

Automated Analysis and Synthesis of Block-Cipher Modes of Operation

Alex J. Malozemoff
Department of Computer Science
University of Maryland
amaloz@cs.umd.edu

Jonathan Katz
Department of Computer Science
University of Maryland
jkatz@cs.umd.edu

Matthew D. Green
Department of Computer Science
Johns Hopkins University
mgreen@cs.jhu.edu

Abstract—Block ciphers such as AES are deterministic, keyed functions that operate on small, fixed-size blocks. Block-cipher *modes of operation* define a mechanism for probabilistic encryption of arbitrary length messages using any underlying block cipher. A mode of operation can be proven secure (say, against chosen-plaintext attacks) based on the assumption that the underlying block cipher is a pseudorandom function. Such proofs are complex and error-prone, however, and must be done from scratch whenever a new mode of operation is developed.

We propose an *automated* approach for the security analysis of block-cipher modes of operation based on a “local” analysis of the steps carried out by the mode when handling a *single* message block. We model these steps as a directed, acyclic graph, with nodes corresponding to instructions and edges corresponding to intermediate values. We then introduce a set of *labels* and *constraints* on the edges, and prove a meta-theorem showing that any mode for which there exists a labeling of the edges satisfying these constraints is secure (against chosen-plaintext attacks). This allows us to reduce security of a given mode to a constraint-satisfaction problem, which in turn can be handled using an SMT solver. We couple our security-analysis tool with a routine that automatically generates viable modes; together, these allow us to synthesize hundreds of secure modes.

Index Terms—symmetric-key encryption, modes of operation, program synthesis

I. INTRODUCTION

Designing and proving security of cryptographic constructions can be difficult, time-consuming, and error-prone. Developing a protocol for some task meeting various requirements involves creativity and hard work; even when only small changes to an existing construction are involved, each change typically requires a cryptographic proof of security to be redone from scratch.

To simplify, streamline, and speed up this process, it would be hugely beneficial to *automate* (parts of) this design process. Being able to automate the analysis of cryptographic constructions would enable rapid examination of proposed candidates with less chance of error; coupling this with a systematic way of generating candidates that satisfy a given set of functional requirements would enable automated synthesis of secure constructions. Here we can draw inspiration from the field of *program synthesis* [1], which has recently witnessed several successes (see Srivastava’s Ph.D. thesis [2] for a recent survey).

The past few years have seen the first steps toward automated analysis and synthesis of cryptographic constructions;

we refer the reader to Section I-A for a review of prior work. In this work we focus on automated analysis and synthesis of *block-cipher modes of operation*. A *block cipher* is a deterministic, keyed function F ; with a key k fixed, the function $F_k : \{0, 1\}^n \rightarrow \{0, 1\}^n$ operates on small, n -bit blocks. (A prominent example is the standardized block cipher AES, which has a 128-bit block size.) Secure¹ encryption using a block cipher requires choosing some *mode of operation* to handle both probabilistic encryption (so that the same message, encrypted twice, does not yield the same ciphertext) as well as messages of arbitrary length. Several secure modes of operation were specified as part of the DES standard in the 1970s, and modes such as CTR and CBC are in widespread use today. More recently, various modes of operation with stronger security properties, or based on newer primitives such as tweakable block ciphers, have been proposed; see <http://csrc.nist.gov/groups/ST/toolkit/BCM> for some examples.

The core idea of our approach to the automated analysis of block-cipher modes of operation is to focus on the operations carried out by a given mode when handling a *single* message block. We model these steps as a directed, acyclic graph in which nodes correspond to atomic operations (such as XORing two values together, evaluating a pseudorandom function on some value, etc.) and edges correspond to intermediate values. We introduce a set of labels for the edges in such graphs, along with a set of constraints on how edges can be labeled. Our central result is a meta-theorem stating that if a given graph, corresponding to some mode, can be labeled while satisfying the constraints, then that mode is secure. We thus reduce the security of any mode of operation to a constraint-satisfaction problem, which in turn can be addressed by an SMT solver.

Our meta-theorem shows that our constraints are *sound*, in that any mode classified as secure is, indeed, secure; it is not *complete*, and so may falsely reject a secure mode. Nevertheless, we show that our constraints are permissive enough to capture all commonly used (secure) modes we are aware of, including CBC, CTR, OFB, CFB, and PCBC.

Based on the above meta-theorem, we implemented a model checker designed to evaluate candidate modes of operation.

¹Here we mean the default notion of CPA-security, a.k.a., security against chosen-plaintext attacks.

Our model checker takes as input a mode specified as a graph, and determines both whether the mode is *correct* (i.e., whether there exists a corresponding decryption algorithm that recovers the message from the ciphertext) and whether it is *secure*. We then use our model checker to synthesize secure modes of operation in the natural way. That is, we implement a program synthesizer that automatically generates candidate modes of operation for analysis, and then filters out those modes that are not identified as being secure by our model checker. Doing this in a naïve manner would be prohibitively slow due to the combinatorial explosion in the number of candidate modes. To deal with this, we use several pruning strategies that allow us to eliminate more than 99.99% of potential candidates from consideration before invoking our model checker. Using our approach we are able to synthesize hundreds of secure modes of operation. Our implementation could easily be integrated with a simple tool assigning scores to different modes (e.g., expressing a preference for minimizing calls to the block cipher or for avoiding integer addition) to synthesize secure modes with desired properties.

Although our approach currently only addresses security against chosen-plaintext attacks, we are hopeful that our ideas can be extended to handle stronger security definitions such as authenticated encryption. More generally, we believe our techniques could find application to the automated design and verification of other cryptographic primitives and protocols.

A. Prior Work

Several works have recently considered automated generation and/or analysis of cryptographic algorithms. Akinyele et al. [3], [4] developed tools for automatically generating batch-verification algorithms for digital-signature schemes [3], for converting schemes using symmetric bilinear groups into ones using asymmetric bilinear groups [4], and for compiling signature schemes to schemes achieving stronger security properties [4]. In all these cases, their tools do not directly analyze or verify security of new constructions; instead, they take as input schemes that are assumed to be secure, and then attempt to apply a fixed set of transformations that have been proven to preserve (or amplify) security. Closer to our results are those of Barthe et al. [5], who propose a technique for the automated analysis of *public-key* encryption schemes based on trapdoor permutations and a small set of other instructions. They also use their analysis tool to synthesize secure schemes. The proof techniques and analysis methods are very different in the public-key setting as compared to the symmetric-key setting we consider here.

The work most similar to our own is that of Gagné et al. [6], [7]. Building on earlier work [8], they propose using a compositional Hoare logic to analyze block-cipher modes of operation. They use an imperative language in which each operation updates a distribution over states of the program, and then use their Hoare logic to reason about the output distribution of modes expressed in this language. Roughly, they show that if, after the execution of a mode, all variables are marked as indistinguishable from uniform, then the mode is

secure. The required property can be proven for a given mode by applying their Hoare-logic rules to the entire sequence of operations used by the mode when encrypting some message. A drawback of their approach is that it can only reason about the encryption of messages of some pre-specified, fixed length, and their proofs provide no “inductive” guarantee about what happens when the mode is used to encrypt messages of other (arbitrary) lengths. Thus, to use their approach one would have to provide a separate proof for all possible message lengths one would expect to encounter in practice. Moreover, our “local” approach is fundamentally different from their “global” perspective. In contrast to Gagné et al., we analyze the operations carried out by a given mode when processing a *single* message block; our meta-theorem guarantees that if that (finite) set of operations has a valid labeling, then it corresponds to a mode of operation that is secure for encrypting messages of arbitrary (polynomial) length. Besides this conceptual difference, our approach also ensures that analysis of a given mode (for messages of unbounded length) can be reduced to the satisfiability of a *finite* set of constraints. This, in turn, can be easily solved (in practice) using an SMT solver. Our approach is thus particularly suited for the synthesis of *new* modes of operation.

B. Outline

We introduce some background material in Section II, and in Section III we discuss our approach to modeling block-cipher modes of operation. Section IV contains our central meta-theorem and proof. Section V describes our implementation of the model checker and a program synthesizer that together can be used to automatically generate secure modes of encryption; the results of our experiments using those tools are described in Section VI. We conclude in Section VII.

II. PRELIMINARIES

We use the standard definitions of private-key encryption schemes and block ciphers (i.e., pseudorandom permutations) [9]. We let n denote the block length of a block cipher F ; thus, for a fixed key k , the function F_k defines a bijection on n -bit strings. Formally we treat n as a security parameter, and so “polynomial” (resp., “negligible”) means “polynomial in n ” (resp., “negligible in n ”). A message/ciphertext block is just an n -bit string.

We utilize a notion of security for private-key encryption called IND\$-CPA [10], which is stronger than indistinguishability against chosen-plaintext attacks in that it requires ciphertexts to be indistinguishable from uniform strings.

Definition 1. Let $\Pi = (\text{Enc}, \text{Dec})$ be an encryption scheme in which the encryption of an ℓ -block message yields an $(\ell + 1)$ -block ciphertext, and let $\$(\cdot)$ be an oracle that, when queried on input $m \in \{0, 1\}^{\ell \cdot n}$, returns a uniform string of length $(\ell + 1) \cdot n$. We say Π is IND\$-CPA secure if the following is negligible for all probabilistic polynomial-time algorithms $\mathcal{A}$:

$$\left| \Pr_{k \leftarrow \{0, 1\}^n} [\mathcal{A}^{\text{Enc}_k(\cdot)}(1^n) = 1] - \Pr [\mathcal{A}^{\$(\cdot)}(1^n) = 1] \right|.$$

Since the issue will arise in our analysis, we remark that the decryption algorithm Dec is irrelevant as far as security is concerned (although it is, of course, critical for correctness).

III. MODELING MODES OF OPERATION

A mode of operation refers to a mechanism for encrypting arbitrary length messages based on any underlying block cipher. For simplicity, we assume all messages being encrypted have length a multiple of the block length. (This is without loss of generality, as messages can be padded unambiguously to such a length using known techniques.) For our purposes in this and the next section, a mode is defined by a pair of efficient algorithms Init and Block that each expect oracle access to a function $F_k : \{0, 1\}^n \rightarrow \{0, 1\}^n$. These algorithms have the following functionality:

- **Init:** This is a probabilistic algorithm that, on input 1^n , outputs an initial ciphertext block $c_0 \in \{0, 1\}^n$ and state information $z_0 \in \{0, 1\}^n$.
- **Block:** This is a deterministic algorithm that takes as input a message block $m \in \{0, 1\}^n$ and state information $z \in \{0, 1\}^n$, and outputs a ciphertext block $c \in \{0, 1\}^n$ and updated state information $z' \in \{0, 1\}^n$.

(The above could be generalized so that, e.g., Block outputs two ciphertext blocks for each message block processed, or both Init and Block output two-block state. While this would require modifications to our constraints, our approach would be able to handle such generalizations.)

With some block cipher F fixed, a mode of operation (Init, Block) induces an encryption algorithm Enc. Specifically, if $\mathbf{m} = m_1 \parallel \dots \parallel m_\ell$ is an ℓ -block message, then $\text{Enc}_k(\mathbf{m})$ does the following:

- 1) Compute $(c_0, z_0) \leftarrow \text{Init}(1^n)$, where Init is given oracle access to F_k .
- 2) For $i = 1, \dots, \ell$, set $(c_i, z_i) := \text{Block}(m_i, z_{i-1})$, where Block is given oracle access to F_k .
- 3) Output the ciphertext $c_0 \parallel c_1 \parallel \dots \parallel c_\ell$.

As an example, consider the CBC mode of operation in Figure 1. In CBC mode, the Init algorithm simply generates a uniform initialization vector (IV); the IV is both output as the initial ciphertext block and also passed along as state information. The Block algorithm XORs the incoming state information z with the current message block m , passes the result through a block cipher F_k , and uses the resulting value $F_k(m \oplus z)$ as both the next ciphertext block and the updated state.

In this paper, a mode (Init, Block) is *secure* if the induced encryption scheme is IND\$-CPA secure whenever F is a secure block cipher. (Here, we do not concern ourselves with decryption since we are only interested in security. For a given mode to be meaningful, it will also need to possess a corresponding decryption algorithm. When we synthesize modes, we check for existence of a suitable decryption algorithm in addition to verifying security.)




Fig. 1. CBC mode.

A. Viewing Modes as Graphs

The Init and Block algorithms constituting a mode of encryption can be viewed as directed, acyclic graphs, where the nodes correspond to instructions (such as $\oplus$ or application of F_k), and the edges correspond to (intermediate) values, which are all n -bit strings. A node applies its associated operation on the value(s) contained on its ingoing edge(s), with the result(s) of the operation assigned to its outgoing edge(s). As an example, consider the graph version of CBC mode in Figure 2. The Init portion of the graph corresponds to the generation of a random IV (the GENRAND node) which is then duplicated (the DUP node) and output as part of the ciphertext (through the OUT node) as well as passed on to the Block algorithm (through the NEXTIV node). The Block portion of the graph contains a START node, which is how the state information is provided as input (this information comes from the NEXTIV node of either the Init algorithm or a previous invocation of Block), as well as an M node, which is how the current message block is provided as input. In the case of CBC mode, those input values are XORed and the result is then passed through F_k . That resulting value is duplicated, and again one copy is output as the next block of the ciphertext (via OUT) and another copy is passed on as state (via NEXTIV).

From now on, we view a mode of encryption as being defined by a pair of directed, acyclic graphs, corresponding to the Init and Block algorithms, whose nodes are labeled with instructions from the list below and which also satisfy certain constraints described next. The instructions we support include:

- **DUP:** Has in-degree one and out-degree two. Replicates the ingoing value on both its outgoing edges.
- **GENRAND:** Has in-degree zero and out-degree one. Generates a uniform value in $\{0, 1\}^n$.
- **M:** Has in-degree zero and out-degree one. This corresponds to the current message block being provided as input. This instruction appears exactly once in the Block graph, and does not appear in the Init graph.
- **PRF and PRP:** Have in-degree one and out-degree one. These instructions both apply the block cipher to the value on the ingoing edge and assign the result to the outgoing edge; PRP is invertible but PRF need not be. With regard to security, we rely on the PRF/PRP switch-




Fig. 2. CBC mode expressed as a graph.

ing lemma and treat both PRF and PRP instructions as pseudorandom functions (rather than permutations). For this reason, in the rest of this and the next section we only speak of a PRF instruction. We will later treat PRF and PRP differently with regard to *correctness*.

- **XOR**: Has in-degree two and out-degree one. XORs the values on its two ingoing edges and assigns the result to its outgoing edge.
- **NEXTIV**: Has in-degree one and out-degree zero. (The outgoing edge from NEXTIV in Figure 2 is conceptual.) The value on the ingoing edge is treated as the state that is output by Init or Block. This instruction appears exactly once in each of the Init and Block graphs.
- **START**: Has in-degree zero and out-degree one. (The ingoing edge to START in Figure 2 is conceptual.) The value on the outgoing edge is the state information provided as input to Block. This instruction appears exactly once in the Block graph, and does not appear in the Init graph.
- **OUT**: Has in-degree one and out-degree zero. The value on its ingoing edge represents the next ciphertext block that is output. This instruction appears exactly once in each of the Init and Block graphs.

The above provide a “basic” instruction set that we consider in the main body of this paper. In the appendix we additionally consider the following instruction:

- **INC**: Has in-degree one and out-degree one. Increments the value on its ingoing edge (modulo 2^n), and assigns the result to its outgoing edge.

The above instructions (including INC) are sufficient to model all the standard modes such as ECB, OFB, CFB, CBC, CTR, and PCBC.

B. Edge Labels and Constraints

In the previous section we defined what it means for a pair of directed, acyclic graphs to correspond to a *legal* mode. But not all legal modes are secure! In this section we introduce a set of edge labels, and also describe a set of constraints that a valid labeling must satisfy. Our meta-theorem states that if a legal mode of operation has a valid edge labeling, then that mode is secure.

We now give the details of the edge labeling for some (legal) mode of operation specified by directed, acyclic graphs (Init, Block). Let G denote the union of Init and Block, and let E denote the total number of edges in G . A *label* is a 3-tuple (fam, type, flags) where

- **fam** $\subseteq \{1, \dots, E\}$ represents the set of *families* to which the edge belongs. Two edges e_1, e_2 with families $\text{fam}_1, \text{fam}_2$, respectively, are *related* if $\text{fam}_1 \cap \text{fam}_2 \neq \emptyset$.
- **type** $\in \{R, \perp\}$ represents a “type,” where, intuitively, R denotes “random” and $\perp$ denotes “adversarially controlled.” We impose the ordering $\perp < R$.
- **flags** $\in \{0, 1\}^2$ is a bit-vector denoting whether the edge can be used as input to OUT or PRF, respectively. We index into this bit-vector by using the notation flags.OUT and flags.PRPF . We define the (partial) ordering $\text{flags} \leq \text{flags}'$ iff $\text{flags.OUT} \leq \text{flags}'.OUT$ and $\text{flags.PRPF} \leq \text{flags}'.PRPF$.

In trying to label G , we begin by assigning each edge in G a set of families using the following, deterministic procedure that first handles edges in Init before moving on to the edges in Block:

- 1) While there is an instruction node whose ingoing edges have all been assigned a set of families, assign its outgoing edge(s) a set of families using the following rules:
 - **DUP**: The outgoing edges are both assigned the same set of families as the ingoing edge.
 - **PRF**: Let int be the least unused value in $\{1, \dots, E\}$. The outgoing edge is assigned the set of families $\{\text{int}\}$.
 - **XOR**: Let the ingoing edges have sets of families fam and fam' . The outgoing edge is assigned the set of families $\text{fam} \cup \text{fam}'$.
- 2) Otherwise, we choose an instruction node with in-degree 0, with preference for the START node if its outgoing edge has not yet been labeled. The outgoing edge from this node is assigned $\{\text{int}\}$, where int is the least unused value in $\{1, \dots, E\}$.

Once each edge in G has been assigned a set of families, we then try to assign a **type** and **flags** to each edge in G subject to the following constraints:

- **START**: Say there is a NEXTIV node in G (whether in Init or in Block) whose ingoing edge has type **type** and flags **flags**. Then the outgoing edge from START has $\text{type}' \leq \text{type}$ and $\text{flags}' \leq \text{flags}$.
- **GENRAND**: The outgoing edge gets type R and flags = 11.

- **M**: The outgoing edge gets type $\perp$ and flags = 00.
- **DUP**: Say the input edge has type type and flags flags . The two output edges get types $\text{type}_1 = \text{type}_2 = \text{type}$, and flags flags_1 and flags_2 such that the following hold: (1) $\text{flags}_1 \& \text{flags}_2 = 00$ and (2) $\text{flags}_1 | \text{flags}_2 = \text{flags}$, where $\&$ and $|$ refer to bit-wise AND and OR, respectively.
- **PRF**: We require that the ingoing edge to PRF has type R and flags $\text{PRF} = 1$. The outgoing edge gets type R and flags = 11.
- **XOR**: We require that the ingoing edges to an XOR node are unrelated, and that at least one of them has type R. The outgoing edge gets type R. Say the input edges have flags flags_1 and flags_2 , respectively. Then:
 - 1) If both ingoing edges have type R, the outgoing edge gets flags $\text{flags}_1 | \text{flags}_2$.
 - 2) If only one edge (say, the first) has type R, the outgoing edge gets flags flags_1 .
- **OUT**: We require the ingoing edge to an OUT node to have type R and flags $\text{OUT} = 1$.




Fig. 3. A valid labeling for CBC mode.

A labeling of G that satisfies all of the above constraints is called a *valid labeling*. One can verify that the graphs corresponding to each of CBC, CFB, and OFB modes have valid labelings,² whereas ECB mode (which is known to be

²If we include the INC instruction (and add constraints for this instruction as discussed in the appendix), then CTR mode has a valid labeling also.

insecure) does not. See Figure 3 for a valid labeling of CBC mode.

We now briefly describe some intuition behind these labels and constraints. We define two types for the edge labels: “random” (a.k.a., R) and “adversarially controlled” (a.k.a., $\perp$). For instructions like PRF and XOR, we want a guarantee that the output from these instructions is random. Thus, we do not, for example, allow adversarially controlled values as input into a PRF node, as the output is not “random” to an attacker who can query the PRF node on the same input multiple times.

There are only two instructions that introduce values: GENRAND produces a random value, and M produces an adversarially controlled value (as the adversary supplies the message to be encrypted). For random values, we want to enforce that they are used in ways that do not void their “randomness.” This is where the flags bit-vector comes into play: by accounting for whether a random value can be input into an OUT node or a PRF node, we prevent a random value from both being output as ciphertext *and* input into a PRF node, which can lead to insecure schemes. (This does not mean there do not exist secure schemes which have this property; however, our tool does not allow such schemes.)

In the next section we prove that *any* mode that has a valid labeling is secure. Importantly, determining the existence of a valid labeling for a given mode requires reasoning over a finite graph G and a finite set of constraints, and we can therefore use an SMT solver to check whether a valid labeling is possible.

IV. META-THEOREM AND PROOF

Theorem 1. *Let (Init, Block) be a legal mode of operation. If it has a valid labeling, then it is secure. In other words, when the mode is instantiated with any secure block cipher, the resulting encryption scheme is IND\$-CPA secure.*

Proof: Consider the encryption scheme that results from instantiating Init/Block with a secure block cipher F . The experiment in which a polynomial-time adversary interacts with this scheme proceeds as follows:

- A uniform key for the block cipher is chosen.
- The attacker may then (adaptively) specify messages to be encrypted. After an ℓ -block message $\mathbf{m} = m_1 || \dots || m_\ell$ is specified, we imagine creating a connected graph G consisting of one copy of Init and ℓ copies of Block by simply adding an edge from a NEXTIV node of Init/Block to the START node of the subsequent copy of Block. To encrypt the message, we begin by assigning the i th message block m_i to be the value on the outgoing edge of the M-instruction in the i th copy of Block. We then iteratively assign values to the edges in G (based on the instruction at each node) until every edge is assigned a value. (The outgoing edge from a NEXTIV node or a START node is assigned the same value as the ingoing edge to that node.) This, in particular, assigns values to the ingoing edges to every OUT node, and so defines an $(\ell + 1)$ -block ciphertext that is given to the attacker.

Using a standard hybrid argument, we may replace the block cipher (and choice of uniform key) with a function chosen uniformly from the space of all functions from n -bit inputs to n -bit outputs. This means that PRF instructions in the above process can now be handled as follows: If the value on the ingoing edge to a PRF node was previously input to a PRF node (whether in the course of encrypting the current message, or during encryption of some prior message), then assign to the outgoing edge the same output value used previously. Otherwise, assign to the outgoing edge a uniform value in $\{0,1\}^n$. Once we make this substitution, the entire experiment becomes information-theoretic and we no longer need to be concerned with the running time of the attacker. (However, we continue to assume a polynomial upper bound on the number of messages the attacker can request to be encrypted as well as the block-length of each such message.) As the adversary is computationally unbounded we may, without loss of generality, treat the attacker as deterministic. The probability space of the attacker's interaction is then taken only over choice of the random function and any internal randomness of the mode (i.e., in GENRAND instructions). For each edge in all the graphs G (corresponding to all the messages whose encryption is requested by the attacker), the value assigned to that edge is a random variable. Let $\text{val}(e)$ be the random variable corresponding to the value assigned to some edge e . If E is a set of edges, then $\text{val}(E)$ denotes the multiset $\{\text{val}(e) \mid e \in E\}$. We show that the joint distribution of the random variables corresponding to all the ingoing edges to OUT nodes—and hence the joint distribution of all ciphertext blocks returned to the attacker—is statistically close to uniform. This implies the scheme is IND\$-CPA secure.

Fix some valid labeling of Init and Block. As discussed above, each time the attacker requests an encryption of some ℓ -block message we imagine creating a graph G consisting of one copy of Init and ℓ copies of Block. The valid labeling of Init and Block naturally extends to a labeling of this graph G ; the edges in Init are labeled in the obvious way, and the corresponding edges in each copy of Block are assigned the same label. (The only special case is an edge from a NEXTIV node in one block to a START node in the next block, which we assign the same label as the ingoing edge to that NEXTIV node.)

Consider the step-by-step process by which edges in G are assigned values when a message $\mathbf{m}$ is encrypted. (Although we speak here of “assigning values,” we will continue to treat these as random variables and so “assigning a value” is merely conceptual.) Edges are assigned values in topological order; i.e., the value of an edge is not assigned unless values have been assigned to its parents. We say an edge is *active* if it has been assigned a value but its children have not. (When we begin encrypting a message, the only active edges are the outgoing edges from all the M-nodes.) The only exceptions are ingoing edges to OUT nodes as well as the ingoing edge to the final NEXTIV node; these become active when they are assigned a value, and remain active until the relevant instruction is processed. (Processing an OUT instruction just

means declaring the ingoing edge to that instruction as no longer being active, but instead being “output” as part of the ciphertext. Processing the final NEXTIV node means giving the computed ciphertext to the adversary.) For simplicity, we assume all edges in a given copy of Block are assigned a value before assigning a value to the outgoing edge of the NEXTIV node in that copy of Block.

When we are done encrypting a message we give the resulting ciphertext to the attacker, who then chooses the next message to be encrypted. We then create another graph and assign values as above. We stress that previous graphs are maintained: even though none of their edges are active, we may still refer to the values taken by edges in previous graphs.

At any step in the above process, define the following sets of edges:

- PRF_a = edges of type R with $\text{flags.PRIF} = 1$
that are active
- OUT_a = edges of type R with $\text{flags.OUT} = 1$
that are active
- OUT = edges output for the message currently
being processed

These sets change as instructions are processed. We prove that with all but negligible probability the following invariants hold at the end of each step:

- **Invariant 1:** For any $S \subseteq PRF_a$, the random variables $\text{val}(S)$ are jointly uniform, even conditioned on the values on all active edges unrelated to edges in S and the values of all edges previously used as input to a PRF instruction.
- **Invariant 2:** For any $S \subseteq OUT_a$, the random variables $\text{val}(S)$ are jointly uniform, even conditioned on the values of all previous ciphertext blocks (including those in OUT) and the values on all active edges unrelated to edges in S .
- **Invariant 3:** The random variables $\text{val}(OUT)$ are jointly uniform, even conditioned on the values of all previous ciphertexts.

Invariant 3 implies that when encryption of the given message is done, the resulting ciphertext is (statistically close to) uniform, which is what we wanted to show. The other two invariants are used as part of proving that Invariant 3 is maintained. (Intuitively, Invariant 1 enforces that the inputs into every PRF instruction are uniform, and thus distinct with high probability; this, in turn, means that all the outputs of the PRF instructions are uniform and independent. Invariant 2 enforces that values on ingoing edges to OUT instructions are uniform, and thus when we “output” these values the resulting ciphertext is also uniform.)

Clearly the above invariants all hold when beginning to encrypt some message. (At that point in time, PRF_a , OUT_a , and OUT are all empty.) Now, assume the invariants hold before executing some instruction; we show that, with high probability, they continue to hold after executing the instruction. We let PRF_a , OUT_a , and OUT denote the sets

in question before the instruction, and let $\widehat{PRF}_a$, $\widehat{OUT}_a$, and $\widehat{OUT}$ denote the (possibly modified) sets following the instruction. (Observe that OUT only changes when processing an OUT instruction. Thus, Invariant 3 trivially holds for all other instructions, and we only consider it when we analyze the OUT instruction.)

- **NEXTIV**: The ingoing and outgoing edges have the same label and the same value; thus, the invariants trivially continue to hold.
- **START**: We consider Invariant 1; the argument is identical for Invariant 2.

Let e , labeled $(fam, type, flags)$, be the ingoing edge to this instruction and let e' , labeled $(fam', type', flags')$, be the outgoing edge. The labeling constraints imply that $type' \leq type$ and $flags' \leq flags$. Note further that the only other active edges at the time this instruction is applied are outgoing edges from M nodes, none of which is related to either e or e' , and none of which is in PRF_a or $\widehat{PRF}_a$. If $e' \notin \widehat{PRF}_a$ then Invariant 1 trivially holds after the instruction is applied. Otherwise, the fact that $type' \leq type$ and $flags' \leq flags$ means that $e \in PRF_a$. Since e and e' have the same value and Invariant 1 holds before the instruction is applied, the invariant continues to hold after the instruction is applied.

- **DUP**: We consider Invariant 1; the argument is identical for Invariant 2.

Let e_0 be the ingoing edge to this instruction, and let e_1 and e_2 be the two outgoing edges. If $e_0 \notin PRF_a$, then $e_1, e_2 \notin \widehat{PRF}_a$; this means that $PRF_a = \widehat{PRF}_a$ and—since e_0, e_1 , and e_2 all have the same families—the invariant continues to hold. If $e_0 \in PRF_a$, then exactly one of e_1 or e_2 is in $\widehat{PRF}_a$. Since the values of e_0, e_1 , and e_2 are the same, and all these edges have the same families, the invariant again continues to hold.

- **GENRAND**: Because of the way outgoing edges of a **GENRAND** instruction are labeled, this instruction adds an edge to PRF_a and OUT_a . Since the value on this edge is uniform and independent of any previous values, the invariants continue to hold.
- **PRF**: Using Invariant 1, the value on the ingoing edge to the **PRF** instruction is uniform, even conditioned on the values of all previous ingoing edges to a **PRF** instruction. Therefore, with all but negligible probability, the value on this edge is distinct from all values previously used as input to a **PRF** instruction. When this is the case, the value on the outgoing edge is uniform and both invariants continue to hold just as for a **GENRAND** instruction.
- **XOR**: We consider Invariant 1; the argument is identical for Invariant 2.

Let e_0 and e'_0 , labeled $(fam_0, type_0, flags_0)$ and $(fam'_0, type'_0, flags'_0)$, respectively, be the two ingoing edges, and let e_1 , labeled $(fam_1, type_1, flags_1)$, be the outgoing edge. Note that e_0 and e'_0 must be unrelated and at least one of e_0 and e'_0 must have type R . Also, $fam_1 = fam_0 \cup fam'_0$ and so anything related to e_0 or e'_0

is also related to e_1 .

If $e_1 \notin \widehat{PRF}_a$, then $e_0, e'_0 \notin PRF_a$; thus $\widehat{PRF}_a = PRF_a$ and the invariant continues to hold.

Otherwise, $e_1 \in \widehat{PRF}_a$ and at least one of e_0 or e'_0 (say e_0) is in PRF_a . By Invariant 1 we have that $val(e_0)$ is uniform even conditioned on $val(e'_0)$. Thus, $val(e_1) = val(e_0) \oplus val(e'_0)$ is uniform and the invariant continues to hold.

- **OUT**: Let e be the ingoing edge. This instruction reduces the number of active edges, so clearly Invariant 1 and Invariant 2 continue to hold.

The labeling constraints ensure that $e \in OUT_a$. Also, note that $\widehat{OUT} = OUT \cup \{e\}$. Using Invariant 2, we have that $val(e)$ is uniform even conditioned on $val(OUT)$, and thus Invariant 3 continues to hold.

This completes the proof. $\blacksquare$

It is possible to extract from our proof a bound on the *concrete* security of the mode of operation as a function of the number (and type) of instructions in **Init/Block** as well as the total number of message blocks encrypted.

V. IMPLEMENTATION

We have developed a prototype model checker for verifying whether a given mode has a valid labeling, as well as a synthesizer for generating modes. The code is written in OCaml, and is freely available from the authors. The implementation includes use of the **INC** instruction, with the modified constraints as described in the appendix.

A. Model Checker

The core of our system is a model checker designed to evaluate a candidate encryption mode. Our model checker takes as input a description of the mode of operation (defined in terms of its **Init** and **Block** routines) and determines both whether the corresponding encryption algorithm is *correct* (i.e., whether there exists an efficient decryption algorithm) and whether it is *secure* (i.e., whether the mode is **IND\$-CPA-secure**).

1) *Correctness*: We have developed an algorithm that takes as input a mode of operation, and determines whether decryption of that mode is possible (with knowledge of the key). This is useful for pruning out uninteresting modes. The algorithm works by examining the description of **Block** to see whether each message block can be recovered from the ciphertext. To do this, we proceed as follows:

- If the value on the outgoing edge of the M node in **Block** (i.e., the current message block) can be recovered from the ingoing edge to the **OUT** node (i.e., the current ciphertext block), the mode is decryptable. (This is not typical, but is true for, e.g., **ECB** mode.)
- Otherwise, we check **Block** to see whether the requisite message block can be recovered from both the current ciphertext block and the value on the outgoing edge of the **START** node (i.e., the state passed from one invocation of **Init/Block** to the next). If not, we reject the mode as being undecryptable.

If so, then we need to verify that the state can be recovered from the previous ciphertext block(s). So:

- We check the Init algorithm to see whether the value on the ingoing edge to the NEXTIV node can be recovered from the value on the ingoing edge to the OUT node. If not, we reject the mode as being undecryptable.
- We check the Block algorithm to see whether the value on the ingoing edge to the NEXTIV node can be recovered from the value on the ingoing edge to the OUT node and the value on the outgoing edge of the START node.

Our algorithm is sound, and successfully identifies the standard modes as being decryptable.

To perform each of the above checks, we recursively progress through the graphs defining the Init/Block algorithms in the natural way. We begin, say, with the value on the ingoing edge to the OUT node known. If we know the value on any edge of a DUP node, then we can derive the value on all other edges incident on that node. If we know the values on any *two* edges incident on an XOR node, we can derive the value on the remaining edge. In this way we determine whether it is possible to derive the value on, say, the outgoing edge of the M node.

As an example, consider again the graph representation of CBC mode depicted in Figure 2. From the current ciphertext block (i.e., the value on the ingoing edge of the OUT node), we can recover the value on the ingoing edge to the DUP node and then, using the fact that the PRP node is invertible, derive the value on the outgoing edge of the XOR node. This is not enough, by itself, to recover the message block; however, if we are additionally given the state (i.e., the value on the outgoing edge of the START node), then we can recover the message block. We then need to verify that, in both Init and Block, the state (i.e., the ingoing edge to the NEXTIV node) can be derived from the ciphertext block.

2) *Security*: Theorem 1 states that if a mode has a valid labeling, it is a secure mode of operation. Thus, determining if a mode is secure reduces to a constraint-satisfaction problem in which the goal is to assign a valid label to all the edges in the Init/Block graphs while satisfying the constraints in Section III-B. Viewing the problem in this way lets us use an SMT solver to derive a correct satisfying assignment if one exists. In our implementation, we use the Z3 Theorem Prover³.

We check for a valid labeling in several phases. Because assigning families to edges is deterministic and straightforward, we do this first. After this we check whether two ingoing edges to any XOR node (if one exists in the graph) are related. If there are, a valid labeling cannot exist and we reject the mode as being (potentially) insecure. Otherwise, we proceed to a second phase in which we encode each constraint from Section III-B as an assertion in Z3. That is, we encode each edge of the graph as a variable, and each node as an

assertion which forces the proper labels on the input and output variables of that node. We then run this through Z3, which tells us whether the expression is satisfiable or not given the asserted constraints.

Figure 4 presents an example encoding of a DUP node. Lines 1–6 define the labels for the outgoing edges. We encode the edge’s *type* as an integer, treating $\perp$ as 0 and R as 2. (In the appendix we add a third type which we treat as 1.) Each bit in *flags* is treated as a Boolean value. In the referenced figure, *dup_l_type*, *dup_l_flag_out*, and *dup_l_flag_prf* denote *type*, *flags.OUT*, and *flags.PRF*, respectively, for the left outgoing edge of the DUP node, and likewise, *dup_r_type*, etc., denote the corresponding values for the right outgoing edge. In this example, the ingoing edge to the DUP node comes from GENRAND, whose labeling is denoted by *genrand_type*, etc.

Lines 7–11 define the constraints. As an example, in Line 7 we “assert” that *dup_l_type*, *dup_r_type*, and *genrand_type* are all equal. This equates to the constraint on DUP as defined in Section III-B. Likewise, Lines 8–11 enforce the correct behavior of the *flags* variable for each edge.

B. Program Synthesizer

We also developed a synthesizer for generating secure modes. Given a fixed size N (where by “size” here we mean the number of instructions in Block), one could naïvely iterate over all possible (valid) combinations of N instructions, feeding each resulting mode into our model checker to determine whether the mode is secure or not. Such an approach would be exceedingly slow, as the majority of instruction combinations result in trivially insecure (or undecryptable) schemes. Thus, we implement several simple but effective pruning mechanisms to remove incorrect or uninteresting modes from being analyzed. This pruning process eliminates more than 99.99% of the possible modes from being passed through the model checker, greatly improving the running time of our synthesizer.

We now describe some of our pruning techniques. One approach looks at the current instruction being added to the generated mode, and checks the previous program instruction(s) against a set of disallowed neighboring instructions. One such example is an XOR instruction applied to the two outgoing edges of a DUP instruction, which can never have a valid labeling. Another pruning strategy disallows two back-to-back PRP instructions; while this is not invalid per se, it only produces uninteresting modes (as no security is gained by the second PRP instruction).

VI. RESULTS

We ran our synthesizer to generate modes having Block algorithms containing up to 10 instructions⁴. We do not currently synthesize an Init algorithm, but instead used a

³Available at <https://z3.codeplex.com>.

⁴ In our implementation we actually encode modes using a stack-based language which includes some additional instructions for moving items around the stack. When we say “up to 10 instructions,” we include these extra instructions. However, the instruction count in Table I does *not* include these extra instructions, and thus denotes the actual size of the mode.

```

1 (declare-const dup_l_type Int)
2 (declare-const dup_l_flag_out Bool)
3 (declare-const dup_l_flag_prf Bool)
4 (declare-const dup_r_type Int)
5 (declare-const dup_r_flag_out Bool)
6 (declare-const dup_r_flag_prf Bool)
7 (assert (= dup_l_type dup_r_type genrand_type))
8 (assert (= (and dup_l_flag_out dup_r_flag_out) false))
9 (assert (= (and dup_l_flag_prf dup_r_flag_prf) false))
10 (assert (= (or dup_l_flag_out dup_r_flag_out) genrand_flag_out))
11 (assert (= (or dup_l_flag_prf dup_r_flag_prf) genrand_flag_prf))

```

Fig. 4. Example Z3 encoding of a DUP instruction.

# Instructions	Valid	Decryptable	Secure
1-6	0	0	0
7	78	67	7
8	661	515	24
9	3467	1823	115
10	5136	1187	163
Total	9342	3592	309

TABLE I
NUMBER OF VALID, DECRYPTABLE, AND SECURE MODES BY NUMBER OF INSTRUCTIONS.

fixed Init algorithm which generates a uniform IV and both outputs it and passes it on (as state) to the Block algorithm (cf. Figure 2). Our experiments include the INC instruction, which increments its input modulo 2^n , and the more complex constraints introduced in the appendix to handle this additional instruction.

Table I shows our results. In the context of this table, a “valid” mode is defined as a mode whose acyclic graph representation is connected, contains at least one PRF or PRP node, contains exactly one M, START, OUT, and NEXTIV node, and contains no GENRAND nodes. A “decryptable” mode is a valid mode which passes our decryption check. Finally, a “secure” mode is one that also has a valid labeling. We also include a check for removing duplicate modes; that is, modes that have different graph layouts but are equivalent. (E.g., in one the left outgoing edge of a DUP node is output, and in the other the right outgoing edge of a DUP node is output.) Thus, the results in Table I represent *distinct* nodes.

We discovered 309 secure modes (out of 9342 valid modes). The modes containing seven instructions constitute the well-known CBC, OFB, and CFB modes. The modes containing eight instructions include CTR mode, as well as variants of CBC, OFB, and CFB mode with an additional PRF or PRP instruction introduced. We also synthesize PCBC mode.

In terms of performance, we found that we can synthesize secure modes of operation with ≤ 10 instructions in around 10 minutes on a standard laptop. However, we believe this can be greatly improved. Due to the fact that the Z3 OCaml bindings were not available at the time of this writing, we had to write our Z3 input to a file and then run a separate process to check the result. This additional process creation constitutes a large portion (nearly half) of the running time of

our synthesizer. Porting our approach to use the Z3 OCaml bindings, once available, should greatly improve the overall running time of our tool.

VII. CONCLUSION AND FUTURE WORK

We have introduced a method for reasoning about modes of operation using only “local” analysis of a single block of the given mode. We model modes as graphs and develop a labeling and constraint system on the edges of the graph, and show that if a mode can be correctly labeled then it is secure. Using this meta-theorem, we developed a model checker and synthesizer for both automatically verifying whether a mode is secure and automatically generating new modes. With these tools, we discovered 309 unique secure modes, many of which have never been studied before in the literature.

As future work, we plan to investigate whether it is possible to adapt our approach to the automated analysis and synthesis of authenticated encryption schemes and/or message authentication codes. We hope that the general framework presented in this paper can be applied to these primitives as well as other cryptographic tools.

In terms of extending the present work, it would be useful to couple our approach with automated generation of a proof of security in EasyCrypt [11]. One could also explore adding operations such as concatenation or field multiplication to our language.

ACKNOWLEDGMENTS

Work of Alex Malozemoff was conducted with Government support awarded by DoD, Air Force Office of Scientific Research, National Defense Science and Engineering Graduate (NDSEG) Fellowship, 32 CFR 168a. Work of Jonathan Katz was done for Exelis under contract number N00173-11-C-2045 to NRL. Work of Matthew Green was supported by the U.S. Defense Advanced Research Projects Agency (DARPA) and the Air Force Research Laboratory (AFRL) under contract FA8750-11-2-0211. The authors thank M. Hicks and J. Foster for comments on an earlier draft of this work.

REFERENCES

- [1] Z. Manna and R. Waldinger, “A deductive approach to program synthesis,” *ACM Transactions on Programming Languages and Systems*, vol. 2, no. 1, pp. 90–121, 1980.
- [2] S. Srivastava, “Satisfiability-based program reasoning and program synthesis,” Ph.D. dissertation, University of Maryland, 2010.

- [3] J. A. Akinyele, M. Green, S. Hohenberger, and M. W. Pagano, “Machine-generated algorithms, proofs and software for the batch verification of digital signature schemes,” in *19th ACM Conference on Computer and Communications Security*, T. Yu, G. Danezis, and V. D. Gligor, Eds. ACM Press, 2012, pp. 474–487.
- [4] J. A. Akinyele, M. Green, and S. Hohenberger, “Using SMT solvers to automate design tasks for encryption and signature schemes,” in *20th ACM Conference on Computer and Communications Security*, A.-R. Sadeghi, V. D. Gligor, and M. Yung, Eds. ACM Press, 2013, pp. 399–410.
- [5] G. Barthe, J. M. Crespo, B. Grégoire, C. Kunz, Y. Lakhnech, B. Schmidt, and S. Z. Béguelin, “Fully automated analysis of padding-based encryption in the computational model,” in *20th ACM Conference on Computer and Communications Security*, A.-R. Sadeghi, V. D. Gligor, and M. Yung, Eds. ACM Press, 2013, pp. 1247–1260.
- [6] M. Gagné, P. Lafourcade, Y. Lakhnech, and R. Safavi-Naini, “Automated security proof for symmetric encryption modes,” in *13th Asian Computing Science Conference*, A. Datta, Ed. Springer, 2009, pp. 39–53.
- [7] —, “Automated verification of block cipher modes of operation, an improved method,” in *5th International Symposium on Foundations and Practice of Security*, ser. Lecture Notes in Computer Science, J. García-Alfaro, F. Cuppens, N. Cuppens-Boulahia, A. Miri, and N. Tawbi, Eds., vol. 7743. Springer, 2012, pp. 23–31.
- [8] J. Courant, M. Daubignard, C. Ene, P. Lafourcade, and Y. Lakhnech, “Towards automated proofs for asymmetric encryption schemes in the random oracle model,” in *15th ACM Conference on Computer and Communications Security*, P. Ning, P. F. Syverson, and S. Jha, Eds. ACM Press, 2008, pp. 371–380.
- [9] J. Katz and Y. Lindell, *Introduction to Modern Cryptography*. CRC Press, 2007.
- [10] P. Rogaway, “Nonce-based symmetric encryption,” in *Fast Software Encryption—FSE 2004*, ser. Lecture Notes in Computer Science, B. K. Roy and W. Meier, Eds., vol. 3017. Springer, 2004, pp. 348–359.
- [11] G. Barthe, B. Grégoire, S. Héraud, and S. Z. Béguelin, “Computer-aided security proofs for the working cryptographer,” in *Advances in Cryptology—Crypto 2011*, ser. Lecture Notes in Computer Science, P. Rogaway, Ed., vol. 6841. Springer, 2011, pp. 71–90.

APPENDIX

In this appendix, we modify the labels and constraints to support the INC instruction, and prove a version of Theorem 1 in this setting.

Edge labels. Recall we have a mode of encryption specified by directed, acyclic graphs Init and Block. Let G denote their union, and let E denote the total number of edges in G . Every edge is labeled with a 3-tuple (fam, type, flags) where

- fam is as before.
- type $\in \{R, U, \perp\}$ represents the “type” of the edge. Intuitively, R denotes “random,” U denotes “unique,” and $\perp$ denotes “adversarially controlled.” We impose the ordering $\perp < U < R$.
- flags $\in \{0, 1\}^4$ is a bit-vector. The first three bits denote whether the edge can be used as input to OUT, PRF, or INC, respectively, and we index into this bit-vector using the notation flags.OUT, flags.PRF, and flags.INC. In addition, we have a flag flags.INCd that, intuitively, indicates whether the given edge or one of its related edges ever had flags.INC = 1 (either currently or in the past).

Constraints. In trying to label G , we begin by assigning each edge in G a set of families as in Section III-B, with the following rule for INC nodes:

- INC: The outgoing edge is assigned the same set of families as the incoming edge.

Once each edge in G has been assigned a set of families, we then try to assign a type and flags to each edge in G subject to the following constraints:

- START: Say there is a NEXTIV node in G (whether in Init or in Block) whose ingoing edge has type type and flags flags. Then the outgoing edge from START has type' $\leq$ type, flags' $\leq$ flags, and flags'.INCd = flags.INCd.
- GENRAND: The outgoing edge gets type R and flags.PRF = flags.OUT = 1. The value of flags.INC can be arbitrary, and we set flags.INCd = flags.INC.
- M: The outgoing edge gets type $= \perp$ and flags = 0000.
- DUP: Say the input edge has type type and flags flags. The two output edges get types type₁ = type₂ = type and flags flags₁ and flags₂ such that: (1) flags₁.OUT & flags₂.OUT = 0, and (2) flags₁.OUT | flags₂.OUT = flags.OUT (the rules for flags.PRF and flags.INC are analogous), and (3) flags₁.INCd = flags₂.INCd = flags.INCd.
- INC: We require that the ingoing edge to INC has type $\in \{U, R\}$ and flags.INC = 1. The outgoing edge gets type' = U, flags'.OUT = 0, and flags'.INC = flags'.PRF = flags'.INCd = 1.
- PRF: We require that the ingoing edge to PRF has type $\in \{U, R\}$ and flags.PRF = 1. The outgoing edge is treated just like the output of GENRAND.
- XOR: We require that the ingoing edges to an XOR node are unrelated, at least one of them has type R, and both have flag.INCd = 0. The outgoing edge gets type R. Say the input edges have flags flags₁ and flags₂, respectively. Then:

- 1) If both ingoing edges have type R, the outgoing edge gets flags flags' such that flags'.OUT = flags₁.OUT | flags₂.OUT (the rule for flags'.PRF is analogous), and flags'.INC = flags'.INCd = 0.
- 2) If only one edge (say, the first) has type R, then flags'.OUT = flags₁.OUT and flags'.PRF = flags₁.PRF but flags'.INC = flags'.INCd = 0.

- OUT: We require the ingoing edge to an OUT node to have type R and flags.OUT = 1.

One can verify that the following invariants hold for any labeling satisfying the above constraints:

- **Invariant 1:** If an edge has flags.INC = 1, then it also has flags.INCd = 1.
- **Invariant 2:** If edge e has flags.INCd = 1 and edge e' has flags'.INCd = 0, then e and e' are unrelated.

We rely on these when proving the following theorem.

Theorem 2. *Let (Init, Block) be a legal mode of operation specified using the instructions above (including INC). If the mode has a valid labeling, then it is a secure. In other words, when the mode is instantiated with any secure block cipher, the resulting encryption scheme is IND\$-CPA secure.*

Proof: We assume the reader has read the proof of the

analogous theorem in Section IV, and thus we jump straight to the details.

Fix some valid labeling of Init and Block. Each time the attacker requests encryption of some ℓ -block message we imagine creating a graph G consisting of one copy of Init and ℓ copies of Block. The valid labeling of Init and Block naturally extends to a labeling of this graph G ; the edges in Init are labeled in the obvious way, and the corresponding edges in each copy of Block are assigned the same label. (The only special case is an edge from a NEXTIV node in one block to a START node in the next block, which we assign the same label as the ingoing edge to that NEXTIV node.)

Consider the step-by-step process by which edges in G are assigned values when a message $\mathbf{m}$ is encrypted. (Although we speak here of “assigning values,” these are really random variables and so “assigning a value” is conceptual.) Edges are assigned values in topological order; i.e., the value of an edge is not assigned unless values have been assigned to its parents. We say an edge is *active* if it has been assigned a value but its children have not. (When we begin encrypting a message, the only active edges are the outgoing edges from all the M-nodes.) The only exceptions are ingoing edges to OUT nodes and the ingoing edge to the final NEXTIV node; these become active when they are assigned a value, and remain active until the relevant instruction is processed. (Processing an OUT instruction just means declaring the ingoing edge to that instruction as no longer being active, but instead being “output” as part of the ciphertext. Processing the final NEXTIV node means giving the computed ciphertext to the adversary.) For simplicity, we assume all edges in a given copy of Block are assigned a value before assigning a value to the outgoing edge of the NEXTIV node in that copy of Block.

When we are done encrypting a message we give the resulting ciphertext to the attacker, who then chooses the next message to be encrypted. We then create another graph and assign values as above. We stress that previous graphs are maintained: even though none of their edges are active, we may still refer to the values taken by edges in previous graphs.

At any step in the above process, define the following sets of edges in the graph:

PRF = edges with $\text{type} \in \{U, R\}$ and $\text{flags}.PRF = 1$ that are either active or have entered a PRF node in the past
 PRF^* = active edges with $\text{type} = R$, $\text{flags}.PRF = 1$, and $\text{flags}.INCd = 0$
 INC = active edges with $\text{type} \in \{U, R\}$ and $\text{flags}.INC = 1$
 OUT_a = active edges with $\text{type} = R$ and $\text{flags}.OUT = 1$
 OUT = edges output for the message currently being processed

Let T be a bound on the number of instructions in the graph G defined earlier. (Note that T must be polynomial, since the number of message blocks is assumed to be polynomial.) We

prove that with all but negligible probability the following invariants hold at the end of each step:

- **Invariant 3:** For all $e, e' \in PRF$, $\text{val}(e) \neq \text{val}(e')$. (Formally, $\text{val}(e)$ and $\text{val}(e')$ are random variables, and what we mean here is that with all but negligible probability they take on different values.)
- **Invariant 4:** After executing the i th instruction, for all $e \in INC$ and $e' \in PRF$:

$$\text{val}(e) \notin \{\text{val}(e') - T + i, \dots, \text{val}(e') - 1\}.$$

Moreover, for all distinct $e, e' \in INC$:

$$\text{val}(e) \notin \{\text{val}(e') - T + i, \dots, \text{val}(e')\}.$$

(Formally, $\text{val}(e)$ and $\text{val}(e')$ are random variables, and what we mean formally in the first case—and analogously for the second case—is that with all but negligible probability the value x taken by $\text{val}(e)$ is not in the set $\{x' - T + i, \dots, x' - 1\}$, where x' is the value taken by $\text{val}(e')$.)

- **Invariant 5:** For all $e \in PRF^*$ and any set E of active edges unrelated to e , $\text{val}(e)$ is uniform even conditioned on $\text{val}(E)$ and $\text{val}(PRF \setminus \{e\})$.
- **Invariant 6:** (a) Random variables $\text{val}(OUT)$ are jointly uniform, even conditioned on the values of all previous ciphertext blocks. Moreover, (b) for any $S \subseteq OUT_a$, random variables $\text{val}(S)$ are jointly uniform, even conditioned on the values of all previous ciphertext blocks, $\text{val}(OUT)$, and the values on all active edges unrelated to edges in S .

Note that Invariant 6(a) proves the theorem.

Clearly, the invariants hold at the beginning of the experiment. Now, assume the invariants hold before executing an instruction; we show that, with high probability, they continue to hold after executing the instruction. As in the proof of Theorem 1, we let “unhatted” variables denote the sets in question before executing the instruction, and let “hatted” variables denote the relevant sets following the instruction. We now consider the possible instructions:

- **Begin processing the next message:** Note there are no active edges before this instruction, and the active edges after this instruction are just the set of outgoing edges from M-nodes. $\widehat{PRF} = PRF$ and $\widehat{PRF}^* = \widehat{INC} = \widehat{OUT}_a = \widehat{OUT} = \emptyset$. Invariant 3 holds by the inductive assumption, and the remaining invariants trivially hold.
- **NEXTIV:** The ingoing and outgoing edges have the same label and value; thus, the invariants are not affected.
- **START:** Let e , with type type and flags flags , be the ingoing edge and e' , with type type' and flags flags' , be the outgoing edge. Edges e and e' have the same value but possibly different labels. However, the constraints ensure that $\text{type}' \leq \text{type}$, $\text{flags}' \leq \text{flags}$, and $\text{flags}.INCd = \text{flags}'.INCd$. Thus, $\text{val}(\widehat{PRF}) \subseteq \text{val}(PRF)$ and $\text{val}(\widehat{INC}) \subseteq \text{val}(INC)$, even when viewed as multisets, and thus Invariant 3 and Invariant 4

continue to hold. $OUT = \widehat{OUT}$, so Invariant 6(a) continues to hold.

When processing the **START** instruction, the only active edges—besides e itself—are outgoing edges from M -nodes, which are not in OUT_a or PRF^* and do not have $flags.INC = 1$. It is thus easy to verify the remaining invariants:

- If $e' \notin \widehat{PRF^*}$, then $\widehat{PRF^*} = \emptyset$ and Invariant 5 is trivial; otherwise, $e \in PRF^*$ and by induction the invariant holds.
- If $e' \notin \widehat{OUT_a}$, then $\widehat{OUT_a} = \emptyset$ and Invariant 6(b) is trivial; otherwise, $e \in OUT_a$ and by induction the invariant holds.
- **DUP**: Let e_0 be the ingoing edge and e_1, e_2 the outgoing edges.
 - Invariant 3: If $e_0 \notin PRF$, then $e_1, e_2 \notin \widehat{PRF}$. If $e_0 \in PRF$, then $e_0 \notin \widehat{PRF}$ and exactly one of e_1 or e_2 is in $\widehat{PRF}$. Since $\text{val}(e_0) = \text{val}(e_1) = \text{val}(e_2)$, this means that $\text{val}(PRF) = \text{val}(\widehat{PRF})$ (even viewed as multisets), so Invariant 3 continues to hold.
 - Invariant 4: Arguing as above, $\text{val}(\widehat{INC}) = \text{val}(INC)$ even viewed as multisets. Since $\text{val}(PRF) = \text{val}(\widehat{PRF})$ also, the invariant continues to hold.
 - Invariant 5: Fix $\hat{e} \in \widehat{PRF^*}$ and a set $\hat{E}$ of active edges unrelated to $\hat{e}$. There are two cases to consider:
 - * $\hat{e} \in \{e_1, e_2\}$. Then $e_0 \in PRF^*$. Note that $\text{val}(e_0) = \text{val}(\hat{e})$. Since $\hat{e}$ and e_0 are related, all edges in $\hat{E}$ are unrelated to e_0 , and thus by induction $\text{val}(e_0)$ is uniform even conditioned on $\text{val}(\hat{E})$ and $\text{val}(PRF \setminus \{e_0\})$. Since $\text{val}(e_0) = \text{val}(\hat{e})$, the invariant continues to hold.
 - * $\hat{e} \notin \{e_1, e_2\}$. In this case, $\hat{e} \in PRF^*$. If $\hat{E} \cap \{e_1, e_2\} \neq \emptyset$, then let $\tilde{E}$ be the same as $\hat{E}$ except with e_0 in place of e_1 and/or e_2 ; otherwise, let $\tilde{E} = \hat{E}$. Clearly, $\text{val}(\tilde{E}) = \text{val}(\hat{E})$ and the invariant continues to hold.
 - Invariant 6: We have $OUT = \widehat{OUT}$ and so Invariant 6(a) is trivial. Note that $e_0 \in OUT_a$ iff exactly one of e_1 or e_2 is in $\widehat{OUT_a}$; thus, $\text{val}(OUT_a) = \text{val}(\widehat{OUT_a})$, even when viewed as multisets. Since e_1 and e_2 are related, and e_0, e_1 , and e_2 are all related to the same edges, Invariant 6(b) continues to hold as well.
- **GENRAND**: This instruction adds an edge to PRF and OUT_a , and possibly PRF^* and INC . Since the value of this edge is uniform and independent of any previous edges, and the edge is not related to any other edges, all the invariants continue to hold (with high probability).
- **PRF**: By Invariant 3, the value on the ingoing edge to the PRF instruction is distinct from all values previously used as input to a PRF instruction. Thus, the value on the outgoing edge is uniform, and all the invariants continue to hold (with high probability) just as in the case of a

GENRAND instruction.

- **INC**: Denote the input edge by e_0 and the output edge by e_1 . Note that $e_0 \in INC$, $e_0 \notin PRF^*$ (by Invariant 1) and $e_1 \in \widehat{INC} \cap \widehat{PRF}$, $e_1 \notin PRF^*$. Because Invariant 4 holds before this step, Invariant 3 and Invariant 4 hold after this step.
 - Invariant 5: Note first that $\widehat{PRF^*} = PRF^*$. Fix $e \in \widehat{PRF^*} = PRF^*$ and a set of active edges $\hat{E}$ unrelated to e . Let E denote the active edges unrelated to e . The set $\hat{E} \cup (\widehat{PRF} \setminus \{e\})$ is identical to $E \cup (PRF \setminus \{e\})$ except that e_1 is in the former (since $e_1 \in \widehat{PRF}$ but $e_1 \neq e$) and e_0 is in the latter (since $e_0 \in E$ by Invariant 2). Since e_1 is a deterministic function of e_0 , the invariant continues to hold.
 - Invariant 6: Invariant 6(a) clearly continues to hold. Since $\widehat{OUT_a} \subseteq OUT_a$ and $\text{val}(e_1)$ is a deterministic function of $\text{val}(e_0)$, Invariant 6(b) also continues to hold.
- **XOR**: Let e_0, e'_0 be the two ingoing edges, and e_1 the outgoing edge. We know that e_0, e'_0 are unrelated, $flags_0.INC = flags'_0.INC = 0$, and at least one of e_0, e'_0 (say, e_0) has $type_0 = R$.
 - Invariant 3: If $flags_1.PR = 0$ the invariant trivially continues to hold. Otherwise, there are two cases to consider:
 - * $flags_1.PR = 1$ and $flags_0.PR = 1$: In this case, $e_0 \in PRF^*$. Let E be the set of active edges unrelated to e_0 , and note $e'_0 \in E$. Invariant 5 implies that $\text{val}(e_0)$ is uniform even conditioned on $\text{val}(e'_0)$ and $\text{val}(PRF \setminus \{e_0\})$. Thus $\text{val}(e_1) = \text{val}(e_0) \oplus \text{val}(e'_0)$ is uniform conditioned on $\text{val}(\widehat{PRF} \setminus \{e_1\})$, and so Invariant 3 continues to hold with all but negligible probability.
 - * $flags_1.PR = 1$ and $flags_0.PR = 0$: Here it must be the case that e'_0 has type R and $flags'_0.PR = 1$. An argument as above, swapping e'_0 and e_0 , shows that Invariant 3 continues to hold except with negligible probability.
 - Invariant 4: Note that $\widehat{INC} = INC$. If $flags_1.PR = 0$, the invariant trivially continues to hold. Otherwise, $e_1 \in \widehat{PRF^*}$ and so at least one of e_0 or e'_0 (say e_0) is in PRF^* . Invariant 2 says that every edge in INC is unrelated to e_0 . Arguing as above (using Invariant 5) gives that $\text{val}(e_1)$ is uniform even conditioned on $\text{val}(\widehat{INC})$ and so Invariant 4 continues to hold.
 - Invariant 5: Fix $\hat{e} \in \widehat{PRF^*}$ and a set $\hat{E}$ of active edges unrelated to $\hat{e}$. There are two cases to consider:
 - * $\hat{e} = e_1$. As above, one can show that $\text{val}(\hat{e}) = \text{val}(e_1) = \text{val}(e_0) \oplus \text{val}(e'_0)$ is uniform even conditioned on $\text{val}(\hat{E})$ and $\text{val}(\widehat{PRF} \setminus \{\hat{e}\})$.
 - * $\hat{e} \neq e_1$. Thus, $\hat{e} \in PRF^*$. If $e_1 \in \hat{E}$, then let $\tilde{E}$ be the same as $\hat{E}$ except with e_0 and e'_0 in

place of e_1 ; otherwise, let $\tilde{E} = \hat{E}$. By induction, $\text{val}(\hat{e})$ is uniform even conditioned on $\text{val}(E)$ and $\text{val}(PRF \setminus \{\hat{e}\})$, and thus we conclude that $\text{val}(\hat{e})$ is also uniform even conditioned on $\text{val}(\hat{E})$ and $\text{val}(\widehat{PRF} \setminus \{\hat{e}\})$.

- Invariant 6: Invariant 6(a) clearly holds. As for Invariant 6(b), an argument as above shows the $\text{val}(e_1)$ is uniform even conditioned on all the required values, and so this invariant continues to hold.
- OUT: This instruction reduces the number of active edges, so Invariant 3, Invariant 4, and Invariant 5 trivially continue to hold. Invariant 6 continues to hold because of how OUT_a is defined.

The above prove the theorem. ■