

Subverting Decryption in AEAD^{*}

Marcel Armour¹(✉) and Bertram Poettering²[0000–0001–6525–5141]

¹ Royal Holloway, University of London, Egham, UK

² IBM Research – Zurich, Switzerland

marcel.armour.2017@rhul.ac.uk

poe@zurich.ibm.com

Abstract. This work introduces a new class of Algorithm Substitution Attack (ASA) on Symmetric Encryption Schemes. ASAs were introduced by Bellare, Paterson and Rogaway in light of revelations concerning mass surveillance. An ASA replaces an encryption scheme with a subverted version that aims to reveal information to an adversary engaged in mass surveillance, while remaining undetected by users. Previous work posited that a particular class of AEAD scheme (satisfying certain correctness and uniqueness properties) is resilient against subversion. Many if not all real-world constructions – such as GCM, CCM and OCB – are members of this class. Our results stand in opposition to those prior results. We present a potent ASA that generically applies to *any* AEAD scheme, is undetectable in all previous frameworks and which achieves successful exfiltration of user keys. We give even more efficient *non-generic* attacks against a selection of AEAD implementations that are most used in practice. In contrast to prior work, our new class of attack targets the decryption algorithm rather than encryption. We argue that this attack represents an attractive opportunity for a mass surveillance adversary. Our work serves to refine the ASA model and contributes to a series of papers that raises awareness and understanding about what is possible with ASAs.

Keywords: Algorithm Substitution Attacks · Privacy · Symmetric Encryption · Mass Surveillance

1 Introduction

The Snowden revelations in 2013 exposed that mass surveillance is a reality. They also showed that even sophisticated adversaries with large resources have been unable to break well established cryptographic primitives and hardness assumptions, shifting their focus to circumventing cryptography. Together, these

* The research of Armour was supported by the EPSRC and the UK government as part of the Centre for Doctoral Training in Cyber Security at Royal Holloway, University of London (EP/P009301/1). The research of Poettering was supported by the European Union’s Horizon 2020 project FutureTPM (779391). The full version of this article is available at <https://eprint.iacr.org/2019/987> [3].

two facts suggest that the study of subverted implementations of cryptographic primitives and protocols is a fruitful area of research; Rogaway has gone so far as to call it a moral imperative [23]. The reader is referred to the survey by Schneier et al. [28], which provides a broad overview of subversion of cryptography, with some useful case studies. The idea that an adversary may embed a backdoor or otherwise tamper with the implementation or specification of a cryptographic scheme or primitive predates the Snowden revelations, and was initiated in a line of work by Young and Yung that they named *kleptography* [30,31]. This area of study can be traced back to Simmons’ work on *subliminal channels*, e.g. [29], undertaken in the context of nuclear non-proliferation during the Cold War. In the original conception, kleptography considered a saboteur who designs a cryptographic algorithm whose outputs are computationally indistinguishable from the outputs of an unmodified trusted algorithm. The saboteur’s algorithm should leak private key data through the output of the system, which was achieved using the same principles as Simmons’ earlier subliminal channels.

PRECEDING WORK. Post-Snowden, work in this area was reignited by Bellare, Paterson and Rogaway (BPR) [8], who formalised study of so-called *algorithm substitution attacks* (ASAs) through the specific example of symmetric encryption schemes. In abstract terms, the adversary’s goal in an ASA is to create a subverted implementation of a scheme that breaks some aspect of security (such as IND-CPA) while remaining undetected by the user. There is a tension for ‘Big Brother’ between mounting a successful attack and being detected; clearly an attack that simply replaces the encryption algorithm with one that outputs the messages in plaintext would be devastating yet trivially detectable. BPR stipulate that subverted schemes should at the very least decrypt correctly (according to the unmodified specification) in order to have some measure of resistance to detection, going on to define the success probability of a mass surveillance adversary in carrying out a successful attack, as well as the advantage of a user in detecting that an attack is taking place. BPR [8] demonstrate an attack against randomized schemes that relies on influencing the randomness generated in the course of encryption. Their attack applies to a sub-class of randomized schemes satisfying a property they call ‘coin-injectivity’. Lastly, BPR also establish a positive result that shows that under certain assumptions, it is possible for authenticated encryption schemes to provide resistance against subversion attacks.

Degabriele, Farshim and Poettering (DFP) [12] critiqued the definitions and underlying assumptions of BPR. Their main insight is that perfect decryptability —as mandated by BPR— is a very strong requirement and artificially limits the adversary’s set of available strategies. In practice, a subversion with negligible failure probability should be considered effectively correct.³ As DFP note, decryption failures may happen for reasons other than subverted encryption, and if they occur sporadically may easily go unnoticed. DFP demonstrate how this can be achieved with an input-triggered subversion, where the trigger is some input

³ This is analogous to the fundamental notion in cryptography that a symmetric encryption scheme be considered secure even in the presence of adversaries with negligible advantage.

(message, associated data, nonce, or a combination thereof) that is difficult to guess, making detection practically impossible.

Bellare, Jaeger and Kane (BJK) [6] improved on the attack of BPR, giving an attack which is effective against all randomized schemes. Whereas the attack of BPR is stateful and so vulnerable to detection through state reset, the BJK attack is stateless. BJK furthermore formalised that the desired outcome of an ASA from the point of view of a mass surveillance adversary is successful key recovery.

In concurrent work, we study the effects of subverting the receiver in the setting of message authentication codes [1,2]. Using similar techniques as in the current report, we provide ASAs that result in successful key exfiltration and thus universal forgeries.

CONTRIBUTIONS. Our work continues a line of investigation that serves to raise awareness of what is possible with ASAs, and highlights the importance of work countering subverted implementations. We consider ASAs from a new perspective that leads to results of practical importance. Recall that BPR established a covert channel through ciphertexts by manipulating the randomness generation; their model stipulated perfect decryptability, which resulted in their definitions being fragile. DFP identified this and proposed tolerating a (minimal) compromise of correctness, allowing trigger messages. We note that attacks employing trigger messages appear trivial to plant in formal security abstractions like IND-CPA where the adversary has full control over encrypted messages, associated data, and nonces. In practice, however, it is certainly questionable that adversaries have enough influence on any of the three to conduct DFP style attacks, as messages are chosen in special formats mandated by applications, nonces are implemented via counters, etc. We remove these dependencies, complementing the DFP approach, by attacking from a different angle: leaving perfect correctness intact, we (minimally) limit ciphertext integrity and establish a covert channel through decryption error events. Concretely, we manipulate the decryption algorithm to accept certain bogus ciphertexts. This requires the surveillance adversary to be able to observe whether a decryption implementation outputs a message or rejects the ciphertext. In many practical scenarios this is a mild assumption, for example if a decryption error results in a packet being dropped and automatically retransmitted. Furthermore, a subverted decryption algorithm could go beyond this by e.g. influencing timing information in future messages sent to the network. We conclude that this attack represents an attractive and easy to implement opportunity for a mass surveillance adversary.

Our results stand in opposition to previous work [8,12,6] which proposed subversion resilience of a large class of AEAD schemes to which many if not all real-world constructions such as GCM, CCM and OCB belong, as long as their nonces are generated deterministically via a shared state maintained by both encryptor and decryptor.⁴ The key observation to resolve this apparent contradiction is that previous work has assumed, besides explicitly spelled out

⁴ The members of this class of schemes are deterministic and satisfy certain technical correctness and uniqueness properties.

requirements like uniqueness of ciphertexts and perfect decryptability, implicit notions such as integrity of ciphertexts. In the ASA setting for AEAD where undermining the confidentiality of a scheme is the key goal of an adversary, it seems just as natural to assume that the adversary is also willing to compromise the integrity guarantees as well.

RELATED WORK. We outlined the key publications on ASAs against symmetric encryption schemes above. Other works, briefly described here, consider subversion on different primitives and in different contexts. Berndt and Liskiewicz [9] reunite the fields of cryptography and steganography. Ateniese, Magri and Venturi [4] study ASAs on signature schemes. In a series of work, Russell, Tang, Yung and Zhou [24,25,26,27] consider ASAs on one-way functions, trapdoor one-way functions and key generation as well as defending randomized algorithms against ASAs. Goh, Boneh, Pinkas and Golle [18] show how to add key recovery to the SSL/TLS and SSH protocols. Dodis, Ganesh, Golovnev, Juels and Ristenpart [13] provide a formal treatment of backdooring PRGs, another form of subversion. Armour and Poettering [1,2] study subversion options for message authentication schemes (MAC). Cryptographic reverse firewalls [21,14,20] represent an architecture to counter ASAs via trusted code in network perimeter filters. Fischlin and Mazaheri show how to construct ASA-resistant encryption and signature algorithms given initial access to a trusted base scheme [17]. Fischlin, Janson and Mazaheri [16] show how to immunize (keyed and unkeyed) hash functions against subversion. Bellare, Kane and Rogaway [7] explore using large keys to prevent key exfiltration in the symmetric encryption setting. Bellare and Hoang [5] give public key encryption schemes that defend against the subversion of random number generators.

Camenisch, Drijvers and Lehmann [11] consider Direct Anonymous Attestation (DAA) in the presence of a subverted Trusted Platform Module (TPM). We note that subversion attacks on cryptographic primitives (on DAA, but just as well on message authentication as considered in the present article) manifest a major attack vector in particular against embedded cryptographic hardware modules like TPMs. This is because the main goal of such modules is to serve as a root of trust in exposed devices for which losing system integrity could be fatal. Subverting a TPM can thus have severe implications. As TPMs are widely available today, including for being embedded into virtually every modern PC, subverting them seems to be a promising option to conduct mass surveillance.

STRUCTURE. We first recall (Sect. 2) standard definitions for symmetric encryption schemes and their security. We next give definitions (Sect. 3) that provide a general framework in which to study ASAs. These have been refined and extended from prior work, crucially including the decryption oracle which had been ignored by previous work. Sect. 4 details our new type of attack, together with formal theorems quantifying the ability of an adversary to exfiltrate keys and the ability of the subversion to go undetected. We give two versions of our ASA: one for a passive adversary (the adversarial model considered by previous work), which we extend to a second ASA requiring an active trigger: a modified ciphertext provided to the decryption algorithm. We discuss the results

of a proof-of-concept implementation in Sect. 5. Lastly, Sect. 6 explains how our attacks can be leveraged to compromise the security of popular practical schemes even more effectively, demonstrating how powerful ASAs become when conducted outside the clearly demarcated boundaries of a formal model. Concretely, we give evidence that ASAs against standardized AEAD constructions like GCM or OCB3 can be even more damaging than our attacks from Sect. 4.

2 Notation and Definitions

NOTATION. For a natural number $k \in \mathbb{N}$, we let $[k] = \{0, 1, \dots, k - 1\}$. We refer to an element $x \in \{0, 1\}^*$ as a string, and denote its length by $|x|$. By ε we denote the empty string. The set of strings of length ℓ is denoted $\{0, 1\}^\ell$. In addition we denote by $\perp \notin \{0, 1\}^*$ a reserved special symbol. For $x \in \{0, 1\}^*$, we let $x[i]$ denote the i -th bit of x , with the convention that we count from 0, i.e., we have $x = x[0] \dots x[|x| - 1]$. For two strings x, x' we denote by $x \parallel x'$ their concatenation. If S is a finite set, then $s \leftarrow_{\S} S$ denotes choosing s uniformly at random from S . If \mathcal{A} is a randomized algorithm, we write $y \leftarrow_{\S} \mathcal{A}(x)$ to indicate that it is invoked on input x (and fresh random coins), and the result is assigned to variable y . In security games we write $\mathcal{A}^{\mathcal{O}_1, \dots, \mathcal{O}_c} \implies 1$ to denote the event that the adversary outputs 1 after being given access to the c oracles.

In Appendix A, we recall standard definitions for (length-preserving) pseudo-random functions and permutations.

2.1 Symmetric Encryption

We focus on the likely most widespread and practically useful encryption primitive: Authenticated Encryption with Associated Data (AEAD). We recall standard definitions of (deterministic) nonce-based AEAD, as per [22].

AEAD. A symmetric encryption scheme Π providing authenticated encryption with associated data is a triple of algorithms $(\Pi.\text{Gen}, \Pi.\text{Enc}, \Pi.\text{Dec})$. Associated to Π are two parameters, $\Pi.\text{kl}$ and $\Pi.\text{nl}$, representing the key length and the nonce length. The key generation algorithm $\Pi.\text{Gen}$ is a probabilistic algorithm that takes as input the key length $\Pi.\text{kl}$ and returns a key $k \in \{0, 1\}^{\Pi.\text{kl}}$. Often $\Pi.\text{Gen}$ is taken as the algorithm choosing k uniformly at random from $\{0, 1\}^{\Pi.\text{kl}}$. The encryption algorithm $\Pi.\text{Enc}$ is deterministic and takes key k , message m , associated data d and nonce $n \in \{0, 1\}^{\Pi.\text{nl}}$ to deterministically obtain ciphertext $c \leftarrow \Pi.\text{Enc}(k, m, d; n)$. Decryption algorithm $\Pi.\text{Dec}$ is deterministic and $\Pi.\text{Dec}(k, c, d; n)$ returns either a message m or the special symbol \perp . For simplicity, we assume that $|\Pi.\text{Enc}(k, m, d; n)|$ is an affine function of the form $|m| + \tau$ where τ is some constant associated to the encryption scheme (all practical encryption schemes are of this type). We call τ the *stretch* of the encryption scheme. Lastly, where the context is clear, we drop the prefix Π .

Definition 1. A symmetric encryption scheme Π is said to be δ -correct if for all tuples $(m, d; n)$ it holds that:

$$\Pr [m \neq m' \mid k \leftarrow_{\$} \text{Gen}(kl), c \leftarrow \text{Enc}(k, m, d; n), m' \leftarrow \text{Dec}(k, c, d; n)] \leq \delta.$$

If $\delta = 0$ the scheme is referred to as being perfectly correct.

The classic privacy notion used for AEAD is indistinguishability from random bits under an adaptive chosen-plaintext-and-nonce attack, utilising standard game-based definitions. For the authenticity notion, we consider adversaries that aim to create (strong) forgeries. Security notions are as in [22]. Intuitively, the scheme provides confidentiality if the privacy advantage of any realistic adversary is negligible and authenticity if the forging advantage of any realistic adversary is negligible.

Definition 2. The privacy advantage of an adversary \mathcal{A} is given by

$$\text{Adv}_{\Pi}^{\text{priv}}(\mathcal{A}) = \Pr \left[\mathcal{A}^{\text{Enc}(k, \cdot, \cdot; \cdot)} \implies 1 \mid k \leftarrow_{\$} \text{Gen}(kl) \right] - \Pr \left[\mathcal{A}^{\$(\cdot, \cdot; \cdot)} \implies 1 \right],$$

where the $\$$ oracle returns $c \leftarrow_{\$} \{0, 1\}^{|m|+\tau}$ for any query $\$(m, d; n)$. We assume that \mathcal{A} is nonce-respecting; that is, \mathcal{A} does not make two queries with the same nonce.

Definition 3. The authenticity advantage of an adversary \mathcal{A} is given by

$$\text{Adv}_{\Pi}^{\text{auth}}(\mathcal{A}) = \Pr \left[\mathcal{A}^{\text{Enc}(k, \cdot, \cdot; \cdot), \text{Dec}(k, \cdot, \cdot; \cdot)} \text{ forges} \mid k \leftarrow_{\$} \text{Gen}(kl) \right],$$

where we say that \mathcal{A} forges if it receives any $m' \neq \perp$ from Dec where we require that $(c, d; n)$ is not the result of an encryption query $(m, d; n)$. We assume that \mathcal{A} is nonce-respecting; that is, \mathcal{A} does not make two encryption queries with the same nonce.

3 ASAs on Symmetric Encryption Schemes

We now outline the framework which will allow us to describe our concrete ASAs in Sect. 4. The aim of an ASA is to replace a given (symmetric encryption) scheme with a compromised version; if the original scheme is denoted Π , we write $\tilde{\Pi}$ for its subversion. The attacker may choose to replace one component of the scheme, or multiple. We model the subverted scheme as having an embedded attacker key which is shared with an external (mass surveillance) adversary. This approach was first used by BPR [8]. From the attacker's perspective, the ASA should be undetectable by the user and result in effective surveillance. We formalise these notions as detectability and key recovery. Our definitions are inherited from prior work [8,6,12]. Whereas previous work assumed that only the encryption algorithm might be subverted, we have generalised the definitions to reflect the possibility that any component (one or multiple) of the symmetric

$\text{Game Det}_{\Pi, \tilde{\Pi}}(\mathcal{D})$ $k_A \leftarrow_{\S} \text{A.Gen}$ $b \leftarrow_{\S} \{0, 1\}, b' \leftarrow_{\S} \mathcal{D}^{\mathcal{O}_{\text{Gen}}, \mathcal{O}_{\text{Enc}}, \mathcal{O}_{\text{Dec}}}$ $\text{return } (b = b')$	$\mathcal{O}_{\text{Enc}}(k, m, d; n)$ $\text{if } (b = 1) \text{ then } c \leftarrow \Pi.\text{Enc}(k, m, d; n)$ $\text{else } c \leftarrow \tilde{\Pi}.\text{Enc}(k_A, k, m, d; n)$ $\text{return } c$
$\mathcal{O}_{\text{Gen}}(\text{kl})$ $\text{if } (b = 1) \text{ then } k \leftarrow_{\S} \Pi.\text{Gen}(\text{kl})$ $\text{else } k \leftarrow_{\S} \tilde{\Pi}.\text{Gen}(k_A, \text{kl})$ $\text{return } k$	$\mathcal{O}_{\text{Dec}}(k, c, d; n)$ $\text{if } (b = 1) \text{ then } m \leftarrow \Pi.\text{Dec}(k, c, d; n)$ $\text{else } m \leftarrow \tilde{\Pi}.\text{Dec}(k_A, k, c, d; n)$ $\text{return } m$

Fig. 1. Game to define the detectability advantage of \mathcal{D} with respect to $\tilde{\Pi}, \Pi$.

encryption scheme could be subverted, and adapted to explicitly consider AEAD schemes. We broadly follow the notational choices of BJK [6].

ASA SYNTAX. An algorithm substitution attack A on a scheme Π consists of a triple $(A.\text{Gen}, A.\text{Ext}, \tilde{\Pi})$, where:

1. The attacker key generation algorithm $A.\text{Gen}$ returns an attacker key $k_A \in \{0, 1\}^{\text{A.kl}}$ for some constant $A.\text{kl}$.
2. $\tilde{\Pi} = (\tilde{\Pi}.\text{Gen}, \tilde{\Pi}.\text{Enc}, \tilde{\Pi}.\text{Dec})$ is a *subverted* symmetric encryption scheme.
 - (a) The subverted key generation algorithm $\tilde{\Pi}.\text{Gen}$ is a probabilistic algorithm that takes as input the key length $\tilde{\Pi}.\text{kl}$ and the attacker key k_A , returning a key $k \in \{0, 1\}^{\tilde{\Pi}.\text{kl}}$.
 - (b) The subverted encryption algorithm $\tilde{\Pi}.\text{Enc}$ takes the attacker key k_A , user key k , message m , associated data d and nonce $n \in \{0, 1\}^{\tilde{\Pi}.\text{nl}}$, outputting ciphertext $c \leftarrow \tilde{\Pi}.\text{Enc}(k_A, k, m, d; n)$.
 - (c) The subverted decryption algorithm $\tilde{\Pi}.\text{Dec}(k_A, k, c, d; n)$ returns either a message m or the special symbol \perp .
3. The key extraction algorithm $A.\text{Ext}$ takes as input k_A and has oracle access to both encryption and decryption oracles in the case of an active adversary, or to a transcript of ciphertexts in the case of a passive adversary. These notions are formalised in the key recovery game in Fig. 2. The output of this algorithm is a key $k \in \{0, 1\}^{\tilde{\Pi}.\text{kl}}$.

We require that $\tilde{\Pi}.\text{kl} = \Pi.\text{kl}$ and $\tilde{\Pi}.\text{nl} = \Pi.\text{nl}$, as the subverted algorithm would otherwise be trivially detected. As in previous work, we assume throughout that the key generation is unsubverted, but we retain a syntax that allows for the more general case.

DETECTABILITY. In the formal notion of detectability, we allow a distinguisher \mathcal{D} to interact with subverted encryption, subverted decryption and (for generality) subverted key generation. We assume that the distinguisher has access to its own reference copy of the unsubverted algorithms. It wins if it can distinguish

between the base scheme and the subverted scheme in the game defined in Fig. 1. The detectability advantage of \mathcal{D} with respect to $\Pi, \tilde{\Pi}$ is given by

$$\text{Adv}_{\Pi, \tilde{\Pi}}^{\text{det}}(\mathcal{D}) = 2 \cdot \Pr \left[\text{Det}_{\Pi, \tilde{\Pi}}(\mathcal{D}) \right] - 1.$$

This definition is adapted from strong undetectability of [6]. Notice that (informally) a ‘hard-to-detect’ subversion of a perfectly correct base scheme necessarily satisfies some correctness condition. To see this, suppose that the subversion does not satisfy δ -correctness: it is detectable with probability at least δ .

KEY RECOVERY. Following [6], recovering the user’s secret key is a strong property for an attacker. We give two flavours of the key recovery game, one for passive adversaries **PassiveKR** and one for active adversaries **ActiveKR**, as given in Fig. 2. In the passive case, we allow the adversary to observe ciphertexts and whether they are rejected. This is formalised through the transcript oracle $\mathcal{O}_{\text{Trans}}$. For the active case, we allow the attacker to generate valid ciphertexts via \mathcal{O}_{Enc} and interact with a decryption oracle \mathcal{O}_{Dec} that reveals whether a submitted ciphertext is rejected. Both games are parametrised by a message sampler algorithm \mathcal{M} . Given its current state σ , \mathcal{M} returns the next message with associated data (m, d) to be encrypted, together with a nonce $n \in \{0, 1\}^{\text{II.nl}}$ and an updated state. It represents the choice of messages made by the sender. For simplicity, we model \mathcal{M} as non-adaptive and nonce-respecting. It could be argued that a more realistic model might take into account that the adversary could influence the user’s choice of messages to be encrypted. However, in constructing attacks we assume the weakest properties of the attacker.

Adversary A wins if $A.\text{Ext}$ recovers the user’s key k after interacting with the subverted encryption scheme. The key recovery advantage of A with respect to $\tilde{\Pi}$ and \mathcal{M} is given by

$$\text{Adv}_{\tilde{\Pi}, \mathcal{M}}^{\text{kr}}(A) = \Pr \left[\text{KR}_{\tilde{\Pi}, \mathcal{M}}(A) \right],$$

where $\text{KR}_{\tilde{\Pi}, \mathcal{M}}(A)$ refers to the appropriate key recovery game according to whether the adversary is passive or active.

4 Mounting Attacks via Decryption Subversion

We now detail our ASAs, first for a passive surveillance adversary and then in the active case. It is easy to see that the attacks are undetectable according to the models in the literature [8,6,12], as the encryption algorithm is not subverted.

Imagine that Alice communicates with Bob. A passive adversary can observe ciphertexts from Alice to Bob. In addition, an active adversary can replace ciphertexts in transmission and submit its own (forged) ciphertexts to Bob. In the passive attack, the decryption algorithm is subverted so that it rejects a fraction of valid ciphertexts, bounded by an attacker controlled parameter. In the active attack, the decryption algorithm is subverted so that it accepts a (similarly bounded) fraction of invalid ciphertexts. The active attack requires

Game $\text{ActiveKR}_{\tilde{\Pi}, \mathcal{M}}$	Game $\text{PassiveKR}_{\tilde{\Pi}, \mathcal{M}}$
$k_A \leftarrow_{\S} \text{A.Gen}$ $k \leftarrow_{\S} \tilde{\Pi}.\text{Gen}(\tilde{\Pi}.\text{kl}), \sigma \leftarrow \varepsilon$ $k' \leftarrow_{\S} \text{A.Ext}^{\mathcal{O}_{\text{Enc}}, \mathcal{O}_{\text{Dec}}}(k_A)$ return $(k' = k)$	$k_A \leftarrow_{\S} \text{A.Gen}$ $k \leftarrow_{\S} \tilde{\Pi}.\text{Gen}(\tilde{\Pi}.\text{kl}), \sigma \leftarrow \varepsilon$ $k' \leftarrow_{\S} \text{A.Ext}^{\mathcal{O}_{\text{Trans}}}(k_A)$ return $(k' = k)$
$\mathcal{O}_{\text{Enc}}()$ $(m, d, n, \sigma) \leftarrow_{\S} \mathcal{M}(\sigma)$ $c \leftarrow \tilde{\Pi}.\text{Enc}(k_A, k, m, d; n)$ return (c, d, n)	$\mathcal{O}_{\text{Trans}}()$ $(m, d, n, \sigma) \leftarrow_{\S} \mathcal{M}(\sigma)$ $c \leftarrow \tilde{\Pi}.\text{Enc}(k_A, k, m, d; n)$ $m \leftarrow \tilde{\Pi}.\text{Dec}(k_A, k, c, d; n)$ return $(c, d, n, (m = \perp))$
$\mathcal{O}_{\text{Dec}}(c, d; n)$ $m \leftarrow \tilde{\Pi}.\text{Dec}(k_A, k, c, d; n)$ return $(m = \perp)$	

Fig. 2. Game to define the key recovery advantage of A with respect to $\tilde{\Pi}$ and \mathcal{M} .

the adversary to send Bob bogus ciphertexts (derived from genuine ciphertexts) that reveal Bob's secret key using decryption errors. Normally, these bogus ciphertexts are unlikely to decrypt correctly, i.e., they would be rejected. In both cases, if the decryptor is subverted then either real ciphertexts (in the passive case) or bogus ciphertexts (in the active case) can either be accepted or rejected, creating via the acceptance/rejection pattern a covert channel that will allow the key to be exfiltrated.

From the point of view of a mass surveillance adversary this is an attractive prospect: having passively collected all communications, triggered by some suspicion they can now target Alice and Bob's communication. By recovering Bob's key they may now decrypt all of the stored communication between Alice and Bob (and indeed from Bob to Alice as well).

We note that both of our attacks are stateless, which not only allows for much easier backdoor implementation from a technical perspective but also should decrease the likelihood that an implemented attack is detected through code review or observing memory usage.

4.1 Attack 1: Passive

Consider the following subversion of a given symmetric encryption scheme $(\Pi.\text{Gen}, \Pi.\text{Enc}, \Pi.\text{Dec})$. Let $\tilde{\Pi}.\text{Gen} = \Pi.\text{Gen}$ and $\tilde{\Pi}.\text{Enc} = \Pi.\text{Enc}$. Let A.Gen choose a key k_A by $k_A \leftarrow_{\S} \{0, 1\}^{\text{A.kl}}$. Algorithms $\tilde{\Pi}.\text{Dec}$ and A.Ext are then specified in Fig. 3. The subverted decryptor $\tilde{\Pi}.\text{Dec}$ takes the same input as $\Pi.\text{Dec}$ together with the attacker key, and utilises a pseudo-random function⁵ F with $F: \{0, 1\}^{\text{A.kl}} \times \{0, 1\}^* \rightarrow [\Pi.\text{kl}] \times \{0, 1\}$. In A.Ext , we use the symbol \star as a ternary

⁵ See Appendix A for definitions of pseudo-random functions and length-preserving pseudo-random permutations.

Algorithm $\tilde{\Pi}.\text{Dec}(k_A, k, c, d; n)$	Algorithm $\text{A.Ext}^{\mathcal{O}_{\text{Trans}}}(k_A)$
1: parse $F(k_A, c \parallel d)$ as $i \parallel b$	1: $\forall i \in [\Pi.\text{kl}], \text{key}[i] \leftarrow \star$
2: if $k[i] = b$ and $B(\delta)$ then	2: while $\exists j: \text{key}[j] = \star$ do
3: return \perp	3: $(c, d, v) \leftarrow \mathcal{O}_{\text{Trans}}$
4: else	4: if $(v = 1)$ then
5: return $\Pi.\text{Dec}(k, c, d; n)$	5: parse $F(k_A, c \parallel d)$ as $i \parallel b$
	6: if $(\text{key}[i] = \star)$ then
	7: $\text{key}[i] \leftarrow b$
	8: return key

Fig. 3. Passive ASA against AEAD

symbol (neither 0 nor 1) to keep track of which key bits have been collected. In line 2 of the algorithm for $\tilde{\Pi}.\text{Dec}$, we write $B(\delta)$ to denote a Bernoulli trial which returns 1 with probability δ . Key extractor A.Ext takes as input the attacker key and the transcript, consisting of triples (c, d, n, v) where v is a bit representing whether or not the ciphertext decrypts to \perp .

Theorem 1. *Let Π be a perfectly-correct symmetric encryption scheme and let $\ell = \Pi.\text{kl}$. Let $\tilde{\Pi}.\text{Dec}$ and A.Ext be defined as in Fig. 3. Let \mathcal{M} be a message sampling algorithm, and $F: \{0, 1\}^{\text{A.kl}} \times \{0, 1\}^* \rightarrow [\ell] \times \{0, 1\}$ be a PRF with $\text{Adv}_F^{\text{prf}}(\mathcal{F}) < \epsilon$ for all efficient adversaries \mathcal{F} . Then*

- (1) $\text{Adv}_{\tilde{\Pi}, \mathcal{M}}^{\text{kr}}(\text{A}) \geq 1 - \ell e^{-\frac{q\delta}{2\ell}}$, where q is the number of queries that A.Ext makes to the transcript oracle.
- (2) For all distinguishers \mathcal{D} , $\text{Adv}_{\tilde{\Pi}, \tilde{\Pi}}^{\text{det}}(\mathcal{D}) \leq \frac{\delta q}{2}(1 + \epsilon)$ where \mathcal{D} makes q queries to its decryption oracle.

Proof of (1). We use a combinatorial argument. Notice that this is essentially a coupon collection problem. We are looking for the probability that every key bit has been exfiltrated. If we fix i key bits that are not exfiltrated, there are $\binom{\ell}{i}$ ways to choose those fixed key bits. The probability that (at least) i of the key bits have not been exfiltrated is given by $\binom{\ell}{i} \left(1 - \frac{i\delta}{2\ell}\right)^q$. Using the principle of inclusion exclusion, the probability that no key bit has not been exfiltrated is given by

$$\begin{aligned}
\text{Adv}_{\tilde{\Pi}, \mathcal{M}}^{\text{kr}}(\text{A}) &= \sum_{i=0}^{\ell} (-1)^i \binom{\ell}{i} \left(1 - \frac{i\delta}{2\ell}\right)^q \\
&\geq 1 - \ell \left(1 - \frac{\delta}{2\ell}\right)^q \\
&\geq 1 - \ell e^{-\frac{q\delta}{2\ell}}. \quad \square
\end{aligned}$$

Proof of (2). Clearly, the only way to distinguish between Π and $\tilde{\Pi}$ is to observe $\tilde{\Pi}.\text{Dec}$ output \perp . Thus in order to distinguish, \mathcal{D} must find $(m, d; n)$ such that

$\perp = \mathcal{O}_{\text{Dec}}(k, c, d; n)$ for $c \leftarrow \Pi.\text{Enc}(k, m, d; n)$. This reduces to \mathcal{D} finding some $c \parallel d$ such that $F(k_A, c \parallel d) = i \parallel k[i]$ for some index i . Call this event W . Notice that for any F it holds that for all k_A, c, d we have $F(k_A, c \parallel d) = i \parallel b$ for some index i and bit b .

We note that $\Pr[W] \leq \Pr[\text{PRF}_F(\mathcal{F})]$ for all PRF adversaries \mathcal{F} . If not, it would be possible for \mathcal{F} to act as a challenger to \mathcal{D} and win its prf game whenever W occurs. Thus,

$$\begin{aligned}
\text{Adv}_{\Pi, \tilde{\Pi}}^{\text{det}}(\mathcal{D}) &= \Pr[\text{Det}_{\Pi, \tilde{\Pi}}(\mathcal{D}) \mid b = 1] + \Pr[\text{Det}_{\Pi, \tilde{\Pi}}(\mathcal{D}) \mid b = 0] - 1 \\
&= 1 - (1 - \delta \cdot \Pr[W])^q \\
&\leq 1 - (1 - \delta \cdot \Pr[\text{PRF}_F(\mathcal{D})])^q \\
&\leq 1 - \left(1 - \frac{\delta}{2}(1 + \text{Adv}_F^{\text{prf}}(\mathcal{D}))\right)^q \\
&\leq 1 - \left(1 - \frac{\delta}{2}(1 + \epsilon)\right)^q \\
&\leq \frac{\delta q}{2}(1 + \epsilon). \quad \square
\end{aligned}$$

REMARK. Whereas (un)detectability does depend on the security of the PRF, the PRF can be quite weak without much impacting the adversary’s key recovery advantage. If the base scheme Π ’s ciphertexts are indistinguishable from random (IND\$), then the PRF could simply choose the first $\lceil \log(\ell) \rceil + 1$ many bits of the ciphertext. This seems paradoxical, as strong privacy security is usually a desirable property but here it allows a simpler ASA to be successful.

We note that in practice, the subverted decryption algorithm $\tilde{\Pi}.\text{Dec}$ can be made more effective in a number of ways. Indeed, the model is very conservative and in practice it may be possible for A.Ext to observe a number of distinguishable error messages following [10].

4.2 Attack 2: Active

Consider algorithms $\tilde{\Pi}.\text{Dec}$ and A.Ext as specified in Fig. 4. The adversary A.Ext crafts special messages using a length-preserving pseudo-random permutation E under the attacker key⁶. We let $E : \{0, 1\}^{\text{A.kl}} \times \{0, 1\}^* \rightarrow \{0, 1\}^*$. The security of E will determine how easily the distinguisher \mathcal{D} will be able to recreate a special message to trigger $\tilde{\Pi}$. Furthermore, as in the passive attack, $\tilde{\Pi}.\text{Dec}$ makes use of a PRF F to determine whether or not to reject submitted ciphertexts. We let $F : \{0, 1\}^{\text{A.kl}} \times \{0, 1\}^* \rightarrow [\Pi.\text{kl}] \times \{0, 1\}$. Although the notation implies keys are the same, we assume independent behaviour of F, E .⁷ We analyse this construction in the formal model defined by game $\text{ActiveKR}_{\tilde{\Pi}, \mathcal{M}}$ in Fig. 2.

⁶ See Appendix A for the definition of a length-preserving PRP.

⁷ Using only one key is just a trick to keep the notation compact.

Algorithm $\tilde{\Pi}.\text{Dec}(k_A, k, c, d; n)$	Algorithm $A.\text{Ext}^{\mathcal{O}_{\text{Enc}}, \mathcal{O}_{\text{Dec}}}(k_A)$
1: $m \leftarrow \tilde{\Pi}.\text{Dec}(k, c, d; n)$	1: $\forall i \in [\Pi.\text{kl}], \text{key}[i] \leftarrow \star$
2: if $(m \neq \perp)$ then	2: while $\exists j: \text{key}[j] = \star$ do
3: return m	3: $(c, d; n) \leftarrow \mathcal{O}_{\text{Enc}}()$
4: $\tilde{c} \leftarrow E^{-1}(k_A, c)$	4: parse $F(k_A, c \parallel d)$ as $i \parallel b$
5: $m \leftarrow \tilde{\Pi}.\text{Dec}(k, \tilde{c}, d; n)$	5: if $(\text{key}[i] = \star)$ then
6: if $(m = \perp)$ then	6: $\tilde{c} \leftarrow E(k_A, c)$
7: return \perp	7: $v \leftarrow \mathcal{O}_{\text{Dec}}(\tilde{c}, d; n)$
8: parse $F(k_A, \tilde{c} \parallel d)$ as $i \parallel b$	8: $\text{key}[i] \leftarrow b \oplus v$
9: if $(k[i] = b)$ then	9: return key
10: return m	
11: else	
12: return \perp	

Fig. 4. Active ASA against AEA

Theorem 2. Let Π be a perfectly-correct symmetric encryption scheme and let $\ell = \Pi.\text{kl}$. Let $\tilde{\Pi}.\text{Dec}$ and $A.\text{Ext}$ be defined as in Fig. 4. Let \mathcal{M} be a message sampling algorithm. Let $\ell = \Pi.\text{kl}$ and $\text{Adv}_{\Pi}^{\text{auth}} < \epsilon$. Let $F: \{0, 1\}^{\text{A.kl}} \times \{0, 1\}^* \rightarrow [\ell] \times \{0, 1\}$ be a PRF with $\text{Adv}_F^{\text{prf}}(\mathcal{F}) < 1$ for all efficient adversaries \mathcal{F} . Let E be a lp-PRP with $E: \{0, 1\}^{\text{A.kl}} \times \{0, 1\}^* \rightarrow \{0, 1\}^*$ and $\text{Adv}_E^{\text{prp}}(\mathcal{F}') < \epsilon'$ for all efficient PRP adversaries \mathcal{F}' . Then

- (1) $\text{Adv}_{\Pi, \mathcal{M}}^{\text{kr}}(A) \geq 1 - \ell e^{-\frac{q}{\ell}(1-\epsilon)}$, where $A.\text{Ext}$ makes exactly $\Pi.\text{kl}$ calls to the decryption oracle and q calls to the encryption oracle.
- (2) For every distinguisher \mathcal{D} , $\text{Adv}_{\Pi, \tilde{\Pi}}^{\text{det}}(\mathcal{D}) \leq \frac{q}{2^r} + \epsilon'$, where \mathcal{D} makes q queries to its decryption oracle.

Proof of (1). We use the same combinatorial argument as in Thm. 1. This time, the probability that (at least) i of the key bits have not been correctly exfiltrated is given by $\binom{\ell}{i} \left[\left(1 - \frac{i}{\ell}\right) + \frac{\alpha i}{2\ell} \right]^q$. Here α is the probability that $\tilde{\Pi}.\text{Dec}(k, \tilde{c}, d; n) \neq \perp$ given that $F^{-1}(k_A, \tilde{c}) = j \parallel k[j]$ for j in the set of indices being counted. We note that $\text{Adv}_{\Pi}^{\text{auth}} \geq \alpha$.

$$\begin{aligned}
\text{Adv}_{\Pi, \mathcal{M}}^{\text{kr}}(A) &= \sum_{i=0}^{\ell} (-1)^i \binom{\ell}{i} \left[\left(1 - \frac{i}{\ell}\right) + \frac{\alpha i}{2\ell} \right]^q \\
&\geq 1 - \ell \left(1 + \frac{1}{\ell} \left(\frac{\alpha}{2} - 1 \right) \right)^q \\
&\geq 1 - \ell e^{-\frac{q}{\ell}(1-\frac{\alpha}{2})} \\
&\geq 1 - \ell e^{-\frac{q}{\ell}(1-\epsilon)}. \quad \square
\end{aligned}$$

Proof of (2). As in Thm. 1, the only way to distinguish between Π and $\tilde{\Pi}$ is by observing $\tilde{\Pi}.\text{Dec}$ accepting a forged ciphertext. To do this, the distinguisher \mathcal{D} must find some ciphertext c with associated data d such that $F(k_A, \tilde{c} \parallel d) = i \parallel k[i]$

for some $i \in [\ell]$ and where $\tilde{c} = E^{-1}(k_A, c)$. Noting that $\text{Adv}_F^{\text{prf}}(\mathcal{F}) < 1$, we thus obtain

$$\Pr \left[\text{Det}_{\Pi, \tilde{\Pi}}(\mathcal{D}) \mid b = 0 \right] \leq \Pr \left[\begin{array}{l} \mathcal{D} \text{ finds } c \text{ with } E^{-1}(k_A, c) = \tilde{c} \text{ for some } \tilde{c} \\ \text{with } \Pi.\text{Dec}(k, \tilde{c}, d; n) \neq \perp, \text{ for some } d, n \end{array} \right]$$

Consider the following game, which we will refer to as the pre-image game. For $b \in \{0, 1\}$ we define experiment b as follows:

1. The challenger initially sets $C \leftarrow \emptyset$ and responds to query c_i in the following way:
 - if $(b = 0)$ then set $c'_i \leftarrow_{\$} \{0, 1\}^{|c_i|} \setminus C$, update $C \leftarrow^{\cup} c'_i$ and return c'_i
 - if $(b = 1)$ then return $c'_i \leftarrow E^{-1}(k_A, c_i)$.
2. The adversary \mathcal{D} submits a sequence of queries c_1, c_2, \dots, c_q to the challenger and receives c'_i for $i \in [q]$.

For $b \in \{0, 1\}$, let W_b be the event that \mathcal{D} outputs 1 in experiment b ; \mathcal{D} outputs 1 if for some d, n , $\Pi.\text{Dec}(k, c'_i, d; n) \neq \perp$. The advantage of \mathcal{D} in the pre-image game is clearly less than its advantage in distinguishing a lp-PRP from a random length preserving permutation. To see this, given \mathcal{D} with some advantage playing the pre-image game we can construct an adversary \mathcal{B} acting as a challenger to \mathcal{D} such that \mathcal{B} outputs 1 in the distinguishing game $\text{PRP}_E(\mathcal{B})$ whenever \mathcal{D} does in the pre-image game. Thus,

$$\Pr[W_0] - \Pr[W_1] \leq \text{Adv}_E^{\text{PRP}}(\mathcal{B}).$$

Noting that $\Pr[W_1] = \frac{q}{2^\tau}$, where τ is the stretch of the encryption scheme, we conclude that

$$\text{Adv}_{\Pi, \tilde{\Pi}}^{\text{det}}(\mathcal{D}) \leq \Pr[W_0] \leq \Pr[W_1] + \text{Adv}_E^{\text{PRP}}(\mathcal{B}) \leq \frac{q}{2^\tau} + \epsilon'. \quad \square$$

5 Implementation

We implemented our attacks in proof-of-concept Python code to verify their functionality and effectiveness.⁸ The particular AEAD scheme we attack is AES-GCM [15], using black-box access to the implementation provided by [32]. We simulated both active and passive attacks 10,000 times, and recorded the number of queries for successful extraction of a 128-bit key (thus, $\ell = 128$). Messages, nonces and associated data were generated using the `random.getrandbits` method from the `Crypto.Random` library. The plots below (Figs. 5 and 6) show the distribution (in blue) of the recorded number of queries q , and (in red) the cumulative success probability as a function of q . Our results confirm the theoretical estimates from Thm. 1 and 2; in particular, the exponential success rate. While the attacks have different application and success profiles, both reliably recover keys.

⁸ We are happy to share our source code. Please contact the authors.

PASSIVE. The expected number of calls to the transcript oracle for successful exfiltration is given by $\frac{2\ell}{\delta} \sum_{i=1}^{\ell} \frac{1}{i}$ (see proof of Thm. 1). We set $\delta = 0.1$ for illustration. This gives us an expected value of $q = 13910$ compared to the recorded mean of 13920.59. Alternatively, the result from Thm. 1 gives a key recovery advantage of $\approx 1/2$ with $q = 14000$, compared to the recorded median of 13380. The discrepancy is due to the exponential approximation in the proof.




Fig. 5. Results of running an implementation of the passive attack 10,000 times. Key length $\ell = 128$, and parameter $\delta = 0.1$. **Left axis:** The blue histogram shows the distribution of the number of queries required for successful key exfiltration. The data has been sorted into 50 bins. **Right axis:** The red curve shows the cumulative probability of successful key exfiltration against q .

ACTIVE. We assume that for AES-GCM, $\text{Adv}_{\Pi}^{\text{auth}} \approx 0$ and set $\epsilon = 0$. The expected number of encryption calls for successful exfiltration is then $\ell \sum_{i=1}^{\ell} \frac{1}{i}$ (see proof of Thm. 2). This gives an expected value of $q = 696$ compared to the recorded mean of 695.05. Alternatively, the result from Thm. 2 gives a key recovery advantage of $\approx 1/2$ with $q = 710$ compared to the recorded median of 670. Again, the difference is due to exponential approximation.




Fig. 6. Results of running an implementation of the active attack 10,000 times with key length $\ell = 128$. **Left axis:** The blue histogram shows the distribution of the number of queries required for successful key exfiltration. The data has been sorted into 50 bins. **Right axis:** The red curve shows the cumulative probability of successful key exfiltration against q .

6 Breaking Security without Extracting the Full Key

The attacks presented in Sect. 4 are generic in that they are independent of the targeted AEAD scheme. Our approach, in common with previous work, was to extract the full key with which the AEAD instance is operated. Message recovery follows immediately by the definition of correctness. From this it is tempting to conclude that choosing longer keys, e.g. 256 bits instead of 128 in the case of AES-based encryption, gives better security against ASAs. (This approach is generally explored in big key cryptography [7].) In this section we show that this intuition is not necessarily correct. As we detail, many current AEAD schemes have inner building blocks that maintain their own secret values, and scaling up key sizes does not automatically also increase the sizes of these internal values. Note that ASAs in the style proposed in the previous section could easily be adapted to leak this internal information instead of the key. As the recovery of such values might not always directly lead to full message recovery, the assessment of whether the resulting overall attack is more or less effective than our generic attacks has to be made on a per scheme basis. We

exemplify this on the basis of two of the currently best-performing AES-based AEAD schemes: GCM [15] and OCB3 [19]. In both cases, the size of the crucial internal value and the block size of the cipher have to coincide and the latter value is fixed to 128 bits for AES (independently of key size).

AES-GCM. We consider the following abstraction of GCM. The AEAD key k is used directly to create an instance E of the AES blockcipher. To encrypt a message m with respect to associated data d and nonce n , E is operated in counter mode, giving a pad $E(n+1) \parallel E(n+2) \parallel \dots$, where a specific nonce encoding ensures there are no collisions between counter values of different encryption operations. The first part c_1 of the ciphertext $c = c_1c_2$ is obtained by XORing the pad into the message, and finally the authentication tag c_2 is derived by computing $c_2 \leftarrow E(n) + H_h(d, c_1)$. Here H_h is an instance of a universal hash function H indexed (that is, keyed) with the 128-bit value $h = E(0^{128})$. Concretely, $H_h(d, c_1) = \sum_{i=1}^l v_i h^{l-i+1}$, where coefficients v_1, \dots, v_l are such that a prefix $v_1 \dots v_j$ is a length-padded copy of the associated data d , the middle part $v_{j+1} \dots v_{l-1}$ is a length-padded copy of ciphertext component c_1 , and the last item v_l is an encoding of the lengths of d and c_1 . The addition and multiplication operations deployed in this computation are those of a specific representation of the Galois field $\text{GF}(2^{128})$.

In executing a practical algorithm substitution attack against AES-GCM, it might suffice to leak the value h (which has length 128 independently of the AES key length, and furthermore stays invariant across encryption operations). The insight is that if the key of a universal hash function is known, then it becomes trivial to compute collisions. Concretely, assume the adversary is provided with the AES-GCM encryption $c = c_1c_2 = \text{Enc}(k, m, d; n)$ for unknown k, m but chosen d, n . Then by the above we have $c_2 = R + \sum_{i=1}^j v_i h^{l-i+1}$ where the coefficients $v_1 \dots v_j$ are an encoding of d and R is some residue. If, having been successfully leaked by the ASA, the internal value h is known, by solving a linear equation it is easy to find an associated data string $d' \neq d$, $|d'| = |d|$, such that for its encoding $v'_1 \dots v'_j$ we have $\sum_{i=1}^j v'_i h^{l-i+1} = \sum_{i=1}^j v_i h^{l-i+1}$. Overall this means that we have found $d' \neq d$ such that $\text{Enc}(k, m, d'; n) = c = \text{Enc}(k, m, d; n)$. In a CCA attack the adversary can thus query for the decryption of c with associated data d' and nonce n , and thus fully recover the target message m . We finally note that this attack can be directly generalized to one where also the c_1 and c_2 components are modified, resulting in the decryption of a message $m' \neq m$ for which the XOR difference between $m = m'$ is controlled by the adversary.

OCB3. Multiple quite different versions of the OCB encryption scheme exist, but a common property is that the associated data input is incorporated via ‘ciphertext translation’ [22]. To encrypt a message m under key k with associated data d and nonce n , in a first step the message m is encrypted with a pure AE scheme (no AD!) to an intermediate ciphertext $c^* \leftarrow \text{Enc}^*(k, m; n)$. Then to obtain the final ciphertext c , a pseudo-random function value $F_k(d)$ of the associated data string is XORed into the trailing bits of c^* . Concretely, in OCB3 we have $F_k(d) = \sum_{i=1}^l E(v_i + C_i)$ where all addition operations are XOR

combinations of 128 bit values, $E(\cdot)$ stands for AES enciphering with key k , values v_1, \dots, v_l represent a length-padded copy of associated data d , and coefficients C_1, \dots, C_l are (secret) constants deterministically derived from the value $L = E(0^{128})$.

In the context of an ASA we argue that it is sufficient to leak the 128 bit value L . The attack procedure is, roughly, as in the AES-GCM case. Assume the adversary is provided with the OCB3 encryption $c = \text{Enc}(k, m, d; n)$ for unknown k, m but chosen d, n , and assume the adversary knows L and thus C_1, \dots, C_l . Now let $1 \leq s < t \leq l$ be any two indices, let $\Delta = C_s + C_t$ and let $d' \neq d$, $|d'| = |d|$, be the associated data string with encoding v'_1, \dots, v'_l such that we have $v'_s = v_t + \Delta$ and $v'_t = v_s + \Delta$ and $v'_i = v_i$ for all $i \neq s, t$. Then we have $E(v'_s + C_s) = E(v_t + \Delta + C_s) = E(v_t + C_t)$ and $E(v'_t + C_t) = E(v_s + \Delta + C_t) = E(v_s + C_s)$, which leads to $F_k(d) = F_k(d')$ and ultimately $\text{Enc}(k, m, d'; n) = \text{Enc}(k, m, d; n)$. In a CCA attack environment, this can immediately be leveraged to the full recovery of m . As in the AES-GCM case, we note that many variants of our attack exist (against all versions of OCB), including some that manipulate message bits in a controlled way.

7 Conclusion

This work examines subversion attacks against decryption only, providing two examples of a new class of Algorithm Substitution Attack that provides a mass surveillance adversary with a powerful and attractive strategy to compromise the confidentiality of mass communication. Previous models of ASA against symmetric encryption only considered subverting the encryption algorithm, and seemed to suggest that decryption could only be subverted together with encryption (and that analysing such “total subversion” is uninteresting, as this gives an adversary too much power).

Acknowledgements

Thanks to Jeroen Pijnenburg and Fabrizio De Santis for their early comments on this paper. Thanks also to the anonymous reviewers.

References

1. Armour, M., Poettering, B.: Substitution attacks against message authentication. IACR Trans. Symmetric Cryptol. 2019(3), 152–168 (9 2019), <https://tosc.iacr.org/index.php/ToSC/article/view/8361>
2. Armour, M., Poettering, B.: Substitution attacks against message authentication. Cryptology ePrint Archive, Report 2019/989 (2019), <http://eprint.iacr.org/2019/989>
3. Armour, M., Poettering, B.: Subverting decryption in AEAD. Cryptology ePrint Archive, Report 2019/987 (2019), <http://eprint.iacr.org/2019/987>

4. Ateniese, G., Magri, B., Venturi, D.: Subversion-resilient signature schemes. In: Ray, I., Li, N., Kruegel, C. (eds.) *ACM CCS 2015: 22nd Conference on Computer and Communications Security*. pp. 364–375. ACM Press (Oct 2015)
5. Bellare, M., Hoang, V.T.: Resisting randomness subversion: Fast deterministic and hedged public-key encryption in the standard model. In: Oswald, E., Fischlin, M. (eds.) *Advances in Cryptology – EUROCRYPT 2015, Part II. Lecture Notes in Computer Science*, vol. 9057, pp. 627–656. Springer, Heidelberg (Apr 2015)
6. Bellare, M., Jaeger, J., Kane, D.: Mass-surveillance without the state: Strongly undetectable algorithm-substitution attacks. In: Ray, I., Li, N., Kruegel, C. (eds.) *ACM CCS 2015: 22nd Conference on Computer and Communications Security*. pp. 1431–1440. ACM Press (Oct 2015)
7. Bellare, M., Kane, D., Rogaway, P.: Big-key symmetric encryption: Resisting key exfiltration. In: Robshaw, M., Katz, J. (eds.) *Advances in Cryptology – CRYPTO 2016, Part I. Lecture Notes in Computer Science*, vol. 9814, pp. 373–402. Springer, Heidelberg (Aug 2016)
8. Bellare, M., Paterson, K.G., Rogaway, P.: Security of symmetric encryption against mass surveillance. In: Garay, J.A., Gennaro, R. (eds.) *Advances in Cryptology – CRYPTO 2014, Part I. Lecture Notes in Computer Science*, vol. 8616, pp. 1–19. Springer, Heidelberg (Aug 2014)
9. Berndt, S., Liskiewicz, M.: Algorithm substitution attacks from a steganographic perspective. In: Thuraisingham, B.M., Evans, D., Malkin, T., Xu, D. (eds.) *ACM CCS 2017: 24th Conference on Computer and Communications Security*. pp. 1649–1660. ACM Press (Oct / Nov 2017)
10. Boldyreva, A., Degabriele, J.P., Paterson, K.G., Stam, M.: On symmetric encryption with distinguishable decryption failures. In: Moriai, S. (ed.) *Fast Software Encryption – FSE 2013. Lecture Notes in Computer Science*, vol. 8424, pp. 367–390. Springer, Heidelberg (Mar 2014)
11. Camenisch, J., Drijvers, M., Lehmann, A.: Anonymous attestation with subverted TPMs. In: Katz, J., Shacham, H. (eds.) *Advances in Cryptology – CRYPTO 2017, Part III. Lecture Notes in Computer Science*, vol. 10403, pp. 427–461. Springer, Heidelberg (Aug 2017)
12. Degabriele, J.P., Farshim, P., Poettering, B.: A more cautious approach to security against mass surveillance. In: Leander, G. (ed.) *Fast Software Encryption – FSE 2015. Lecture Notes in Computer Science*, vol. 9054, pp. 579–598. Springer, Heidelberg (Mar 2015)
13. Dodis, Y., Ganesh, C., Golovnev, A., Juels, A., Ristenpart, T.: A formal treatment of backdoored pseudorandom generators. In: Oswald, E., Fischlin, M. (eds.) *Advances in Cryptology – EUROCRYPT 2015, Part I. Lecture Notes in Computer Science*, vol. 9056, pp. 101–126. Springer, Heidelberg (Apr 2015)
14. Dodis, Y., Mironov, I., Stephens-Davidowitz, N.: Message transmission with reverse firewalls—secure communication on corrupted machines. In: Robshaw, M., Katz, J. (eds.) *Advances in Cryptology – CRYPTO 2016, Part I. Lecture Notes in Computer Science*, vol. 9814, pp. 341–372. Springer, Heidelberg (Aug 2016)
15. Dworkin, M.J.: SP 800-38D: Recommendation for block cipher modes of operation: Galois/Counter Mode (GCM) and GMAC. US National Institute of Standards and Technology (2007)
16. Fischlin, M., Janson, C., Mazaheri, S.: Backdoored hash functions: immunizing HMAC and HKDF. In: *2018 IEEE 31st Computer Security Foundations Symposium (CSF)*. pp. 105–118. IEEE (2018)

17. Fischlin, M., Mazaheri, S.: Self-guarding cryptographic protocols against algorithm substitution attacks. In: 2018 IEEE 31st Computer Security Foundations Symposium (CSF). pp. 76–90. IEEE (2018)
18. Goh, E.J., Boneh, D., Pinkas, B., Golle, P.: The design and implementation of protocol-based hidden key recovery. In: Boyd, C., Mao, W. (eds.) ISC 2003: 6th International Conference on Information Security. Lecture Notes in Computer Science, vol. 2851, pp. 165–179. Springer, Heidelberg (Oct 2003)
19. Krovetz, T., Rogaway, P.: The OCB authenticated-encryption algorithm (2014), <https://tools.ietf.org/html/rfc7253>
20. Ma, H., Zhang, R., Yang, G., Song, Z., Sun, S., Xiao, Y.: Concessive online/offline attribute based encryption with cryptographic reverse firewalls - secure and efficient fine-grained access control on corrupted machines. In: López, J., Zhou, J., Soriano, M. (eds.) ESORICS 2018: 23rd European Symposium on Research in Computer Security, Part II. Lecture Notes in Computer Science, vol. 11099, pp. 507–526. Springer, Heidelberg (Sep 2018)
21. Mironov, I., Stephens-Davidowitz, N.: Cryptographic reverse firewalls. In: Oswald, E., Fischlin, M. (eds.) Advances in Cryptology – EUROCRYPT 2015, Part II. Lecture Notes in Computer Science, vol. 9057, pp. 657–686. Springer, Heidelberg (Apr 2015)
22. Rogaway, P.: Authenticated-encryption with associated-data. In: Atluri, V. (ed.) ACM CCS 2002: 9th Conference on Computer and Communications Security. pp. 98–107. ACM Press (Nov 2002)
23. Rogaway, P.: The moral character of cryptographic work. Cryptology ePrint Archive, Report 2015/1162 (2015), <http://eprint.iacr.org/2015/1162>
24. Russell, A., Tang, Q., Yung, M., Zhou, H.S.: Cliptography: Clipping the power of kleptographic attacks. In: Cheon, J.H., Takagi, T. (eds.) Advances in Cryptology – ASIACRYPT 2016, Part II. Lecture Notes in Computer Science, vol. 10032, pp. 34–64. Springer, Heidelberg (Dec 2016)
25. Russell, A., Tang, Q., Yung, M., Zhou, H.S.: Destroying steganography via amalgamation: Kleptographically CPA secure public key encryption. Cryptology ePrint Archive, Report 2016/530 (2016), <http://eprint.iacr.org/2016/530>
26. Russell, A., Tang, Q., Yung, M., Zhou, H.S.: Generic semantic security against a kleptographic adversary. In: Thuraisingham, B.M., Evans, D., Malkin, T., Xu, D. (eds.) ACM CCS 2017: 24th Conference on Computer and Communications Security. pp. 907–922. ACM Press (Oct / Nov 2017)
27. Russell, A., Tang, Q., Yung, M., Zhou, H.S.: Correcting subverted random oracles. In: Shacham, H., Boldyreva, A. (eds.) Advances in Cryptology – CRYPTO 2018, Part II. Lecture Notes in Computer Science, vol. 10992, pp. 241–271. Springer, Heidelberg (Aug 2018)
28. Schneier, B., Fredrikson, M., Kohno, T., Ristenpart, T.: Surreptitiously weakening cryptographic systems. Cryptology ePrint Archive, Report 2015/097 (2015), <http://eprint.iacr.org/2015/097>
29. Simmons, G.J.: The prisoners’ problem and the subliminal channel. In: Chaum, D. (ed.) Advances in Cryptology – CRYPTO’83. pp. 51–67. Plenum Press, New York, USA (1983)
30. Young, A., Yung, M.: The dark side of “black-box” cryptography, or: Should we trust capstone? In: Kobitz, N. (ed.) Advances in Cryptology – CRYPTO’96. Lecture Notes in Computer Science, vol. 1109, pp. 89–103. Springer, Heidelberg (Aug 1996)

31. Young, A., Yung, M.: Kleptography: Using cryptography against cryptography. In: Fumy, W. (ed.) Advances in Cryptology – EUROCRYPT’97. Lecture Notes in Computer Science, vol. 1233, pp. 62–74. Springer, Heidelberg (May 1997)
32. Zhu, B.: AES-GCM-Python (2013), https://github.com/bozhu/AES-GCM-Python/blob/master/aes_gcm.py

A Pseudo-random Functions and Permutations

We recall standard notions of pseudo-random functions and permutations.

Definition 4. A keyed pseudo-random function (PRF) for range R is an efficiently computable function $F: \{0, 1\}^\ell \times \{0, 1\}^* \rightarrow R$ taking a key $L \in \{0, 1\}^\ell$ and input $s \in \{0, 1\}^*$ to return an output $F(L, s) \in R$. Consider game $\text{PRF}_F(\mathcal{F})$ in Fig. 7 associated to F and adversary \mathcal{F} . Let

$$\text{Adv}_F^{\text{prf}}(\mathcal{F}) = 2 \cdot \Pr[\text{PRF}_F(\mathcal{F})] - 1$$

be the prf advantage of adversary \mathcal{F} against function F . Intuitively, the function is pseudo-random if the prf advantage of any realistic adversary is negligible.

Definition 5. A keyed length-preserving pseudo-random permutation (lp-PRP) is an efficiently computable function E where $E: \{0, 1\}^\ell \times \{0, 1\}^* \rightarrow \{0, 1\}^*$ takes a key $L \in \{0, 1\}^\ell$ and input $s \in \{0, 1\}^*$ to return an output $E(L, s) \in \{0, 1\}^{|s|}$. We require that any keyed instance of E is a permutation on $\{0, 1\}^n$ for all $n \in \mathbb{N}$ and also that its inverse E^{-1} is efficiently computable. Consider game $\text{PRP}_E(\mathcal{F})$ in Fig. 7 associated to E and adversary \mathcal{F} . Let

$$\text{Adv}_E^{\text{prp}}(\mathcal{F}) = 2 \cdot \Pr[\text{PRP}_E(\mathcal{F})] - 1$$

be the prp advantage of adversary \mathcal{F} against function E . Intuitively, the permutation is pseudo-random if the prp advantage of any realistic adversary is negligible.

Game $\text{PRF}_F(\mathcal{F})$	Game $\text{PRP}_E(\mathcal{F})$
$L \leftarrow_{\mathcal{S}} \{0, 1\}^\ell, S \leftarrow \emptyset$ $b \leftarrow_{\mathcal{S}} \{0, 1\}, b' \leftarrow \mathcal{F}^{\mathcal{O}_{\text{Fun}}}$ return $(b = b')$	$L \leftarrow_{\mathcal{S}} \{0, 1\}^\ell, S \leftarrow \emptyset$ $b \leftarrow_{\mathcal{S}} \{0, 1\}, b' \leftarrow \mathcal{F}^{\mathcal{O}_{\text{Perm}}}$ return $(b = b')$
$\mathcal{O}_{\text{Fun}}(s)$ if $(b = 1)$ then $y_s \leftarrow F(L, s)$ else if $s \notin S$ then $y_s \leftarrow_{\mathcal{S}} R$ $S \leftarrow S \cup \{s\}$ return y_s	$\mathcal{O}_{\text{Perm}}(s)$ if $(b = 1)$ then $y_s \leftarrow E(L, s)$ else if $s \notin S$ then $y_s \leftarrow_{\mathcal{S}} \{0, 1\}^{ s }$ $S \leftarrow S \cup \{s\}$ return y_s

Fig. 7. Game to define prf and prp advantage of \mathcal{F} with respect to F, E .