.⁷

Theorem 3.1 *Given an exact Max#SAT oracle, Algorithm 1 maximizes in expectation the number of candidate plaintext messages ruled out at each iteration.*

A proof of Theorem 3.1 appears in the full version of this paper [15].

Remarks. Note that a greedy adaptive attack may not be globally optimal. It is hypothetically possible to modify the algorithm, allowing it to reason over multiple oracle queries simultaneously (in fact, Phan *et al.* discuss such a generalization in their side channel work [58]). We find that this is computationally infeasible in practice. Finally, note also that our proof assumes an exact Max#SAT oracle. In practice, this will likely be realized with a probably approximately correct instantiation, causing the resulting attack to be a probably approximately greedy attack.

3.2 Realizing the Max#SAT Oracle

Realizing Algorithm 1 in practice requires that we provide a concrete subroutine that can solve specific instances of Max#SAT. We now address techniques for approximately solving this problem.

Realization from Fremont *et al.* Fremont *et al.* [35] propose an approximate algorithm called MaxCount that can be used to instantiate our attack algorithms. The MaxCount algorithm is based on repeated application of approximate counting and sampling algorithms [23, 24, 25], which can in turn be realized using a general SAT solver.

⁷Our concrete implementation in §4 uses the solver to enumerate each of the known and unknown bits of M^* .

While MaxCount is approximate, it can be tuned to provide a high degree of accuracy that is likely to be effective for our attacks. Unfortunately, the Fremont *et al.* solution has two significant downsides. First, to achieve the discussed bounds requires parameter selections which induce infeasible queries to the underlying SAT solver. Fremont *et al.* address this by implementing their algorithm with substantially reduced parameters, for which they demonstrate good empirical performance. However, even the reduced Fremont *et al.* approach still requires numerous calls to a solver. Even conducting a single approximate count of solutions to the constraint systems in our experiments could take hours to days, and such counts might occur several times in a single execution of MaxCount.

A more efficient realization. To improve the efficiency of our implementations, we instead realize a more efficient optimization algorithm we name FastSample. This algorithm can be used in place of the Max#SAT subroutine calls in Algorithm 1. Our algorithm can be viewed as being a subset of the full MaxCount algorithm of Fremont *et al.*

The FastSample algorithm operates over a constraint system $\phi(S, M_0 \parallel M_1)$, and returns a concrete value \mathbf{S} that (heuristically) maximizes the number of solutions for the bitvectors M_0, M_1 . It does this by first conjecturing some value s , and sampling a series of $2s$ low-density parity hash functions of the form $H : \{0, 1\}^n \rightarrow \{0, 1\}$ (where n is the maximum length of M_0 or M_1). It then modifies the constraint system by adding s such hash function constraints to each of M_0, M_1 , and asking the solver to find a solution to the modified constraint system. If a solution is found (resp. not found) for a specific s , FastSample adjusts the size of s upwards (resp. downwards) until it has found the maximal value of s that produces a satisfying assignment, or else is unable to find an assignment even at $s = 0$.

The goal of this approach is to identify a malleation string \mathbf{S} as well as the largest integer s such that at least 2^s solutions can be found for each of M_0, M_1 . To improve the accuracy of this approach, we employ a technique originally pioneered by Gomes *et al.* [37] and modify each SAT query to include multiple *trials* of this form, such that only a fraction $\delta + 1/2$ of the trials must succeed in order for \mathbf{S} to be considered valid. The parameters t, δ are adjustable; we evaluate candidate values in §5.

Unlike Fremont *et al.* (at least, when implemented at full parameters) our algorithm does not constitute a sound realization of a Max#SAT solver. However, empirically we find that our attacks using FastSample produce query counts that are close to the optimal possible attack.

More critically, our approach is capable of identifying a candidate malleation string in seconds on the constraint systems we encountered during our experiments.

Additional algorithms. Our algorithms employ an abstract subroutine `AdjustSize` that is responsible for updating the conjectured set size s in our optimization loop:

$$(b_{\text{continue}}, s', Z') \leftarrow \text{AdjustSize}(b_{\text{success}}, n, s, Z)$$

The input bit b_{success} indicates whether or not a solution was found for a conjectured size s , while n provides a known upper-bound. The history string $Z \in \{0, 1\}^*$ allows the routine to record state between consecutive calls. `AdjustSize` outputs a bit b_{continue} indicating whether the attack should attempt to find a new solution, as well as an updated set size s' . If `AdjustSize` is called with $s = \perp$, then s' is set to an initial set size to test, $b_{\text{continue}} = \text{TRUE}$, and $Z' = Z$.

Finally, the subroutine `ParityConstraint`(n, l) constructs l randomized parity constraints of weight k over a bitvector $b = b_1 b_2 \dots b_n$ where $k \leq n$ denotes the number of bit indices included in a parity constraint (i.e. the parity constraints come from a family of functions $H(b) = \bigoplus_{i=1}^n b_i \cdot a_i$ where $a \in \{0, 1\}^n$ and the hamming weight of a is k).

Algorithm 1: DeriveAttack

Input: Machine-readable description of F , $\text{Maul}_{\text{plain}}$; target ciphertext C^* ; initial constraints G_0 ;
Output: M^* or a model of the remaining plaintext candidates
Procedure:
 $i \leftarrow 1$;
do
 Define $\phi(S, M_0 \| M_1)$ as $[G_{i-1}(M_0) = 1 \wedge G_{i-1}(M_1) = 1 \wedge F(\text{Maul}_{\text{plain}}(M_0, S)) = 0 \wedge F(\text{Maul}_{\text{plain}}(M_1, S)) = 1]$;
 $S \leftarrow \text{Max\#SAT}(\phi, S, M_0 \| M_1)$;
 if $S \neq \perp$ **then**
 $r \leftarrow \mathcal{O}_{\text{dec}}(\text{Maul}_{\text{ciph}}(C^*, S))$;
 Define $G_i(M)$ as
 $[G_{i-1}(M) \wedge (F(\text{Maul}_{\text{plain}}(M, S)) = r)]$;
 $i \leftarrow i + 1$;
while $S \neq \perp$;
return `SolveForPlaintext`(G_i);

Algorithm 2: FastSample

Input: ϕ a constraint system over abstract bitvectors $S, M_0 \| M_1$; n the maximum length of (each of) M_0, M_1 ; m the maximum length of S ; t number of trials; δ fraction of trials that must succeed

Output: $S \in \{0, 1\}^m$

Procedure:

```
(b_continue, s, Z) ← AdjustSize(FALSE, n, ⊥, ε);
// define t symbolic copies of the
// abstract bitvectors M_0, M_1, and a
// new constraint system ϕ'
{M_{1,0}, ..., M_{t,0}} ← M_0;
{M_{1,1}, ..., M_{t,1}} ← M_1;
Define ϕ'(S, {M_{1,0}, ..., M_{t,0}}, {M_{1,1}, ..., M_{t,1}}) as
ϕ(S, M_{1,0} \| M_{1,1}) ∧ ... ∧ ϕ(S, M_{t,0} \| M_{t,1});
while b_continue do
    // Construct 2t s-bit parity
    // constraints
    for i ← 1 to t do
        ℋ_{i,0} ← ParityConstraint(n, s);
        ℋ_{i,1} ← ParityConstraint(n, s)
    // Query the solver
    S ←
    SATSolve({S, {M_{1,0}, ..., M_{t,0}}, {M_{1,1}, ..., M_{t,1}}):
        ∃ℛ_0 ⊆ [1, t]: |ℛ_0| ≥
        [(0.5 + δ)t], ∀j ∈ ℛ_0: ℋ_{j,0}(M_{j,0}) = 1 ∧
        ∃ℛ_1 ⊆ [1, t]: |ℛ_1| ≥
        [(0.5 + δ)t], ∀j ∈ ℛ_1: ℋ_{j,1}(M_{j,1}) = 1 ∧
        ϕ'(S, {M_{1,0}, ..., M_{t,0}}, {M_{1,1}, ..., M_{t,1}})});

    if S == unsat then
        b_success = FALSE;
        (b_continue, s, Z) ←
        AdjustSize(b_success, n, s, Z);
return S
```

4 Prototype Implementation

We now describe our prototype implementation, which we call Delphinium. We designed Delphinium as an extensible toolkit that can be used by practitioners to evaluate and exploit real format oracles.

4.1 Architecture Overview

Figure 3 illustrates the architecture of Delphinium. The software comprises several components:

Attack orchestrator. This central component is respon-

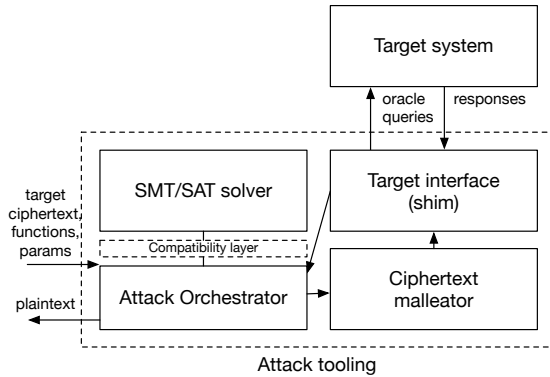


Figure 3: Architecture of Delphinium.

sible for executing the core algorithms of the attack, keeping state, and initiating queries to both the decryption oracle and SMT/SAT solver. It takes the target ciphertext C^* and a description of the functions F and $\text{Maul}_{\text{plain}}$ as well as the attack parameters t, δ as input, and outputs the recovered plaintext.

SMT/SAT solver. Our implementation supports multiple SMT solver frameworks (STP [1] and Z3 [49]) via a custom compatibility layer that we developed for our tool. To improve performance, the orchestrator may launch multiple parallel instances of this solver.

In addition to these core components, the system incorporates two user-supplied modules, which can be customized for a specific target:

Ciphertext malleator. This module provides a working implementation of the malleation function $\text{Maul}_{\text{ciph}}^{\Pi}$. We realize this module as a Python program, but it can be implemented as any executable compatible with the expected interface.⁸

Target interface (shim). This module is responsible for formatting and transmitting decryption queries to the target system. It is designed as a user-supplied module in recognition of the fact that this portion will need to be customized for specific target systems and communication channels.

As part of our prototype implementation, we provide working examples for each of these modules, as well as a test harness to evaluate attacks locally.

4.2 Implementation Details

Realizing our algorithms in a practical tool required us to solve a number of challenging engineering problems

⁸The interface requires input of a ciphertext and a malleation string, with output the mauled ciphertext.

and to navigate limitations of existing SAT/SMT solvers.

Test Harness. For our experiments in §5 we developed a test harness to implement the Ciphertext Malleator and Target Interface shim. This test harness implements the code for mauling and decrypting M^* locally using a given malleation string S .

Selecting SAT and SMT solvers. In the course of this work we evaluated several SMT and SAT solvers optimized for different settings. Seeking the best of a few worlds, we use Z3 for formula manipulation and Crypto-MiniSAT as a solving backend, bridged by CNF formula representations. Refer to the full version of this paper [15] for discussion and challenges of the solvers we evaluated.

Low-density parity constraints. Our implementation of model counting requires our tool to incorporate $2t$ s -bit distinct parity functions into each solver query. Each parity constraint comprises an average of $\frac{n}{2}$ exclusive-ORs (where n is the maximum length of M^*), resulting in a complexity increase of tens to hundreds of gates in our SAT queries. To address this, we adopted an approach used by several previous model counting works [32, 77]: using low-density parity functions. Each such function of these samples k random bits of the input string, with k centered around $\log_2(n)$. As a further optimization, we periodically evaluate the current constraint formula G_i to determine if any bit of the plaintext has been fixed. We omit fixed bits from the input to the parity functions, and reduce both n and k accordingly.

Implementing AdjustSize. Because SAT/SMT queries are computationally expensive, we evaluate a few strategies for implementing AdjustSize which minimize time spent solving. We omit discussion of these strategies for brevity; refer to the full version of this paper [15].

Describing malleation. To avoid making users reimplement basic functionality, Delphinium provides built-in support for several malleation functions. These include simple stream ciphers, stream ciphers that support truncation (from either the left or the right side), and CBC mode encryption. The design of these malleation functions required substantial extensions to the Delphinium framework.

4.2.1 Implementing Malleation Functions

Truncation. Support for truncation requires Delphinium to support plaintexts of variable length. This functionality is not natively provided by the bitvector interfaces used in most solvers. We therefore modify the solver values to encode message length in addition to

content. This necessitates changes to the interface for F . We accomplish this by treating the first $\log_2(n)$ bits of each bitvector as a *length field* specifying how long the message is and by having every implementation of F decode this value prior to evaluating the plaintext. To properly capture truncation off either end of a message, the malleation bitvector is extended by $2\log_2(n)$ so the lowest order $\log_2(n)$ bits of the malleation bitvector specify how many bits should be truncated off the low order bits of the plaintext and the next $\log_2(n)$ bits specify what should be truncated from high order bits of the message. For ease of implementation, in some schemes the n bits following the truncation describe the length field of the plaintext. This allows for easily expressing the exclusive-OR portion of our malleation without bit-shifting and allows encoding extension. Some schemes, such as stream ciphers, only enable truncation off one side of the message, and so in this case we add a constraint to the formula which disallows truncation off the low order bits of a message. This is because truncation off the high order bits would imply a misalignment of the ciphertext with the keystream, causing decryption to produce effectively randomized plaintext.

Truncation for Block Cipher Modes. In block cipher modes such as CTR, CFB, and OFB, an attacker also has the ability to increment the nonce and truncate off blocks of ciphertext.⁹ To capture this capability, in the malleation function we additionally constrain the malleation string to express truncation off the high order bits of a message (earlier blocks of ciphertext), provided the number of bits being truncated is a multiple of the block size.

CBC Mode. In contrast with stream ciphers, $\text{Maul}_{\text{plain}}^{\Pi_{\text{CBC}}}$ is not equal to $\text{Maul}_{\text{ciph}}^{\Pi_{\text{CBC}}}$ and moreover $\text{Maul}_{\text{plain}}^{\Pi_{\text{CBC}}}$ is significantly more complex. In CBC mode, decryption of a ciphertext block C_i is defined as $P_i = \text{Dec}_k(C_i) \oplus C_{i-1}$ where C_{i-1} denotes the previous ciphertext block. Since the block C_i is given directly to a block cipher, any implementation must account for the fact that modification of the block C_i creates an unpredictable effect on the output P_i , effectively randomizing it via the block cipher.

For a solver to reason over such an effect on the plaintext output, we would need to include constraint clauses corresponding to encryption and decryption, i.e. boolean operations implementing symmetric schemes like AES. To avoid this significant overhead, we instead modify the interface of $\text{Maul}_{\text{plain}}^{\Pi_{\text{CBC}}}$ to output two abstract bitvec-

⁹This is not necessarily possible when dealing with other stream ciphers, due to the keystream being misaligned with the ciphertext.

tors (M, Mask) . Mask represents a *mask string*: any bit j where $\text{Mask}[j] = 1$ is viewed as a *wildcard* in the message vector M . When $\text{Mask}[j] = 0$, the value of the output message is equal to $M[j]$ at that position, and when $\text{Mask}[j] = 1$ the value at position $M[j]$ must be viewed as unconstrained. This requires that we modify F to take (M, Mask) as input. The modified F is able to produce a third value in addition to true and false. This new output value indicates that the format check cannot assign a definite true/false value on this input, due to the uncertainty created by the unconstrained bits.¹⁰ Realizing this formulation requires only minor implementation changes to our core algorithms.

Exclusive-OR and Truncation for CBC. With CBC mode decryption, manipulating a preceding ciphertext block C_{i-1} produces a predictable exclusive-OR in the plaintext block P_i . A message that has been encrypted with a block cipher can also be truncated, provided that truncation is done in multiples of the block size. Therefore, we define malleability for CBC to capture (1) block-wise truncation (from either the left or right side of the ciphertext) and (2) exclusive-OR, where exclusive-OR at index i in one block produces the corresponding bit-flip at index i in the next block of decrypted ciphertext.

Supporting Extension. For encryption schemes that allow truncation off the beginning of a message, an attacker may also be able to fill in the truncated portion with arbitrary ciphertext, even if this ciphertext may decrypt to plaintext unknown to them. If the corresponding portion of the plaintext is not examined by the format check function, the attacker can derive information from such queries (if the portion is checked, the attacker can only learn the result of the check over random bits by nature of ciphers). Thus, we create an additional initial constraint for this special case, which allows extension to the ciphertext, limited to where the corresponding plaintext is not examined by the format function.

4.3 Software

Our prototype implementation of *Delphinium* comprises roughly 4.2 kLOC of Python. This includes the attack orchestrator, example format check implementations, the test harness, and our generic solver Python API which allows for modular swapping of backing SMT solvers, with implementations for Z3 and STP provided. In pursuing this prototype, we submitted various patches to the

¹⁰In practice, we implement the output of F as a bitvector of length 2, and modify our algorithms to use 00 and 01 in place of 0 and 1, respectively.

underlying theory solvers that have since been included in the upstream software projects.

4.4 Extensions

In general, arbitrary functions on fixed-size values can be converted into boolean circuits which SMT solvers can reason over. Existing work in MPC develops compilers from DSLs or a subset of C to boolean circuits which could be used to input arbitrary check format functions easily [34, 53]. Experimenting with these, we find that the circuit representations are very large and thus have high runtime overhead when used as constraints. It is possible that circuit synthesis algorithms designed to decrease circuit size (used for applications such as FPGA synthesis) or other logic optimizers could reduce circuit complexity, but we leave exploring this to future work.

We additionally provide a translation tool from the output format of CMBC-GC [34] to Python (entirely comprised of circuit operations) to enable use of the Python front-end to Delphinium.

5 Experiments

5.1 Experimental Setup

To evaluate the performance of Delphinium, we tested our implementation on several multi-core servers using the most up-to-date builds of Z3 (4.8.4) and CryptoMiniSAT (5.6.8). The bulk of our testing was conducted using Amazon EC2, using compute-optimized `c5d.18xlarge` instances with 72 virtual cores and 144GB of RAM.¹¹ Several additional tests were run a 72-core Intel Xeon E5 CPU with 500GB of memory running on Ubuntu 16.04, and a 96-core Intel Xeon E7 CPU with 1TB of memory running Ubuntu 18.04. We refer to these machines as AWS, E5 and E7 in the sections below.

Data collection. For each experimental run, we collected statistics including the total number of decryption oracle queries performed; the wall-clock time required to construct each query; the number of plaintext bits recovered following each query; and the value of s used to construct a given malleation string. We also recorded each malleation string \mathbf{S} produced by our attack, which allows us to “replay” any transcript after the fact. The total number of queries required to complete an attack provides the clearest signal of attack progress, and we use that as the primary metric for evaluation. However, in some

¹¹We also mounted 900GB of ephemeral EC2 storage to each instance as a temporary filesystem to save CNF files during operation.

cases we evaluate partial attacks using the ApproxMC approximate model counting tool [65]. This tool provides us with an estimate for the total number of remaining candidates for M^* at every phase of a given attack, and thus allows evaluation of partial attack transcripts.

Selecting attack parameters. The adjustable parameters in FastSample include t , the number of counting trials, δ , which determines the fraction of trials that must succeed, and the length of the parity constraints used to sample. We ran a number of experiments to determine optimal values for these parameters across the format functions PKCS7 and a bitwise format function defined in §5.2. Empirically, $\delta = 0.5$, $2 \leq t \leq 5$, and parity functions of logarithmic length are suitable for our purposes. Experiments varying t and comparing parity hash function lengths can be found in the full version of this paper [15]. These tests were performed on AWS.

5.2 Experiments with Stream Ciphers

Because the malleation function for stream ciphers is relatively simple (consisting simply of bitwise exclusive-OR), we initiated our experiments with these ciphers.

Bitwise Encryption Padding. The PKCS #7 encryption standard (RFC 2315) [44] defines a padding scheme for use with block cipher modes of operation. This padding is similar to the standard TLS CBC-mode padding [7] considered by Vaudenay [73]. We evaluate our algorithm on both these functions as a benchmark because PKCS7 and its variants are reasonably complex, and because the human-developed attack is well understood. Throughout the rest of this paper, we refer to these schemes as PKCS7 and TLS-PKCS7.

Setup. We conducted an experimental evaluation of the PKCS #7 attack against a 128-bit stream cipher, using parameters $t = 5$, $\delta = 0.5$. Our experiments begin by sampling a random message M^* from the space of all possible PKCS #7 padded messages, and setting $G_0 \leftarrow F_{\text{PKCS7}}$.¹² This evaluation was performed on AWS, E5, and E7.

Results. Our four complete attacks completed in an average of 1699.25 queries (min. 1475, max. 1994) requiring 1.875 hours each (min. 1.63, max. 2.18). A visualization of the resulting attack appears in the full version of this paper [15]. These results compare favorably to the Vaudenay attack, which requires ~2000 queries in expectation, however it is likely that additional tests would find some examples in excess of this average. As points of comparison, attacks with $t = 3$ resulted in a similar number of

¹²In practice, this plaintext distribution tends to produce messages with short padding.

queries (modulo expected variability over different randomly sampled messages) but took roughly 2 to 3 times as long to complete, and attacks with $t = 1$ reached over 5000 queries having only discovered half of the target plaintext message.

Bitwise Padding. To test our attacks, we constructed a simplified *bit* padding scheme F_{bitpad} . This contrived scheme encodes the bit length of the padding P into the rightmost $\lceil \log_2(n) \rceil$ bits of the plaintext string, and then places up to P padding bits directly to the left of this length field, with each padding bit set to 1. We verified the effectiveness of our attacks against this format using a simple stream cipher. Using the parameters $t = 5$, $\delta = 0.5$ the generated attacks took on average 153 queries (min. 137, max. 178). Figure 1 shows one attack transcript at $t = 5$, $\delta = 0.5$. Additional experiments measuring the effect of t on this format are provided in the full version of this paper [15]. These experiments were run primarily on E5.

Negative result: Cyclic Redundancy Checks (CRCs). Cyclic redundancy checks (CRCs) are used in many network protocols for error detection and correction. CRCs are well known to be malleable, due to the linearity of the functions: namely, for a CRC it is always the case that $\text{CRC}(a \oplus b) = \text{CRC}(a) \oplus \text{CRC}(b)$. To test *Delphinium*'s ability to *rule out* attacks against format functions, we implemented a message format consisting of up to three bytes of message, followed by a CRC-8 and a 5-bit message length field. The format function F_{crc8} computes the CRC over the message bytes, and verifies that the CRC in the message matches the computed CRC.¹³ A key feature of this format is that a valid ciphertext C^* should *not* be vulnerable to a format oracle attack using a simple exclusive-OR malleation against this format, for the simple reason that the attacker can predict the output of the decryption oracle for every possible malleation of the ciphertext (due to the linearity of CRC), and thus no information will be learned from executing a query. This intuition was confirmed by our attack algorithm, which immediately reported that no malleation strings could be found. These experiments were performed on E5.

5.3 Ciphers with Truncation

A more powerful malleation capability grants the attacker to arbitrarily *truncate* plaintexts. In some ciphers, this truncation can be conducted from the low-order bits of the plaintext, simply by removing them from the right

¹³In our implementation we used a simple implementation that does not reflect input and output, or add an initial constant value before or after the remainder is calculated.

side of the ciphertext. In other ciphers, such as CTR-mode or CBC-mode, a more limited left-side truncation can be implemented by modifying the IV of a ciphertext. *Delphinium* includes malleation functions that incorporate all three functionalities.

CRC-8 with a truncatable stream cipher. To evaluate how truncation affects the ability of *Delphinium* to find attacks, we conducted a second attack using the function F_{crc8} , this time using an implementation of AES-CTR supporting truncation. Such a scheme may seem contrived, since it involves an encrypted CRC value. However, this very flaw was utilized by Beck and Trew to break WPA [68]. In our experiment, the attack algorithm was able to recover two bytes of the three-byte message, by using the practical strategy of truncating the message and iterating through all possible values of the remaining byte. Additional CRC experiments can be found in the full version of this paper [15]. These experiments were run primarily on E5.

As this example demonstrates, the level of customization and variation in how software developers operate over encrypted data streams can obfuscate the concrete security of an existing implementation. This illustrates the utility of *Delphinium* since such variation's effect on the underlying scheme does not need to be fully understood by a user, outside of encoding the format's basic operation.

Thumb Embedded ISA. To exercise *Delphinium* against a novel format oracle of notably different structure than those traditionally analyzed (such as padding), we implemented a minimal instruction interpreter for the 16-bit Thumb instruction set architecture (ISA), defined as part of the ARM specification [3], capable of emitting illegal instruction signals. Then, operating over stream-cipher encrypted Thumb instructions and using illegal instructions as a boolean signal, *Delphinium* is able to exploit the exclusive-OR malleation to uncover the top seven bits of each 16-bit instruction, in many cases uncovering nine or more (up to 16) bits of each instruction,¹⁴ in an average of ~ 13.3 queries, with each full attack taking only seconds on E5.

Although limited in a few regards, most notably in the simplification of the format oracle into a boolean signal and the assumption that an attacker could be situated in a way that this signal could be gathered, this attack is timely in that it is inspired by the widespread use of unauthenticated encryption in device firmware updates [31].

¹⁴Such a partial firmware decryption generally leaks the instruction opcode, but not its arguments. This could be very useful to an attacker, for example in fuzzy comparison with compiled open source libraries to determine libraries and their versions used in a given firmware update.

If these updates are delivered over-the-air, they may be susceptible to man-in-the-middle attacks enabling such a decryption oracle. Extensive industry research and a current Internet Draft note that unauthenticated firmware updates are an ongoing problem [31, 54].

This initial result serves both as validation of Delphinium and as creation of an avenue for future work, including the development of a model for a more complex but widespread ISA such as 32-bit ARM [3], perhaps exploiting additional signals such as segmentation faults or side channels in order to capture the capabilities of a sophisticated adversary.

S2N with Exclusive-OR and Truncation. To evaluate a realistic attack on a practical format function, we developed a format checking function for the Amazon `s2n` [2] TLS session ticket format. `s2n` uses 60-byte tickets with a 12-byte header comprising a protocol version, cipher-suite version, and format version, along with an 8-byte timestamp that is compared against the current server clock. Although `s2n` uses authenticated encryption (AES-GCM), we consider a hypothetical scenario where nonce re-use has allowed for message forgery [21, 33].

Our experiments recovered the 8-byte time field that a session ticket was issued at: in one attack run, with fewer than 50 queries. However, the attack was unable to obtain the remaining fields from the ticket. This is in part due to some portions of the message being untouched by the format function, and due to the complexity of obtaining a positive result from the oracle when many bytes are unknown. We determined that a full attack against the remaining bytes of the ticket key is possible, but would leave 16 bytes unknown and would require approximately 2^{50} queries. Unsurprisingly, Delphinium timed out on this attack. These experiments were run on AWS and E5.

5.4 CBC mode

We also used the malleation function for CBC-mode encryption. This malleation function supports an arbitrary number of blocks, and admits truncation of plaintexts from either side of the plaintext.¹⁵ The CBC malleation function accepts a structured malleation string S , which can be parsed as (S', l, r) where l, r are integers indicating the number of blocks to truncate from the message.

To test this capability, we used the PKCS7 format function with a blocksize of $B = 16$ bytes, and a two-block CBC plaintext. (This corresponds to a ciphertext consisting of three blocks, including the Initialization Vector.)

¹⁵In practice, truncation in CBC simply removes blocks from either end of the ciphertext.




Figure 4: A contiguous set of malleation queries made by Delphinium during a simulated CBC attack. The rightmost bits signal truncation (from left or right).

Delphinium generated an attack which took 3441 oracle queries for a random message with four bytes of padding. This compares favorably to the Vaudenay attack, which requires 3588 queries in expectation. Interestingly, Delphinium settled on a more or less random strategy of truncation. Where a human attacker would focus on recovering the entire contents of one block before truncating and attacking the next block of plaintext, Delphinium instead truncates more or less as it pleases: in some queries it truncates the message and modifies the Initialization vector to attack the first block. In other queries it focuses on the second block. Figure 4 gives a brief snapshot of this pattern of malleations discovered by Delphinium. Despite this query efficiency (which we seek to optimize, over wall-clock efficiency), the compute time for this attack was almost a week of computation on E5.

6 Related Work

CCA-2 and format oracle attacks. The literature contains an abundance of works on chosen ciphertext and format oracle attacks. Many works consider the problem of constructing and analyzing authenticated encryption modes [17, 61, 62], or analyzing deployed encryption protocols, *e.g.*, [16]. Among many practical format oracle attacks [10, 13, 36, 42, 45, 47, 57, 59, 60, 76], the Lucky13 attacks [5, 7] are notable since they use a *noisy* timing-based side channel.

Automated discovery of cryptographic attacks. Automated attack discovery on systems has been considered in the past. One line of work [26], [58] focuses on generating public input values that lead to maximum leakage of secret input in Java programs where leakage is defined in terms of channel capacity and shannon entropy. Unlike our work, Pasareanu et al. [26] do not consider an

adversary that makes *adaptive* queries based on results of previous oracle replies. Both [26] and [58] assume leakage results from timing and memory usage side channels.

Using solvers for cryptographic tasks/model counting. A wide variety of cryptographic use cases for theory solvers have been considered in the literature. Soos *et al.* [66] developed CryptoMiniSAT to recover state from weak stream ciphers, an application also considered in [27]. Solvers have also been used against hash functions [50], and to obtain cipher key schedules following cold boot attacks [8]. There have been many model counting techniques proposed in the past based on universal hash functions [37, 77]. However, many other techniques have been proposed in the literature. Several works propose sophisticated multi-query approach with high accuracy [25, 65], resulting in the ApproxMC tool we use in our experiments. Other works examine the complexity of parity constraints [77], and optimize the number of variables that must be constrained over to find a satisfying assignment [41].

7 Conclusion

Our work leaves a number of open problems. In particular, we proposed several optimizations that we were not able to implement in our tool, due to time and performance constraints. Additionally, while we demonstrated the viability of our model count optimization techniques through empirical analysis, these techniques require theoretical attention. Our ideas may also be extensible in many ways: for example, developing automated attacks on protocols with side-channel leakage; on public-key encryption; and on “leaky” searchable encryption schemes, *e.g.*, [38]. Most critically, a key contribution of this work is that it poses new challenges for the solver research community, which may result in improvements both to general solver efficiency, as well as to the performance of these attack tools.

References

- [1] The Simple Theorem Prover (STP). Available at <https://stp.github.io/>.
- [2] Introducing s2n, a New Open Source TLS Implementation. <https://aws.amazon.com/blogs/security/introducing-s2n-a-new-open-source-tls-implementation/>, June 2015.
- [3] ARM Architecture Reference Manual, Mar 2018.
- [4] David Adrian, Karthikeyan Bhargavan, Zakir Durumeric, Pierrick Gaudry, Matthew Green, J. Alex Halderman, Nadia Heninger, Drew Springall, Emmanuel Thomé, Luke Valenta, Benjamin VanderSloot, Eric Wustrow, Santiago Zanella-Béguélin, and Paul Zimmermann. Imperfect Forward Secrecy: How Diffie-Hellman Fails in Practice. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security, CCS '15*, pages 5–17, New York, NY, USA, 2015. ACM.
- [5] Martin R. Albrecht and Kenneth G. Paterson. Lucky Microseconds: A Timing Attack on Amazon’s s2n Implementation of TLS. Cryptology ePrint Archive, Report 2015/1129, 2015. <https://eprint.iacr.org/2015/1129>.
- [6] Martin R. Albrecht, Kenneth G. Paterson, and Gaven J. Watson. Plaintext recovery attacks against SSH. In *Proceedings of the 2009 30th IEEE Symposium on Security and Privacy, SP '09*, pages 16–26, Washington, DC, USA, 2009. IEEE Computer Society.
- [7] Nadhem J. AlFardan and Kenneth G. Paterson. Lucky Thirteen: Breaking the TLS and DTLS Record Protocols. In *IEEE S&P (Oakland) '13*, pages 526–540, 2013.
- [8] Abdel Alim Kamal and Amr M. Youssef. Applications of SAT Solvers to AES key Recovery from Decayed Key Schedule Images. *IACR Cryptology ePrint Archive*, 2010:324, 07 2010.
- [9] Nimrod Aviram, Sebastian Schinzel, Juraj Somorovsky, Nadia Heninger, Maik Dankel, Jens Steube, Luke Valenta, David Adrian, J. Alex Halderman, Viktor Dukhovni, Emilia Käseper, Shaanan Cohny, Susanne Engels, Christof Paar, and Yuval Shavitt. DROWN: Breaking TLS using SSLv2. In *25th USENIX Security Symposium (USENIX Security 16)*, pages 689–706, Austin, TX, 2016. USENIX Association.
- [10] Gildas Avoine and Loïc Ferreira. Attacking GlobalPlatform SCP02-compliant Smart Cards Using a Padding Oracle Attack. *IACR Transactions on Cryptographic Hardware and Embedded Systems*, 2018(2):149–170, May 2018.
- [11] Fahiem Bacchus, Shannon Dalmao, and Toniann Pitassi. DPLL with Caching: A new algorithm for #SAT and Bayesian inference. *Electronic Colloquium on Computational Complexity (ECCC)*, 10, 01 2003.
- [12] L. Bang, N. Rosner, and T. Bultan. Online Synthesis of Adaptive Side-Channel Attacks Based On Noisy Observations. In *2018 IEEE European Symposium on Security and Privacy (EuroS P)*, pages 307–322, April 2018.
- [13] Romain Bardou, Riccardo Focardi, Yusuke Kawamoto, Lorenzo Simionato, Graham Steel, and Joe-Kai Tsay. Efficient Padding Oracle Attacks on Cryptographic Hardware. In *CRYPTO '12*, volume 7417 of LNCS, pages 608–625. Springer, 2012.
- [14] Roberto J. Bayardo, Jr., and J. D. Pehoushek. Counting Models using Connected Components. In *In AAAI*, pages 157–162, 2000.
- [15] Gabrielle Beck, Maximilian Zinkus, and Matthew Green. Using SMT Solvers to Automate Chosen Ciphertext Attacks. Cryptology ePrint Archive, Report 2019/958, 2019. <https://eprint.iacr.org/2019/958>.

- [16] Mihir Bellare, Tadayoshi Kohno, and Chanathip Namprempre. Breaking and Provably Repairing the SSH Authenticated Encryption Scheme: A Case Study of the Encode-then-Encrypt-and-MAC Paradigm. *ACM Trans. Inf. Syst. Secur.*, 7(2):206–241, May 2004.
- [17] Mihir Bellare and Chanathip Namprempre. Authenticated Encryption: Relations among Notions and Analysis of the Generic Composition Paradigm. In Tatsuaki Okamoto, editor, *ASIACRYPT 2000*, pages 531–545, Berlin, Heidelberg, 2000. Springer Berlin Heidelberg.
- [18] Steven M. Bellare. Problem Areas for the IP Security Protocols. In *Proceedings of the 6th Conference on USENIX Security Symposium, Focusing on Applications of Cryptography*, volume 6, pages 21–21, Berkeley, CA, USA, 1996. USENIX Association.
- [19] B. Beurdouche, K. Bhargavan, A. Delignat-Lavaud, C. Fournet, M. Kohlweiss, A. Pironi, P. Strub, and J. K. Zinzindohoue. A Messy State of the Union: Taming the Composite State Machines of TLS. In *2015 IEEE Symposium on Security and Privacy*, pages 535–552, May 2015.
- [20] Elazar Birnbaum and Eliezer L. Lozinskii. The Good Old Davis-Putnam Procedure Helps Counting Models. *J. Artif. Int. Res.*, 10(1):457–477, June 1999.
- [21] Hanno Böck, Aaron Zauner, Sean Devlin, Juraj Somorovsky, and Philipp Jovanovic. Nonce-Disrespecting Adversaries: Practical Forgery Attacks on GCM in TLS. In *10th USENIX Workshop on Offensive Technologies (WOOT 16)*, Austin, TX, 2016. USENIX Association.
- [22] Brice Canvel, Alain Hiltgen, Serge Vaudenay, and Martin Vuagnoux. Password interception in a SSL/TLS channel. In Dan Boneh, editor, *Advances in Cryptology - CRYPTO 2003*, pages 583–599, Berlin, Heidelberg, 2003. Springer Berlin Heidelberg.
- [23] Supratik Chakraborty, Daniel J Fremont, Kuldeep S Meel, Sanjit A Seshia, and Moshe Y Vardi. Distribution-aware sampling and weighted model counting for SAT. In *Twenty-Eighth AAAI Conference on Artificial Intelligence*, 2014.
- [24] Supratik Chakraborty, Dror Fried, Kuldeep S Meel, and Moshe Y Vardi. From weighted to unweighted model counting. In *Twenty-Fourth International Joint Conference on Artificial Intelligence*, 2015.
- [25] Supratik Chakraborty, Kuldeep S. Meel, and Moshe Y. Vardi. Algorithmic Improvements in Approximate Counting for Probabilistic Inference: From Linear to Logarithmic SAT Calls. In *Proceedings of International Joint Conference on Artificial Intelligence (IJCAI)*, 7 2016.
- [26] Pasquale Malacaria Corina S. Pasareanu, Quoc-Sang Phan. Multi-run side-channel analysis using Symbolic Execution and Max-SMT. In *2016 IEEE 29th Computer Security Foundations Symposium (CSF)*. IEEE, June 2016.
- [27] Nicolas T. Courtois, Sean O’Neil, and Jean-Jacques Quisquater. Practical Algebraic Attacks on the Hitag2 Stream Cipher. In Pierangela Samarati, Moti Yung, Fabio Martinelli, and Claudio A. Ardagna, editors, *Information Security*, pages 167–176, Berlin, Heidelberg, 2009. Springer Berlin Heidelberg.
- [28] Martin Davis, George Logemann, and Donald Loveland. A Machine Program for Theorem-proving. *Commun. ACM*, 5(7):394–397, July 1962.
- [29] Martin Davis and Hilary Putnam. A Computing Procedure for Quantification Theory. *J. ACM*, 7(3):201–215, July 1960.
- [30] Jean Paul Degabriele and Kenneth G. Paterson. On the (in)security of IPsec in MAC-then-encrypt configurations. In *Proceedings of the 17th ACM Conference on Computer and Communications Security, CCS ’10*, pages 493–504, New York, NY, USA, 2010. ACM.
- [31] Eclipsium. Perilous Peripherals: The Hidden Dangers Inside Windows & Linux Computers, Feb 2020.
- [32] Stefano Ermon, Carla P. Gomes, Ashish Sabharwal, and Bart Selman. Low-Density Parity Constraints for Hashing-Based Discrete Integration, 2014.
- [33] Niels Ferguson. Authentication Weaknesses in GCM, 05 2005.
- [34] Martin Franz, Andreas Holzer, Stefan Katzenbeisser, Christian Schallhart, and Helmut Veith. CBMC-GC: an ANSI C compiler for secure two-party computations. In *International Conference on Compiler Construction*, pages 244–249. Springer, 2014.
- [35] Daniel Fremont, Markus N. Rabe, and Sanjit A. Seshia. Maximum Model Counting. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*, pages 3885–3892, February 2017.
- [36] Christina Garman, Matthew Green, Gabriel Kaptchuk, Ian Miers, and Michael Rushanan. Dancing on the Lip of the Volcano: Chosen Ciphertext Attacks on Apple iMessage. In *25th USENIX Security Symposium (USENIX Security 16)*, pages 655–672, Austin, TX, 2016. USENIX Association.
- [37] Carla P. Gomes, Ashish Sabharwal, and Bart Selman. Model Counting: A New Strategy for Obtaining Good Bounds. In *Proceedings of the 21st National Conference on Artificial Intelligence - Volume 1, AAAI’06*, pages 54–61. AAAI Press, 2006.
- [38] Paul Grubbs, Marie-Sarah Lacharite, Brice Minaud, and Kenneth G. Paterson. Pump Up the Volume: Practical Database Reconstruction from Volume Leakage on Range Queries. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security, CCS ’18*, pages 315–331, New York, NY, USA, 2018. ACM.
- [39] Holger Hoos and Thomas Stztle. *Stochastic Local Search: Foundations & Applications*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2004.
- [40] Troy Hunt. Fear, uncertainty and the padding oracle exploit in ASP.NET. Available at <https://www.troyhunt.com/fear-uncertainty-and-and-padding-oracle/>, September 2010.
- [41] Alexander Ivrii, Sharad Malik, Kuldeep Meel, and Moshe Vardi. On computing minimal independent support and its applications to sampling and counting. *Constraints*, 21, 08 2015.
- [42] Tibor Jager and Juraj Somorovsky. How to Break XML Encryption. In *ACM CCS ’2011*. ACM Press, October 2011.
- [43] Antoine Joux. Authentication failures in NIST version of GCM. Available at https://csrc.nist.gov/csrc/media/projects/block-cipher-techniques/documents/bcm/comments/800-38-series-drafts/gcm/joux_comments.pdf, January 2006.
- [44] Burt Kaliski. PKCS #7: Cryptographic Message Syntax Version 1.5. RFC 2315, March 1998.
- [45] Dennis Kupser, Christian Mainka, Jörg Schwenk, and Juraj Somorovsky. How to Break XML Encryption – Automatically. In *Proceedings of the 9th USENIX Conference on Offensive Technologies, WOOT’15*, Berkeley, CA, USA, 2015. USENIX Association.

- [46] John Mattsson and Magnus Westerlund. Authentication Key Recovery on Galois/Counter Mode GCM. In *Proceedings of the 8th International Conference on Progress in Cryptology — AFRICACRYPT 2016 - Volume 9646*, pages 127–143, Berlin, Heidelberg, 2016. Springer-Verlag.
- [47] Florian Maury, Jean-Rene Reinhard, Olivier Levillain, and Henri Gilbert. Format Oracles on OpenPGP. In *Topics in Cryptology - CT-RSA 2015, The Cryptographer's Track at the RSA Conference 2015*, pages 220–236, San Francisco, United States, April 2015.
- [48] Robert Merget, Juraj Somorovsky, Nimrod Aviram, Craig Young, Janis Fliegenschmidt, Jörg Schwenk, and Yuval Shavitt. Scalable Scanning and Automatic Classification of TLS Padding Oracle Vulnerabilities. In *Proceedings of the 28th USENIX Conference on Security Symposium, SEC'19*, pages 1029–1046, Berkeley, CA, USA, 2019. USENIX Association.
- [49] Microsoft Research. The Z3 Theorem Prover. Available at <https://github.com/Z3Prover/z3>.
- [50] Ilya Mironov and Lintao Zhang. Applications of SAT Solvers to Cryptanalysis of Hash Functions. In *International Conference on Theory and Applications of Satisfiability Testing (SAT 06)*, volume 4121 of *Lecture Notes in Computer Science*, pages 102–115. Springer, August 2006.
- [51] Chris J. Mitchell. Error oracle attacks on CBC mode: Is there a future for CBC mode encryption? In *Information Security, 8th International Conference, ISC 2005, Singapore, September 20-23, 2005, Proceedings*, pages 244–258, 2005.
- [52] Bodo Möller, Thai Duong, and Krzysztof Kotowicz. This POODLE bites: Exploiting the SSLv3 Fallback. Available at <https://www.openssl.org/~bodo/ssl-poodle.pdf>, September 2014.
- [53] Benjamin Mood, Debayan Gupta, Henry Carter, Kevin Butler, and Patrick Traynor. Frigate: A validated, extensible, and efficient compiler and interpreter for secure computation. In *2016 IEEE European Symposium on Security and Privacy (EuroS&P)*, pages 112–127. IEEE, 2016.
- [54] Brendan Moran, Hannes Tschofenig, and Henk Birkholz. An Information Model for Firmware Updates in IoT Devices. Internet-Draft draft-ietf-suit-information-model-05, IETF Secretariat, January 2020. <http://www.ietf.org/internet-drafts/draft-ietf-suit-information-model-05.txt>.
- [55] Jens Müller, Fabian Ising, Vladislav Mladenov, Christian Mainka, Sebastian Schinzel, and Jörg Schwenk. PDF Insecurity. Available at <https://www.pdf-insecurity.org/encryption/encryption.html>, November 2019.
- [56] M. Naor and M. Yung. Public-key Cryptosystems Provably Secure Against Chosen Ciphertext Attacks. In *Proceedings of the Twenty-second Annual ACM Symposium on Theory of Computing, STOC '90*, pages 427–437, New York, NY, USA, 1990. ACM.
- [57] Kenneth G. Paterson and Arnold K. L. Yau. Padding Oracle Attacks on the ISO CBC Mode Encryption Standard. In *Topics in Cryptology - CT-RSA 2004, The Cryptographers' Track at the RSA Conference 2004, San Francisco, CA, USA, February 23-27, 2004, Proceedings*, 2004.
- [58] Quoc-Sang Phan, Lucas Bang, Corina S. Pasareanu, Pasquale Malacaria, and Tevfik Bultan. Synthesis of Adaptive Side-Channel Attacks. In *2017 IEEE 30th Computer Security Foundations Symposium (CSF)*. IEEE, August 2017.
- [59] Damian Poddebniak, Christian Dresen, Jens Müller, Fabian Ising, Sebastian Schinzel, Simon Friedberger, Juraj Somorovsky, and Jörg Schwenk. Efail: Breaking S/MIME and OpenPGP Email Encryption using Exfiltration Channels. In *27th USENIX Security Symposium (USENIX Security 18)*, pages 549–566, Baltimore, MD, 2018. USENIX Association.
- [60] Juliano Rizzo and Thai Duong. Practical Padding Oracle Attacks. In *Proceedings of the 4th USENIX Conference on Offensive Technologies*, pages 1–8, 2010.
- [61] Phillip Rogaway. Authenticated Encryption with Associated Data. In *CCS '02*. ACM Press, 2002.
- [62] Phillip Rogaway and Thomas Shrimpton. A Provable-Security Treatment of the Key-Wrap Problem. In Serge Vaudenay, editor, *EUROCRYPT 2006*, pages 373–390, Berlin, Heidelberg, 2006. Springer Berlin Heidelberg.
- [63] Tian Sang, Fahiem Bacchus, Paul Beam, Henry Kautz, and Toniann Pitassi. Combining Component Caching and Clause Learning for Effective Model Counting. In *SAT 2004*, 05 2004.
- [64] Michael Smith. What you need to know about BEAST. Available at <https://blogs.akamai.com/2012/05/what-you-need-to-know-about-beast.html>, May 2012.
- [65] Mate Soos and Kuldeep S. Meel. BIRD: Engineering an Efficient CNF-XOR SAT Solver and its Applications to Approximate Model Counting. In *Proceedings of AAAI Conference on Artificial Intelligence (AAAI)*, 1 2019.
- [66] Mate Soos, Karsten Nohl, and Claude Castelluccia. Extending SAT Solvers to Cryptographic Problems. In *International Conference on Theory and Applications of Satisfiability Testing (SAT 2009)*, pages 244–257, 06 2009.
- [67] Larry Stockmeyer. The complexity of approximate counting. In *Proceedings of the fifteenth annual ACM symposium on Theory of computing*, pages 118–126. ACM, 1983.
- [68] Erik Tews and Martin Beck. Practical attacks against WEP and WPA. In *Proceedings of the second ACM conference on Wireless network security*, pages 79–86. ACM, 2009.
- [69] Seinosuke Toda. PP is As Hard As the Polynomial-time Hierarchy. *SIAM J. Comput.*, 20(5):865–877, October 1991.
- [70] Leslie G. Valiant. The Complexity of Enumeration and Reliability Problems. *SIAM J. Comput.*, 8(3):410–421, 1979.
- [71] L.G. Valiant and V.V. Vazirani. NP is as easy as detecting unique solutions. *Theoretical Computer Science*, 47:85 – 93, 1986.
- [72] Mathy Vanhoef and Frank Piessens. Key Reinstallation Attacks: Forcing Nonce Reuse in WPA2. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, CCS '17*, pages 1313–1328, New York, NY, USA, 2017. ACM.
- [73] Serge Vaudenay. Security Flaws Induced by CBC Padding - Applications to SSL, IPSEC, WTLS. In *EUROCRYPT '02*, volume 2332 of *LNCS*, pages 534–546, London, UK, 2002. Springer-Verlag.
- [74] W3C. Web Cryptography API. Available at <https://www.w3.org/TR/WebCryptoAPI/>, January 2017.
- [75] Wei Wei and Bart Selman. A New Approach to Model Counting. In Fahiem Bacchus and Toby Walsh, editors, *Theory and Applications of Satisfiability Testing*, pages 324–339, Berlin, Heidelberg, 2005. Springer Berlin Heidelberg.
- [76] Arnold K. L. Yau, Kenneth G. Paterson, and Chris J. Mitchell. Padding Oracle Attacks on CBC-Mode Encryption with Secret and Random IVs. In *Fast Software Encryption: 12th International Workshop, FSE 2005, Paris, France, February 21-23, 2005, Revised Selected Papers*, pages 299–319, 2005.
- [77] Shengjia Zhao, Sorathan Chaturapruek, Ashish Sabharwal, and Stefano Ermon. Closing the Gap Between Short and Long XORs for Model Counting. In *AAAI 2016*, 12 2016.