

Protecting encrypted cookies from compression side-channel attacks

Janaka Alawatugoda¹, Douglas Stebila^{1,2}, and Colin Boyd³

¹ *School of Electrical Engineering and Computer Science,*

² *School of Mathematical Sciences*

^{1,2} *Queensland University of Technology, Brisbane, Australia*

janaka.alawatugoda@qut.edu.au, stebila@qut.edu.au

³ *Department of Telematics,*

Norwegian University of Science and Technology, Trondheim, Norway

colin.boyd@item.ntnu.no

December 28, 2014

Abstract

Compression is desirable for network applications as it saves bandwidth; however, when data is compressed before being encrypted, the amount of compression leaks information about the amount of redundancy in the plaintext. This side channel has led to successful CRIME and BREACH attacks on web traffic protected by the Transport Layer Security (TLS) protocol. The general guidance in light of these attacks has been to disable compression, preserving confidentiality but sacrificing bandwidth. In this paper, we examine two techniques—heuristic separation of secrets and fixed-dictionary compression—for enabling compression while protecting high-value secrets, such as cookies, from attack. We model the security offered by these techniques and report on the amount of compressibility that they can achieve.

¹This is the full version of a paper published in the *Proceedings of the 19th International Conference on Financial Cryptography and Data Security (FC 2015)* in San Juan, Puerto Rico, USA, January 26–30, 2015, organized by the International Financial Cryptography Association in cooperation with IACR.

Contents

1	Introduction	3
2	Definitions	6
2.1	Encryption and compression schemes	6
2.2	Existing security notions	7
2.3	New security notions	7
2.4	Relations and separations between security notions	8
3	Technique 1: Separating secrets from user inputs	9
3.1	The scheme	9
3.2	CCI security of basic separating-secrets technique	10
3.3	Separating secrets in HTML	10
3.4	Experimental results on separating-secrets in HTML	11
3.5	Discussion	11
4	Technique 2: Fixed-dictionary compression	12
4.1	The scheme	12
4.2	CR security of basic fixed-dictionary technique	12
4.3	Experimental results on fixed-dictionary technique	14
4.4	Discussion	14
5	Conclusion	14
A	Relations and separations between security notions	16
A.1	$\text{IND-CPA} \implies \text{CCI}$	16
A.2	$\text{CCI} \not\Rightarrow \text{IND-CPA}$	18
A.3	$\text{CCI} \implies \text{RCI}$	18
A.4	$\text{RCI} \not\Rightarrow \text{CCI}$	19
A.5	$\text{RCI} \implies \text{CR}$	20
A.6	$\text{CR} \not\Rightarrow \text{RCI}$	21
B	Proof of CCI security of separating-secrets technique	22
C	Analysis of security of fixed-dictionary technique	23
C.1	Probability bounds, no prefix/suffix	23
C.2	Probability bounds, prefix/suffix	26

1 Introduction

To save communication costs, network applications often compress data before transmitting it; for example, the Hypertext Transport Protocol (HTTP) [FR14, §4.2] has an optional mechanism in which a server compresses the body of an HTTP response, most commonly using the gzip algorithm. When encryption is used to protect communication, compression must be applied before encryption (since ciphertexts should look random, they should have little apparent redundancy that can be compressed). In fact, to facilitate this, the Transport Layer Security (TLS) protocol [DR08, §6.2.2] has an optional compression mode that will compress all application data before encrypting it.

While compression is useful for reducing the size of transmitted data, it has had a negative impact when combined with encryption, because the amount of compression acts as a *side channel*. Most research considers side-channels such as timing [Koc96, KSWH98] or power consumption [HMF07], which can reveal information about cryptographic operations and secret parameters.

Compression-based leakage. In 2002, Kelsey [Kel02] showed how compression can act as a form of side-channel leakage. If plaintext data is compressed before being encrypted, the length of the ciphertext reveals information about the amount of compression, which in turn can reveal information about the plaintext. Kelsey notes that this side channel differs from other types of side channels in two key ways: “it reveals information about the plaintext, rather than key material”, and “it is a property of the algorithm, not the implementation”.

Kelsey’s most powerful attack is an *adaptive chosen input attack*: if an attacker is allowed to choose inputs x that are combined with a target secret s and the concatenation $x||s$ is compressed and encrypted, observing the length of the outputs can eventually allow the attacker to extract the secret s . For example, to determine the first character of s , the attacker could ask to have the string $x = \text{prefix*prefix}$ combined with s , then compressed and encrypted, for every possible character $*$; in one case, when $* = s_1$, the amount of redundancy is higher and the ciphertext should be shorter. Once each character of s is found, the attack can be carried out on the next character. The attack is somewhat noisy, but succeeds reasonably often.

Key to this attack is the fact that most compression algorithms (such as the DEFLATE algorithm underlying gzip) are *adaptive*: they adaptively build and maintain a *dictionary* of recently observed strings, and replace subsequent occurrences of that string with a code.

The CRIME and BREACH attacks. In 2012, Rizzo and Duong [RD12] showed how to apply Kelsey’s adaptive chosen input attack against gzip compression as used in TLS, in what they called the *Compression Ratio Info-leak Mass Exploitation (CRIME)* attack. The primary target of the CRIME attack was the user’s cookie in the HTTP header. If the victim visited an attacker-controlled web page, the attacker could use Javascript to cause the victim to send HTTP requests to URLs of the attacker’s choice on a specified server. The attacker could adaptively choose those URLs to include a prefix to carry out Kelsey’s adaptive chosen input attack. Some care is required to ensure the padding does not hide the length with block ciphers, but this can be dealt with. The CRIME attack also applies to compression as used in the SPDY protocol [The].

As a result of the CRIME attack, it was recommended that TLS compression be disabled, and the Trustworthy Internet Movement’s SSL Pulse report for December 2014 finds that just 7.2% of websites have TLS compression enabled [Tru14]; moreover, all major browsers have disabled it.

However, compression is also built into the HTTP protocol: servers can optionally compress the body of HTTP responses. While this excludes the cookie in the header, this attack can still succeed against secret values in the HTTP body, such as anti-cross-site request forgery (CSRF)

tokens. Suggested by Rizzo and Duong, this was demonstrated by Gluck et al. [GHP13] in the *Browser Reconnaissance and Exfiltration via Adaptive Compression of Hypertext (BREACH)* attack.

Mitigation techniques. Gluck et al. [GHP13] discussed several possible mitigation techniques against the BREACH attack, listed in decreasing order of effectiveness:

1. Disabling HTTP compression
2. Separating secrets from user input
3. Randomizing secrets per request
4. Masking secrets (effectively randomizing by XORing with a random nonce)
5. Length hiding (by adding a random number of bytes to the responses)
6. Rate-limiting the requests

Despite the demonstrated practicality of the BREACH attack, support for and use of HTTP compression remains widespread, due in large part to the value of decreasing communication costs and time. In fact, compression is even more tightly integrated into the proposed HTTP version 2 [BPT14] than previous versions. Techniques 2–4 generally require changes to both browsers and web servers. For example, masking secrets such as anti-CSRF tokens requires new mark-up for secrets, which browsers and servers can interpret to apply the randomized masking technique. Techniques 5–6 can be unilaterally applied by web servers, though length hiding can be defeated with statistical averaging, and rate-limiting must find a balance between legitimate requests and information leakage.

Related work. There has been little academic study of compression and encryption. Besides Kelsey’s adaptive chosen input attack and the related CRIME and BREACH attacks, the only relevant work we are aware of is that of Kelley and Tamassia [KT14]. They give a new security notion called *entropy-restricted semantic security* (ER-IND-CPA) for *keyed compression functions* which combine both encryption and compression: compared with the normal indistinguishability under chosen plaintext attack (IND-CPA) security notion, in ER-IND-CPA the adversary should not be able to distinguish between the encryption of two messages that *compress* to the same length. Kelley and Tamassia then show how to construct a cipher based on the LZW compression algorithm by rerandomizing the compression dictionary. Unfortunately, the ER-IND-CPA notion does not capture the CRIME and BREACH attacks, which depend on observing messages that compress to different lengths.

In leakage-resilient security definitions [AGV09, ADW09, DP08, NS09], leakage of the secret key is addressed. This differs from the setting in compression-based side-channel attacks, which addresses leakage of the plaintext. Thus, previous leakage-resilient approaches are not suitable to model compression-based side-channel attacks.

Our contributions. In this work, we study symmetric-key compression-encryption schemes, characterizing the security properties that can be achieved by various mitigation techniques in the face of CRIME- and BREACH-like attacks.

To some extent, the side channel exposed by compression is fundamentally unavoidable: if transmission of data is decreased, nothing can hide the fact that some redundancy existed in the plaintext. Hence, we focus our study on the ability of the attacker to learn specific “high value” secrets embedded in a plaintext, such as cookies or anti-CSRF tokens. In our models, we imagine there is a secret value ck , and the adversary can adaptively obtain encryptions

$$\text{Enc}_k(m' || ck || m'') \tag{1}$$

for prefix m' and suffix m'' of its choice; the attacker’s goal is to learn about ck .

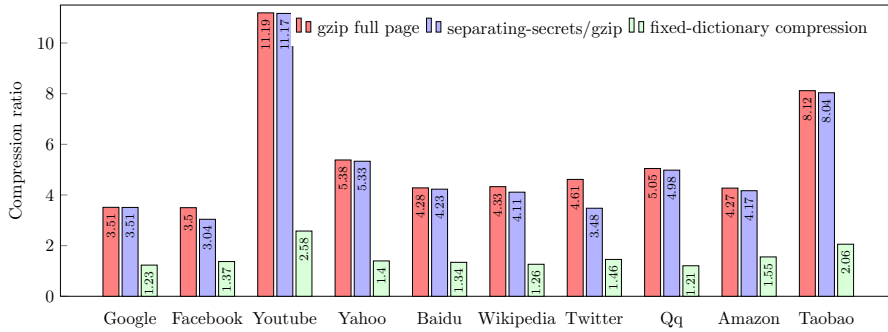


Figure 1: Compression ratios of full page compression versus mitigation techniques

The first mitigation technique we consider is that of *separating secrets*. During compression/encryption, an application-aware filter is applied to the plaintext to separate out any potential secret values from the data, the remaining plaintext is compressed, then the secrets and compressed plaintext are encrypted; after decryption, the inverse of the filter is used to reinsert the secret values in the decompressed plaintext. Assuming the filter fully separates out all secret values, we show that the separating secrets technique is able to achieve a strong notion of protection, which we call *chosen cookie indistinguishability* (CCI): the adversary cannot determine which of two cookies ck_0 and ck_1 of the adversary’s choice was encrypted with messages of the adversary’s choice given ciphertexts as in (1).

The second mitigation technique we consider is the use of a *fixed-dictionary compression scheme*, where the dictionary used for compression does not adapt to the plaintext being compressed, but instead is preselected in advance based on the expected distribution of plaintext messages, for example including common English words like “the” and “and”.¹ We show that, if the secret values are sufficiently high entropy, then fixed-dictionary compression is able to achieve *cookie recovery* (CR) security: if the secret cookie is chosen uniformly at random, the adversary cannot recover the entire secret cookie even given an adaptive message attack as in (1). While cookie recovery security does not meet the “gold standard” of indistinguishability notions for encryption, it may be sufficient for some settings, for example protecting compressed HTTP traffic from CRIME and BREACH attacks that try to recover cookies and anti-CSRF tokens.

We also characterize the relationship among the CCI and CR security notions, as well as an intermediate notion called *random cookie indistinguishability* (RCI) and the ER-IND-CPA notion of Kelley and Tamassia [KT14].

In the separating secrets technique, if the number of secrets extracted by the separating filter is relatively small, then the compressibility generally remains close to that of normal compression of the full plaintext. In the fixed-dictionary compression technique, compressibility suffers quite a bit compared to adaptive techniques on the full plaintext, although if the dictionary is constructed from a corpus of text similar to the plaintext, then some compression can be achieved.

Figure 1 summarizes experimental results comparing compression ratios for these two techniques on the HTML, CSS, and Javascript source code of the top 10 global websites as reported by Alexa Top Sites (<http://www.alexa.com/topsites>). On average, the compression ratio (uncompressed : compressed size) of gzip applied to the full source code was $5.42\times$; applying a separation filter that extracted all values following `value=` in the HTML source code yielded an average compression ratio of $5.20\times$; compression of each page using a fixed dictionary trained on all 10 pages yielded an average compression ratio of $1.55\times$.

¹Sadly “cryptography” is only the 29,697th most-frequently used English word. (http://en.wiktionary.org/wiki/Wiktionary:Frequency_lists/PG/2006/04/20001-30000)

2 Definitions

Notation. If x is a string, then x_i denotes the i th character of x ; $x_{i:\ell}$ denotes the length- ℓ substring of x starting at position i : $x_{i:\ell} = x_i || \dots || x_{i+\ell-1}$. If x and y are strings, then $x \preceq y$ denotes that x is a substring of y . The *index* of x in y is the smallest i such that $y_{i:|x|} = x$ and is denoted by $\text{ind}_y(x)$; if $x \not\preceq y$, we denote $\text{ind}_y(x) = \perp$. The empty string is denoted by ϵ .

2.1 Encryption and compression schemes

Recall the standard definition of an encryption scheme:

Definition 1 (Symmetric-key encryption). A *symmetric-key encryption scheme* Π for message space \mathcal{M} and ciphertext space \mathcal{C} is a tuple of algorithms:

- $\text{KeyGen}() \xrightarrow{\S} k$: A probabilistic *key generation algorithm* that generates a random key k in the keyspace \mathcal{K} .
- $\text{Enc}_k(m) \xrightarrow{\S} c$: A possibly probabilistic *encryption algorithm* that takes as input a key $k \in \mathcal{K}$ and a message $m \in \mathcal{M}$ and outputs a ciphertext $c \in \mathcal{C}$.
- $\text{Dec}_k(c) \rightarrow m'$ or \perp : A deterministic *decryption algorithm* that takes as input a key $k \in \mathcal{K}$ and a ciphertext $c \in \mathcal{C}$, and outputs either a message $m' \in \mathcal{M}$ or an error symbol \perp .

Correctness of symmetric-key encryption is defined in the obvious way: for all $k \xleftarrow{\S} \text{KeyGen}()$ and all $m \in \mathcal{M}$, we require that $\text{Dec}_k(\text{Enc}_k(m)) = m$.

Definition 2 (Compression scheme). A *compression scheme* Γ for message space \mathcal{M} with output space \mathcal{O} is a pair of algorithms:

- $\text{Comp}(m) \xrightarrow{\S} o$: A possibly probabilistic *compression algorithm* that takes as input a message $m \in \mathcal{M}$ and outputs an encoded value $o \in \mathcal{O}$.
- $\text{Decomp}(o) \rightarrow m'$ or \perp : A *decompression algorithm* that takes as input an encoded value $o \in \mathcal{O}$ and outputs a message $m' \in \mathcal{M}$ or an error symbol \perp .

Note that $|\text{Comp}(m)|$ may not necessarily be less than $|m|$; Shannon’s coding theorem implies that no algorithm can encode every message with shorter length, so not all messages may actually be “compressed”: some may increase in length.

Correctness of a compression scheme is again defined in the obvious way: for all $m \in \mathcal{M}$, we require that $\text{Decomp}(\text{Comp}(m)) = m$.

In this paper, we are interested in *symmetric-key compression-encryption schemes*, which formally are just symmetric-key encryption schemes as in Definition 1, but usually have the goal of outputting shorter ciphertexts via some form of compression. Of course, every symmetric-key encryption scheme is also a symmetric-key compression-encryption scheme, with “compression” being the identity function. We will often deal with the following specific, natural composition of compression and symmetric-key encryption:

Definition 3 (Composition of compression and encryption). Let $\Gamma = (\text{Comp}, \text{Decomp})$ be a compression scheme with message space \mathcal{M} and output space \mathcal{O} . Let $\Pi = (\text{KeyGen}, \text{Enc}, \text{Dec})$ be a symmetric-key encryption scheme with message space \mathcal{O} and ciphertext space \mathcal{C} . The symmetric-key compression-encryption scheme $\Pi \circ \Gamma$ constructed from Γ and Π is the following tuple:

$$\begin{aligned} (\Pi \circ \Gamma).\text{KeyGen}() &= \Pi.\text{KeyGen}() \\ (\Pi \circ \Gamma).\text{Enc}_k(m) &= \Pi.\text{Enc}_k(\Gamma.\text{Comp}(m)) \\ (\Pi \circ \Gamma).\text{Dec}_k(c) &= \Gamma.\text{Decomp}(\Pi.\text{Dec}_k(c)) \end{aligned}$$

Note that $\Pi \circ \Gamma$ is itself a symmetric-key encryption scheme with message space \mathcal{M} and ciphertext space \mathcal{C} . If Γ and Π are both correct, then so is $\Pi \circ \Gamma$.

$\text{Exp}_{\Pi}^{\text{IND-CPA}}(\mathcal{A})$ <ol style="list-style-type: none"> 1: $k \xleftarrow{\\$} \Pi.\text{KeyGen}()$ 2: $b \xleftarrow{\\$} \{0, 1\}$ 3: $(m_0, m_1, st) \xleftarrow{\\$} \mathcal{A}^E()$ 4: if $m_0 \neq m_1$, then return \perp 5: $c \leftarrow \Pi.\text{Enc}_k(m_b)$ 6: $b' \xleftarrow{\\$} \mathcal{A}^E(c, st)$ 7: return $(b' = b)$ $\underline{E(m)}$ <ol style="list-style-type: none"> 1: return $\Pi.\text{Enc}_k(m)$ 	$\text{Exp}_{\Pi, \mathcal{L}}^{\text{ER-IND-CPA}}(\mathcal{A})$ <ol style="list-style-type: none"> 1: $k \xleftarrow{\\$} \Pi.\text{KeyGen}()$ 2: $b \xleftarrow{\\$} \{0, 1\}$ 3: $(m_0, m_1, st) \xleftarrow{\\$} \mathcal{A}^E()$ 4: if $m_0 \notin \mathcal{L}$ or $m_1 \notin \mathcal{L}$, then return \perp 5: $c \leftarrow \Pi.\text{Enc}_k(m_b)$ 6: $b' \xleftarrow{\\$} \mathcal{A}^E(c, st)$ 7: return $(b' = b)$ $\underline{E(m)}$ <ol style="list-style-type: none"> 1: return $\Pi.\text{Enc}_k(m)$
---	--

Figure 2: Security experiments for indistinguishability under chosen plaintext attack (IND-CPA, left) and entropy-restricted IND-CPA (ER-IND-CPA, right)

2.2 Existing security notions

The standard security notion for symmetric-key encryption is indistinguishability of encrypted messages. In this paper, we focus on chosen plaintext attack. The security experiment $\text{Exp}_{\Pi}^{\text{IND-CPA}}(\mathcal{A})$ for indistinguishability under chosen plaintext attack (IND-CPA) of a symmetric-key encryption scheme Π against a stateful adversary \mathcal{A} is given in Figure 2. The *advantage* of \mathcal{A} in breaking the IND-CPA experiment for Π is $\text{Adv}_{\Pi}^{\text{IND-CPA}}(\mathcal{A}) = |2 \Pr(\text{Exp}_{\Pi}^{\text{IND-CPA}}(\mathcal{A}) = 1) - 1|$.

Kelley and Tamassia [KT14] give a definition of *entropy-restricted IND-CPA security* which applies to keyed compression schemes Π , and demands indistinguishability of encryptions of messages from the same class $\mathcal{L} \subseteq \mathcal{M}$; typically, \mathcal{L} is the class of messages that encrypt (compress) to the same length under $\Pi.\text{Enc}$, such as:

$$\mathcal{L}_{\ell} = \{m \in \mathcal{M} : |\Pi.\text{Enc}(m)| = \ell\} .$$

The ER-IND-CPA security experiment is given in Figure 2; the corresponding advantage is defined similarly. Kelley and Tamassia note that any IND-CPA-secure symmetric-key encryption scheme Π , combined with any compression scheme Γ , is immediately ER-IND-CPA-secure. As well, it is easily seen that if a symmetric-key encryption scheme is ER-IND-CPA-secure for the class $\mathcal{L}_{\ell} = \{m \in \mathcal{M} : |m| = \ell\}$, then that scheme is also an IND-CPA-secure symmetric-key encryption.

2.3 New security notions

In this paper, we focus on the ability of an attacker to learn about a secret piece of data inside a larger piece of data, where the attacker controls everything except the secret data. We use the term *cookie* to refer to the secret data; in practice, this could be an HTTP cookie in a header, an anti-CSRF token, or some piece of personal information. We will allow the attacker to adaptively obtain encryptions of compressions of data of the form $m' \| ck \| m''$ for a secret cookie ck and adversary-chosen message prefix m' and suffix m'' .

We now present three notions for the security of cookies in the context of compression-encryption schemes:

- *Cookie recovery (CR) security*: A simple, but relatively weak, security notion for symmetric-key compression-encryption schemes: it should be hard for the attacker to *fully recover* a secret value, even given adaptive access to an oracle that encrypts plaintexts of its choosing with the target cookie embedded.
- *Random cookie indistinguishability (RCI) security*: The adversary has to decide which of two randomly chosen cookies was embedded in the encrypted plaintext, given adaptive access to an oracle that encrypts plaintexts of its choosing with the target cookie embedded.

$\text{Exp}_{\Psi, \mathcal{CK}}^{\text{CR}}(\mathcal{A})$	$\text{Exp}_{\Psi, \mathcal{CK}}^{\text{RCI/CCI}}(\mathcal{A})$
1: $k \xleftarrow{\$} \Psi.\text{KeyGen}()$ 2: $ck \xleftarrow{\$} \mathcal{CK}$ 3: $ck' \xleftarrow{\$} \mathcal{A}^{E_1, E_2}()$ 4: return $(ck' = ck)$	1: $k \xleftarrow{\$} \Psi.\text{KeyGen}()$ 2: if RCI then 3: $(ck_0, ck_1) \xleftarrow{\$} \mathcal{CK}$ s.t. $ ck_0 = ck_1 $; $st \leftarrow \perp$ 4: else if CCI then 5: $(ck_0, ck_1, st) \xleftarrow{\$} \mathcal{A}^{E_2}()$ s.t. $ ck_0 = ck_1 $ 6: $b \xleftarrow{\$} \{0, 1\}$ 7: $b' \xleftarrow{\$} \mathcal{A}^{E_1, E_2}(ck_0, ck_1, st)$ 8: return $(b' = b)$
$E_1(m', m'')$ 1: return $\Psi.\text{Enc}_k(m' \ ck \ m'')$	$E_1(m', m'')$ 1: return $\Psi.\text{Enc}_k(m' \ ck_b \ m'')$
$E_2(m)$ 1: return $\Psi.\text{Enc}_k(m)$	$E_2(m)$ 1: return $\Psi.\text{Enc}_k(m)$

Figure 3: Security experiments for cookie recovery (left) and random cookie indistinguishability and chosen cookie indistinguishability (right) attacks

- *Chosen cookie indistinguishability (CCI) security:* Here, the adversary has to decide which of two cookies of the adversary’s choice was embedded in the encrypted plaintext, given adaptive access to an oracle that encrypts plaintexts of its choosing with the target cookie embedded.

These security notions are formalized in the following definition, which refers to the security experiments shown in Figure 3.

Definition 4 (CR, RCI, CCI security). Let Ψ be a symmetric-key compression-encryption scheme. Let \mathcal{A} denote an algorithm. Let \mathcal{CK} denote the cookie space. Let $\text{xxx} \in \{\text{CR}, \text{RCI}, \text{CCI}\}$ be a security notion. Consider the security experiment $\text{Exp}_{\Psi, \mathcal{CK}}^{\text{xxx}}(\mathcal{A})$ in Figure 3. Define $\text{Adv}_{\Psi, \mathcal{CK}}^{\text{CR}}(\mathcal{A}) = \Pr(\text{Exp}_{\Psi, \mathcal{CK}}^{\text{CR}}(\mathcal{A}) = 1)$ as the probability that \mathcal{A} wins the cookie recovery experiment for Ψ and \mathcal{CK} . Similarly, define $\text{Adv}_{\Psi, \mathcal{CK}}^{\text{xxx}}(\mathcal{A}) = |2 \Pr(\text{Exp}_{\Psi, \mathcal{CK}}^{\text{xxx}}(\mathcal{A}) = 1) - 1|$, $\text{xxx} \in \{\text{RCI}, \text{CCI}\}$, as the advantage that \mathcal{A} has in winning the random cookie and chosen cookie indistinguishability experiments.

Remark. The CR, RCI, and CCI security notions intentionally include only the confidentiality of the cookie as a security goal, and not the confidentiality of any non-cookie data in the rest of the message. In most applications it would be desirable to obtain confidentiality of non-cookie data as well, and in many real-world situations, the application layer’s cookie and non-cookie data are jointly sent to the security layer (such as SSL/TLS) for encryption. Our notions do not preclude the scheme from encrypting the non-cookie data as well (and in fact our constructions in Sections 3 and 4 do so). However, it is not possible in general to require confidentiality of the non-cookie data while still allowing it to be compressed, as that brings us back around to the original problem that motivated the work—compression of adversary-provided data can lead to ciphertexts of different lengths that break indistinguishability. This cycle can be broken by demanding some length restriction on the separated non-cookie data, such as in the ER-IND-CPA notion described in Section 2.2, but we omit that complication to focus solely on the security of the high-value secret cookies.

2.4 Relations and separations between security notions

Cookie recovery, being a computational problem rather than a decisional problem, is a weaker security notion. Keeping CR as an initial step, the RCI and CCI notions gradually increase the security afforded to the cookie.

The following relations exist between security notions for symmetric-key compression-encryption schemes:

$$\text{CCI} \implies \text{RCI} \implies \text{CR} .$$

In other words, every scheme that provides chosen cookie indistinguishability provides random cookie indistinguishability, and so on. Moreover, these notions are distinct, and we can show separations between them:

$$\text{CR} \not\Rightarrow \text{RCI} \not\Rightarrow \text{CCI} .$$

Additionally, we can connect our new notions with existing notions:

$$\text{ER-IND-CPA} \implies \text{IND-CPA} \implies \text{CCI} \quad \text{and} \quad \text{CCI} \not\Rightarrow \text{IND-CPA} .$$

(These last relations should be interpreted as follows. A standard (non-compressing) IND-CPA-secure symmetric-key encryption scheme is also CCI-secure. This is not to say, however, that an IND-CPA-secure symmetric-key encryption scheme combined with a compression scheme, such as $\Pi \circ \Gamma$ from Definition 3, is CCI-secure.)

The proofs of these relations and counterexamples for the separations appear in Appendix A.

3 Technique 1: Separating secrets from user inputs

In this section we analyze a mitigation technique against attacks that recover secrets from compressed data: separating secrets from user inputs. The basic idea of separating secrets from user inputs is: given an input, use a filter to separate all the secrets from the rest of the content, including user inputs. Then the rest of the content is compressed, while the secrets are kept uncompressed. This mitigation technique is a generic mitigation technique against a whole class of compression-based side-channel attacks.

3.1 The scheme

Definition 5 (Filter). A *filter* is an invertible (efficient) function $f : \{0, 1\}^* \rightarrow \{0, 1\}^* \times \{0, 1\}^*$.

Given a filter f and a compression scheme Γ , the separating-secrets scheme $\text{SS}_{f,\Gamma}$ is given in Figure 4.

$\text{SS}_{f,\Gamma}.\text{Comp}(m)$	$\text{SS}_{f,\Gamma}.\text{Decomp}(pt)$
1: $(pt_s, pt_{ns}) \leftarrow f(m)$ 2: $pt_{ns} \leftarrow \Gamma.\text{Comp}(pt_{ns})$ 3: return $pt_s pt_{ns}$	1: Parse $pt_s \widetilde{pt_{ns}} \leftarrow pt$ 2: $pt_{ns} \leftarrow \Gamma.\text{Decomp}(\widetilde{pt_{ns}})$ 3: $m \leftarrow f^{-1}(pt_s, pt_{ns})$ 4: return m

Figure 4: Abstract separating-secrets compression scheme SS

Our results will make use of the following two conditions on filters. Intuitively, a filter is *effective* if it removes cookies from an input string, and is *safe* if no prefix/suffix can fool the filter into separating out one cookie but not another.

Definition 6 (Effective filter). Let \mathcal{CK} be a cookie space, and let f be a filter. We say that f is *effective at separating out* \mathcal{CK} if, for all $ck \in \mathcal{CK}$ and all m', m'' , we have that $ck \not\leq y$, where $(x, y) = f(m' || ck || m'')$.

Definition 7 (Safe filter). Let \mathcal{CK} be a cookie space, and let f be a filter. We say that f is *safe for* \mathcal{CK} if, for all $ck_0, ck_1 \in \mathcal{CK}$ such that $|ck_0| = |ck_1|$ and all m', m'' , we have that $|x_0| = |x_1|$ and $y_0 = y_1$, where $(x_0, y_0) = f(m' || ck_0 || m'')$ and $(x_1, y_1) = f(m' || ck_1 || m'')$.

Example cookie space and filter. Let $\lambda \in \mathbb{N}$ and let \mathcal{CK} be the set of alphanumeric strings starting with the literal “secret” and starting and ending with a space (denoted by $_$), i.e., strings matched by the regular expression

$$_ \text{secret} [\text{A-Za-z0-9}]^\lambda _$$

Let f be a filter that uses the above regular expression to separate out secrets. Consider a string of the form $m = m_ck_m_ck_m_ck_ \dots ck_m$, where m_i contains no substring matching the above regular expression and ck_i is a string completely matching the above regular expression (excluding the initial and terminal space $_$). Then $f(m) = (pt_s, pt_{ns})$, where $pt_s = ck_1 \| \dots \| ck_n$ and $pt_{ns} = m_0 \| \tau \| m_1 \| \tau \| \dots \| m_n$, and τ represents a fixed replacement token that can not appear as a substring of any $m \in \mathcal{M}$.

Claim 1. *The above filter f is effective at separating out and safe for the above \mathcal{CK} .*

Proof sketch. Since each cookie begins and ends with a character $_$ which does not appear within the cookie, no prefix or suffix can cause the filter to not separate a cookie.

More precisely, for any $_ck_ \in \mathcal{CK}$ and any m' that contains no substring matching the above regular expression and any $m'' \neq \epsilon$, we have that $f(m' \| _ck_ \| m'') = (ck \| x, m' \| \tau \| y)$, where $(x, y) = f(m'')$. Such an f is effective at separating out \mathcal{CK} since it separates every substring of m this is a cookie into the first component of the output. Moreover, f is safe for \mathcal{CK} by recursively applying the above identity. \square

3.2 CCI security of basic separating-secrets technique

In this section we analyze the security of separating-secrets mitigation technique according to CCI notion. Let $\Pi = (\text{KeyGen}, \text{Enc}, \text{Dec})$ be an IND-CPA-secure symmetric-key encryption scheme and $\text{SS}_{f,\Gamma}$ be the separating-secrets compression scheme given in Figure 4. We consider the security of the resulting symmetric-key compression-encryption scheme $\Pi \circ \text{SS}_{f,\Gamma}$, showing that, if the filter f safely separates out cookies, then breaking chosen cookie indistinguishability of $\Pi \circ \text{SS}_{f,\Gamma}$ is as hard as breaking indistinguishability (IND-CPA) of encryption scheme Π . The proof of Theorem 1 appears in Appendix B.

Theorem 1. *Let Π be a symmetric-key encryption scheme and let Γ be a compression scheme. Let \mathcal{CK} be a cookie space, and let f be a filter that is safe for \mathcal{CK} . Let \mathcal{A} be any adversary against the CCI security of the separating-secrets symmetric-key compression-encryption scheme $\Pi \circ \text{SS}_{f,\Gamma}$, and let q denote the number of queries that \mathcal{A} makes to its E_1 oracle. Then*

$$\text{Adv}_{\Pi \circ \text{SS}_{f,\Gamma}, \mathcal{CK}}^{\text{CCI}}(\mathcal{A}) \leq q \cdot \text{Adv}_{\Pi}^{\text{IND-CPA}}(\mathcal{B}^{\mathcal{A}}) ,$$

where \mathcal{B} is an algorithm, constructed using the adversary \mathcal{A} as described in the proof, against the IND-CPA security of the symmetric-key encryption scheme Π .

3.3 Separating secrets in HTML

Separating secrets from user inputs is a realistic mitigation technique against the BREACH attack: in the application layer, some fields which contain secrets (such as anti-CSRF tokens) can be identified and separated from the HTTP response body. In order to implement separating secrets from user inputs in HTML we need to describe a filter f_{HTML} .

One possible method to separate secrets in HTML is to separate the content assigned to the `value` attribute of HTML elements. Among other uses, the `value` attribute defines the value of a specific field in a form. The HTML code segment of Figure 5 shows inclusion of a secret anti-CSRF token as a hidden `input` field in a web form, which will appear in a HTML response body. By separating the content in the `value` attribute, we separate the anti-CSRF token.

```

<form action="/money_transfer" method="post">
<input type="hidden" name="csrftoken"
      value="OWT4NmQlODE4ODRjN2Q1NTlhMmZlYWE...">
...
</form>

```

Figure 5: HTML code segment showing inclusion of anti-CSRF token in a web form

Table 1: Compression performance (file size in bytes and compression ratio) for separating secrets (Section 3) and fixed dictionary (Section 4) techniques

Website	Uncompressed	gzip full page	Separating secrets	Fixed dictionary
Google.com	145 599	41 455 (3.51×)	41 502 (3.51×)	117 794 (1.23×)
Facebook.com	48 226	13 785 (3.50×)	15 863 (3.04×)	35 036 (1.37×)
Youtube.com	467 928	41 813 (11.19×)	41 893 (11.17×)	181 676 (2.58×)
Yahoo.com	444 408	82 572 (5.38×)	83 342 (5.33×)	318 386 (1.40×)
Baidu.com	74 979	17 519 (4.28×)	17 727 (4.23×)	55 950 (1.34×)
Wikipedia.org	48 548	11 217 (4.33×)	11 809 (4.11×)	38 406 (1.26×)
Twitter.com	57 777	12 520 (4.61×)	16 618 (3.48×)	39 712 (1.46×)
Qq.com	626 297	124 108 (5.05×)	125 747 (4.98×)	519 830 (1.21×)
Amazon.com	234 609	54 922 (4.27×)	56 278 (4.17×)	150 924 (1.55×)
Taobao.com	192 068	23 658 (8.12×)	23 898 (8.04×)	93 410 (2.06×)

The following (case-insensitive) regular expression can be used to separate out quoted anti-CSRF tokens in the `value` attribute of HTML elements:

$$\text{value}\backslash\text{s}*\text{=}\backslash\text{s}*"[A-Za-z0-9]+|\text{value}\backslash\text{s}*\text{=}\backslash\text{s}'[A-Za-z0-9]+'$$

This filter is effective at separating out and safe for the implied set of cookies, in the sense of Definitions 6 and 7.

However, the above regular expression is not perfect, highlighting the challenges of using heuristic techniques to separate out secrets.

First, the above regular expression will also capture the `value` attribute of HTML elements other than hidden `input` elements, such as `option`, which may not need to be treated as secret, so it is not as efficient as it could be.

Second, the above regular expression does not capture anti-CSRF tokens in unquoted `value` attributes, such as `value=OWT4NmQl`, which are allowed by the HTML specification. While it is easy to add an additional term such as `|value\s*=\s*[A-Za-z0-9]+` to the regular expression to capture unquoted attributes, this filter would no longer be effective in the sense of Definition 6: if a cookie is `value=OWT4NmQl`, and the adversary constructs $m' = \text{value=}$, then $m' || ck = \text{value=value=OWT4NmQl}$, and the filter applied to $m' || ck$ would separate out `value=value` as the cookie and leave `=OWT4NmQl` unprotected.

3.4 Experimental results on separating-secrets in HTML

Table 1 shows the result of applying the above regular expression to separate secrets on the top 10 global websites of Alexa Top Sites. As most pages contain little data in `value` attributes, the total amount of space required to transmit the separated secrets plus the remaining data is not much more than when the full page is compressed. (Table 1 also contains performance results of the fixed dictionary technique, to be discussed in Section 4.)

3.5 Discussion

The main drawback of the separating secrets technique is that the separation filter must be application-dependent. We noted already the challenges in using the heuristic regular expression

above to capture anti-CSRF tokens: it may separate out non-secrets as well as secrets (which yields suboptimal compression) and it does not capture unquoted tokens (which is a problem for security).

Moreover, this HTML filter also only captures secrets in a `value` attribute, which does not necessarily capture all values that might be considered sensitive. For example, should the titles of books in a search results page on an shopping site be considered secret? If so, an alternative separation filter would have to be developed. To provide complete certainty, secret separation would require additional markup with which the developer clearly identifies which data should be treated as secret. Otherwise, any sensitive values which are not separated may be compressed together with user inputs and other application data, and hence remain open to the compression-based side-channel.

4 Technique 2: Fixed-dictionary compression

The CRIME and BREACH attacks work because the dictionary constructed by the DEFLATE compression algorithm is adaptive: if the attacker injects a substring of the target secret into the plaintext nearby the secret itself, then the plaintext will compress more because of the repeated substring. Some early compression algorithms were non-adaptive, using a fixed dictionary mechanism. For example, Pike [Pik80] used a fixed dictionary of 205 popular English words and a variable length coding mechanism to compress typical English text at a rate of less than 4 bits per character. Another recent algorithm, Smaz [San09], similarly uses a fixed dictionary consisting of common digrams and trigrams from English and HTML source code, allowing it to compress even very short strings. Because the CRIME and BREACH attacks rely on the adaptivity of the compression dictionary, fixed-dictionary algorithms can offer resistance to such attacks while still providing some compression, albeit not as good as adaptive compression.

In this section, we investigate the use of fixed-dictionary compression in the context of encryption. We describe the basic idea of fixed-dictionary compression. We show that fixed-dictionary compression-encryption schemes can satisfy cookie recovery security for sufficiently large cookies. We then present an example of a modern fixed-dictionary compression algorithm and report on the compression ratios achieved by our algorithm.

4.1 The scheme

In general, fixed-dictionary compression schemes work by advancing through the string x and looking to see if the current substring appears in the dictionary \mathcal{D} : if it does, then an encoding of the index of the substring is recorded, otherwise an encoding of the current substring is recorded. The compression scheme must specify the encoding rules in a way that unambiguously discriminates between the two cases to allow for correct decompression.

An abstract version of a fixed-dictionary fixed-width compression algorithm FD is given in Figure 6. FD checks if the current substring of length w appears in the dictionary \mathcal{D} . If it does, it records the index of the substring in \mathcal{D} and advances w characters. If it does not, it records the next ℓ characters directly, then advances. (Using $\ell > 1$ but $\ell < w$ may be more efficient when it comes to encodings.) One could treat \mathcal{D} either as a set of strings (recording which element is matched) or a long string (recording the starting and ending position of the matching substring); we will use the latter in the rest of this section.

For example, if $\mathcal{D} = \text{“cookierecoveryattack”}$, then $\text{FD}_{\mathcal{D},4,2}.\text{Comp}(\text{“recover the cookie”})$ yields `7ver_the_1ie`.

4.2 CR security of basic fixed-dictionary technique

Let Π be a symmetric-key encryption scheme. Let \mathcal{D} be a dictionary of length d and $\text{FD}_{\mathcal{D},w,\ell}$ be the abstract fixed-dictionary compression scheme in Figure 6.

<u>FD_{\mathcal{D},w,ℓ}.Comp(x)</u>	<u>FD_{\mathcal{D},w,ℓ}.Decomp(y)</u>
1: $y \leftarrow \epsilon$	1: $x \leftarrow \epsilon$
2: $i \leftarrow 1$	2: $i \leftarrow 1$
3: while $i \leq x - w + 1$ do	3: while $i \leq y $ do
4: if $x_{i:w} \preceq \mathcal{D}$ then	4: if y_i encodes an index then
5: $y \leftarrow y \parallel \text{encoding of } \text{ind}_{\mathcal{D}}(x_{i:w})$	5: $x \leftarrow x \parallel \mathcal{D}_{y_{i:w}}$
6: $i \leftarrow i + w$	6: $i \leftarrow i + 1$
7: else	7: else
8: $y \leftarrow y \parallel \text{encoding of } x_{i:\ell}$	8: $x \leftarrow x \parallel \text{decoding of } y_{i:\ell'}$
9: $i \leftarrow i + \ell$	9: $i \leftarrow i + \ell'$
10: return y	10: return x

Figure 6: Abstract fixed-dictionary fixed-width compression scheme FD
Note the simplification that ℓ characters of x are encoded as ℓ' characters of y .

Suppose the cookie space is binary strings of length 8λ , or equivalently byte strings of length λ : $\mathcal{CK} = \{0x00, \dots, 0xFF\}^\lambda$.

If Π is a secure encryption scheme, then, intuitively, the only way the adversary can learn information about the cookie from seeing ciphertexts $\text{Enc}_k(\cdot \parallel ck \parallel \cdot)$ and $\text{Enc}_k(\cdot)$ is from the length of the ciphertext: if some substring of ck appears in the dictionary \mathcal{D} , then ck will compress, and that length difference tells the adversary that the secret cookie is restricted to some subset of \mathcal{CK} matching \mathcal{D} .

The situation is subtler in the full CR experiment: the attacker can provide m' and m'' and get $\text{Enc}_k(\text{Comp}(m' \parallel ck \parallel m''))$. If the last few bytes of m' followed by the first few bytes of ck appear in \mathcal{D} , then the string will compress more. This allows the attacker to carry out a CRIME-like attack on the first few bytes of ck .

For example, let $w = 4$ and suppose $\mathcal{D} = 1234567890\text{ABCDEFGHIJKLMN}\text{OPQRSTUVWXYZ}$ and $\mathcal{CK} = [0\text{-9A-F}]^\lambda$. The attacker can query $m' = 890$, $m' = 90\text{A}$, $m' = 0\text{AB}$, \dots . In exactly one case, the adversary's m' combined with the cookie's first byte will be in the dictionary, telling the adversary ck_1 . For example, if $ck_1 = \text{B}$, then when the adversary queries $m' = 90\text{A}$, the value that is compressed and then encrypted is $m' \parallel ck \parallel m'' = 90\text{AB}\dots$, which is a substring of \mathcal{D} .

While this allows the attacker to recover the first byte or two of the secret cookie with decent probability, it drops off exponentially; a similar argument applies to the last few bytes of the secret cookie. Theorem 2 captures this issue. Theorem 2 only provides quantifiable security when the cookie length n is significantly bigger than the compression window w . Additionally, this type of attack on the first/last few bytes of the cookie precludes *indistinguishable* security, which is why we focus on cookie *recovery* here. (Admittedly, in some settings recovering the first/last few cookie bytes may still be quite damaging.)

Theorem 2. *Let Π be a symmetric-key encryption scheme. Let \mathcal{D} be a dictionary of d words, each of length ℓ . Let w be positive integer. Let $\mathcal{CK} = \Omega^n$. Let \mathcal{A} be any adversary against the cookie recovery security of the fixed-dictionary symmetric-key compression-encryption scheme $\Pi \circ \text{FD}_{\mathcal{D},w,\ell}$. Then*

$$\text{Adv}_{\Pi \circ \text{FD}_{\mathcal{D},w,\ell}}^{\text{CR}}(\mathcal{A}) \leq \text{Adv}_{\Pi}^{\text{IND-CPA}}(\mathcal{B}) + 2^{-\Delta} \ ,$$

where \mathcal{B} is an algorithm, constructed using adversary \mathcal{A} , against the IND-CPA security of the symmetric-key encryption scheme Π , and

$$\Delta \geq \left(1 - d \left(1 - \left(1 - \frac{1}{|\Omega|^w} \right)^{n-3w+1} \right) \right) \cdot \log_2 \left(|\Omega|^{n-2w} - |\Omega|^{n-2w} \cdot d \left(1 - \left(1 - \frac{1}{|\Omega|^w} \right)^{n-3w+1} \right) \right) \ .$$

For example, for cookies of $n = 16$ bytes, with a dictionary of $d = 4000$ words each of length $w = 4$, we have $\Delta \geq 63.999695$. Doubling d gives $\Delta \geq 63.999391$.

The derivation and proof of the formula in Theorem 2 appear in Appendix C.

4.3 Experimental results on fixed-dictionary technique

Table 1 shows the result of applying a fixed-dictionary based compression algorithm to the top 10 global websites of Alexa Top Sites. The 4000-byte dictionary was built from the most common 8-, 16-, and 32-character substrings of the pages. The compression algorithm was based in part on the Smaz [San09] algorithm and was adapted slightly from Figure 6, to allow for variable-length words to be matched. Specifically, when attempting to encode the substring at the current position at line 4 in Figure 6, we first try variable length words in order of decreasing length, checking to see if $w = 18$, then $w = 16$, then \dots , then $w = 4$ characters can be found in the dictionary. This requires the encoding to include both index and length of the dictionary substring.

- To encode a dictionary word at index $0 \leq j < 4096$ of length $w = 2w' + 4, 0 \leq w' \leq 7$, store 16 bits: $1 \parallel [12\text{-bit encoding of } j] \parallel [3\text{-bit encoding of } w']$
- To encode 2 lower-ASCII characters $z_1 z_2$, store 16 bits: $00 \parallel [7\text{-bit encoding of } z_1] \parallel [7\text{-bit encoding of } z_2]$
- To encode 1 byte z , store 16 bits: $01000000 \parallel [8\text{-bit encoding of } z]$

4.4 Discussion

The main drawback of the fixed dictionary mitigation technique is that in practice it achieves relatively poor—albeit non-zero—compression compared with adaptive compression techniques. However, it does not rely on application-dependent or heuristic techniques for separating secrets.

5 Conclusion

In this paper we introduced theoretical models to analyze compression-based side-channel attacks on high-value secrets embedded inside messages: the notions of cookie recovery (CR) security, random cookie indistinguishability (RCI), and chosen cookie indistinguishability (CCI). Each notion allows an attacker adaptive access to an oracle which encrypts chosen plaintexts alongside a target secret.

The simple, but relatively weak, CR security notion is sufficient to model real-world compression-based side-channel attacks such as CRIME and BREACH that aim to recover the target secret. The CCI security notion addresses stronger situations where the adversary has to decide which of two secrets of the adversary’s choice was embedded in the encrypted plaintext, even given adaptive access to an oracle that encrypts plaintexts of its choosing with the target secret embedded.

The most secure countermeasure to compression-based side-channel attacks remains to disable compression. As implementers seem loathe to do so—indeed, compression is even more heavily embedded in current drafts of HTTP version 2 [BPT14, §10.6] than it was in previous versions—techniques for safely compressing data that may be partially adversarially controlled are of significant importance. While compression inherently leaks information about redundancy in plaintext, some compression techniques, such as the separating secrets and fixed dictionary approaches treated in this paper, provide some resistance to previous compression-based attacks like CRIME and BREACH. Further cryptographic study of compression seems like a worthwhile research direction, including the investigation of definitions that provide both cookie indistinguishability and some measure of message indistinguishability.

Acknowledgements

The authors acknowledge support by Australian Research Council (ARC) Discovery Project DP130104304.

References

- [ADW09] Joël Alwen, Yevgeniy Dodis, and Daniel Wichs. Leakage-resilient public-key cryptography in the bounded-retrieval model. In Shai Halevi, editor, *CRYPTO 2009, LNCS*, volume 5677, pp. 36–54. Springer, August 2009.
- [AGV09] Adi Akavia, Shafi Goldwasser, and Vinod Vaikuntanathan. Simultaneous hardcore bits and cryptography against memory attacks. In Omer Reingold, editor, *TCC 2009, LNCS*, volume 5444, pp. 474–495. Springer, March 2009.
- [BPT14] Mike Belshe, Roberto Peon, and Martin Thomson. Hypertext Transfer Protocol version 2, November 2014. Internet-Draft. <http://tools.ietf.org/html/draft-ietf-httpbis-http2-16>.
- [DP08] Stefan Dziembowski and Krzysztof Pietrzak. Leakage-resilient cryptography. In *49th FOCS*, pp. 293–302. IEEE Computer Society Press, October 2008.
- [DR08] T. Dierks and E. Rescorla. The Transport Layer Security (TLS) Protocol Version 1.2. RFC 5246 (Proposed Standard), August 2008. URL <http://www.ietf.org/rfc/rfc5246.txt>. Updated by RFCs 5746, 5878, 6176.
- [FR14] R. Fielding and J. Reschke. Hypertext Transfer Protocol (HTTP/1.1): Message Syntax and Routing. RFC 7230 (Proposed Standard), June 2014. URL <http://www.ietf.org/rfc/rfc7230.txt>.
- [GHP13] Yoel Gluck, Neal Harris, and Angelo Prado. SSL, gone in 30 seconds: A BREACH beyond CRIME. In *Black Hat USA 2013*, August 2013. URL <https://www.blackhat.com/us-13/archives.html#Prado>.
- [HMF07] Michael Hutter, Stefan Mangard, and Martin Feldhofer. Power and EM attacks on passive 13.56 mhz RFID devices. In Pascal Paillier and Ingrid Verbauwhede, editors, *CHES 2007, LNCS*, volume 4727, pp. 320–333. Springer, September 2007.
- [Kel02] John Kelsey. Compression and information leakage of plaintext. In Joan Daemen and Vincent Rijmen, editors, *FSE 2002, LNCS*, volume 2365, pp. 263–276. Springer, February 2002.
- [Koc96] Paul C. Kocher. Timing attacks on implementations of Diffie-Hellman, RSA, DSS, and other systems. In Neal Koblitz, editor, *CRYPTO'96, LNCS*, volume 1109, pp. 104–113. Springer, August 1996.
- [KSWH98] John Kelsey, Bruce Schneier, David Wagner, and Chris Hall. Side channel cryptanalysis of product ciphers. In Jean-Jacques Quisquater, Yves Deswarte, Catherine Meadows, and Dieter Gollmann, editors, *ESORICS'98, LNCS*, volume 1485, pp. 97–110. Springer, September 1998.
- [KT14] James Kelley and Roberto Tamassia. Secure compression: Theory & practice. Cryptology ePrint Archive, Report 2014/113, 2014. <http://eprint.iacr.org/2014/113>.

- [NS09] Moni Naor and Gil Segev. Public-key cryptosystems resilient to key leakage. In Shai Halevi, editor, *CRYPTO 2009, LNCS*, volume 5677, pp. 18–35. Springer, August 2009.
- [Pik80] J. Pike. Text compression using a 4 bit coding scheme. *The Computer Journal*, **24**(4):324–330, September 1980. DOI:10.1093/comjnl/24.4.324.
- [RD12] Juliano Rizzo and Thai Duong. The CRIME attack, 2012. Presented at ekoparty '12. <http://goo.gl/mlw1X1>.
- [San09] Salvatore Sanfilippo. Smaz: Small strings compression library, April 2009. <https://github.com/antirez/smaz>.
- [The] The Chromium Projects. SPDY. URL <http://dev.chromium.org/spdy>. <http://dev.chromium.org/spdy>.
- [Tru14] Trustworthy Internet Movement. SSL Pulse, December 2014. <https://www.trustworthyinternet.org/ssl-pulse/>.

A Relations and separations between security notions

This subsection briefly gives the intuition for the proofs of the relations and separations of the security notions; details follow in the rest of the section.

- IND-CPA \implies CCI: A (non-compressing) IND-CPA-secure symmetric-key encryption scheme provides indistinguishability of any pair of equal-length chosen messages, including messages involving a cookie. The proof proceeds by a hybrid argument, making the cookie used in each query made by the adversary to its E_1 oracle independent of the secret bit b .
- CCI $\not\Rightarrow$ IND-CPA: A degenerate scheme that uses a separating-secrets filter to extract secret cookies then encrypt the cookies but not the non-cookie data is CCI-secure but not IND-CPA-secure for the whole message.
- CCI \implies RCI: A straightforward simulation: an adversary who cannot distinguish between encryptions of equal-length cookies of its choosing can also not distinguish between encryptions of randomly chosen equal-length cookies.
- RCI $\not\Rightarrow$ CCI: A counterexample is constructed that uses a separating-secrets filter: an extra ciphertext component c_2 is added, consisting of a point function applied to the separated secrets, where the point function is 1 on a single, publicly known cookie value z . With high probability, two randomly chosen cookies will not match z , so c_2 carries no useful information and the scheme is RCI-secure, but a CCI adversary can choose one cookie that matches z and one that does not, so c_2 allows distinguishing of the chosen cookies.
- RCI \implies CR: A straightforward simulation: an adversary who recovers a cookie given only ciphertexts easily distinguishes encryptions of cookies.
- CR $\not\Rightarrow$ RCI: A counterexample is constructed: an extra ciphertext component c_2 is added, consisting of a random oracle applied to the message. The adversary gets encryptions of $m' \| ck \| m''$ for m', m'' of its choice; without querying the random oracle on exactly $m' \| ck \| m''$, c_2 provides no information to the adversary, so the scheme is CR-secure. However, an RCI adversary can check the random oracle on the two given random cookies, so c_2 allows distinguishing of the given random cookies.

A.1 IND-CPA \implies CCI

Theorem 3. *Let Ψ be an IND-CPA-secure symmetric-key encryption scheme. Then Ψ is also a CCI-secure symmetric-key compression-encryption scheme for any cookie space CK . Formally,*

let \mathcal{A} be an adversary against the CCI security of Ψ , and let q denote the number of queries that \mathcal{A} makes to its E_1 oracle. Then

$$\text{Adv}_{\Psi, \mathcal{CK}}^{\text{CCI}}(\mathcal{A}) \leq q \cdot \text{Adv}_{\Psi}^{\text{IND-CPA}}(\mathcal{B}^{\mathcal{A}}) ,$$

where \mathcal{B} is an algorithm, constructed using the adversary \mathcal{A} as described in the proof, against the IND-CPA security of the underlying symmetric-key encryption scheme Ψ .

Proof. The proof proceeds in a sequence of games, using a hybrid approach. Each Game i proceeds as in the original CCI security experiment, except that the queries to E_1 are answered as in Figure 7. Let Adv^i denote the probability that game i outputs 1.

$E_1(m', m'')$
1: **if** query $\# \leq i$ **then**
2: **return** $\Psi.\text{Enc}_k(m' \| ck_0 \| m'')$
3: **else if** query $\# > i$ **then**
4: **return** $\Psi.\text{Enc}_k(m' \| ck_b \| m'')$

Figure 7: Oracle E_1 used in Game i in proof of Theorem 3.

Game 0. This is the original CCI security game for Π . By definition,

$$\text{Adv}_{\Psi, \mathcal{CK}}^{\text{CCI}}(\mathcal{A}) = \text{Adv}^0 .$$

Transition from Game $(i - 1)$ to Game i , $1 \leq i \leq q$. Each hybrid transition changes how one query is answered; if the adversary's behaviour differs because of the change in answering the query, we can construct a simulator \mathcal{B}_i that wins the IND-CPA game for Ψ , as shown in Figure 8. When the IND-CPA challenger uses $b = 0$, c^* is the encryption of $m' \| ck_{\hat{b}} \| m''$, so \mathcal{B}_i is playing game $(i - 1)$ with \mathcal{A} . When the IND-CPA challenger uses $b = 1$, c^* is the encryption of $m' \| ck_{\hat{b}} \| m''$, so \mathcal{B}_i is playing game i with \mathcal{A} . Thus,

$$|\text{Adv}^{i-1} - \text{Adv}^i| \leq \text{Adv}_{\Psi}^{\text{IND-CPA}}(\mathcal{B}_i^{\mathcal{A}}) .$$

<u>$\mathcal{B}_i^{\mathcal{A}, E}()$</u> 1: $(ck_0, ck_1, st) \xleftarrow{\$} \mathcal{A}^{E_2}()$ s.t. $ ck_0 = ck_1 $ 2: $\hat{b} \xleftarrow{\$} \{0, 1\}$ 3: $b' \xleftarrow{\$} \mathcal{A}^{E_1, E_2}(ck_0, ck_1, st)$ 4: return b'	<u>$E_1(m', m'')$</u> 1: if query $\# < i$ then 2: return $E(m' \ ck_0 \ m'')$ 3: else if query $\# = i$ then 4: Give $(m' \ ck_{\hat{b}} \ m'', m' \ ck_0 \ m'')$ to IND-CPA challenger 5: Receive c^* from IND-CPA challenger 6: return c^* 7: else if query $\# > i$ then 8: return $E(m' \ ck_{\hat{b}} \ m'')$
<u>$E_2(m)$</u> 1: return $E(m)$	

Figure 8: Simulator \mathcal{B}_i used in the proof of Theorem 3

Analysis of Game q . Since the adversary's view is independent of b in Game q , we have

$$\text{Adv}^q = 0 .$$

Conclusion. Combining the above results, we have

$$\text{Adv}_{\Psi, \mathcal{CK}}^{\text{CCI}}(\mathcal{A}) \leq \sum_{i=1}^q \text{Adv}_{\Psi}^{\text{IND-CPA}}(\mathcal{B}_i) = q \cdot \text{Adv}_{\Psi}^{\text{IND-CPA}}(\mathcal{B})$$

(with a small abuse of notation in creating a single \mathcal{B} from the disparate \mathcal{B}_i). \square

A.2 CCI $\not\Rightarrow$ IND-CPA

Theorem 4. *There exists a symmetric-key compression-encryption scheme that is CCI-secure but not IND-CPA-secure.*

Theorem 4 is shown using a degenerate counterexample involving the separating-secrets technique. The basic idea is that we encrypt *only* secret cookies and not the message. Technically, the CCI security definition only requires any confidentiality for the cookie portion of the ciphertext and not the rest of it, so a scheme that extracts and encrypts only the cookies is CCI-secure, but is clearly not IND-CPA-secure.

In particular, let Π be a symmetric-key encryption scheme. Let \mathcal{CK} be the cookie space recognized by $\text{secret}[A\text{-}Za\text{-}z0\text{-}9]^\lambda$ and f be the corresponding filter, as described in Section 3.1. Recall this filter is effective at separating out \mathcal{CK} .

We construct Ψ from Π and f as in Figure 9. We will show that Ψ is CCI-secure, but is not IND-CPA-secure.

<u>Ψ.KeyGen()</u>	<u>Ψ.Enc$_k$(m)</u>	<u>Ψ.Dec$_k$(c)</u>
1: return Π .KeyGen()	1: $pt_s \ pt_{ns} \leftarrow f(m)$ 2: $c_1 \xleftarrow{\$} \Pi$.Enc $_k$ (pt_s) 3: return $c_1 \ pt_{ns}$	1: Parse $c_1 \ pt_{ns} \leftarrow c$ 2: $pt_s \leftarrow \Pi$.Dec $_k$ (c_1) 3: $m \leftarrow f^{-1}(pt_s, pt_{ns})$ 4: return m

Figure 9: Scheme Ψ used in the proof of Theorem 4

Claim 2. Ψ in Figure 9 is CCI-secure, assuming Π is IND-CPA-secure.

Proof sketch of claim. Since f is effective at separating out \mathcal{CK} , only c_1 components carry any information about b . However, c_1 is the encryption of the secrets extracted from m' and m'' as well as ck_b . Since f is safe for \mathcal{CK} , the length of pt_s is the same when derived from either $m' \| ck_0 \| m''$ or $m' \| ck_1 \| m''$. Thus any adversary that can guess the bit b serves as a distinguisher for Π under chosen plaintext attack. \square

Claim 3. Ψ in Figure 9 is not IND-CPA-secure.

Proof sketch of claim. The construction of a successful \mathcal{A} against the IND-CPA security of Ψ is straightforward. \mathcal{A} picks two distinct messages m'_0 and m'_1 that do not match the regular expression defining f , and gives these as the challenge messages to the IND-CPA challenger for Ψ . The resulting ciphertext will have an empty c_1 component (since neither m'_0 nor m'_1 has any value that will be separated out), and the second ciphertext component is encrypted, so the adversary receives back m'_b directly. The adversary then can immediately provide a correct guess of b . \square

A.3 CCI \Rightarrow RCI

Theorem 5. *Let Ψ be a CCI-secure symmetric-key compression-encryption scheme. Then Ψ is also an RCI-secure symmetric-key compression-encryption scheme. Formally, let \mathcal{CK} be a cookie*

space, and let \mathcal{A} be an algorithm against the RCI security of Ψ . Then, for the algorithm \mathcal{B} given in Figure 10,

$$\text{Adv}_{\Psi, \mathcal{CK}}^{\text{RCI}}(\mathcal{A}) \leq \text{Adv}_{\Psi, \mathcal{CK}}^{\text{CCI}}(\mathcal{B}^{\mathcal{A}}) .$$

Proof. The proof proceeds via direct simulation. We are given an adversary \mathcal{A} against the RCI security of Ψ . We must construct an adversary \mathcal{B} against the CCI security of Ψ ; note that \mathcal{B} will have access to oracles E_1 and E_2 described in the CCI security experiment. The simulator \mathcal{B} is constructed in Figure 10. Notice in particular that \mathcal{B} uses the CCI challenger's E_1 and E_2 oracles to answer \mathcal{A} 's queries.

$\mathcal{B}^{\mathcal{A}, E_1, E_2}()$

- 1: $(ck_0, ck_1) \xleftarrow{\$} \mathcal{CK}$ s.t. $|ck_0| = |ck_1|$
- 2: Give (ck_0, ck_1) to the CCI challenger
- 3: $b' \xleftarrow{\$} \mathcal{A}^{E_1, E_2}(ck_0, ck_1, \perp)$
- 4: **return** b'

Figure 10: Simulator used in the proof of Theorem 5.

\mathcal{B} 's simulation of the RCI experiment to \mathcal{A} is perfect. If \mathcal{A} 's guess b' of the b is correct in the RCI experiment, then it is also correct for the CCI experiment, and similarly when the guess is wrong. This yields the bound in the theorem. \square

A.4 RCI $\not\Rightarrow$ CCI

Theorem 6. *There exists a symmetric-key compression-encryption scheme that is RCI-secure but not CCI-secure.*

Theorem 6 is shown using a counterexample involving the separating-secrets technique and a point function involving a hard-coded publicly known string. The basic idea is that we append to each ciphertext a single bit representing the output of the point function on the cookie(s) in the message. Since randomly chosen cookies are highly unlikely to be the same as the hard-coded value in the point function, this extra bit provides no useful information for an RCI adversary, but a chosen-cookie adversary could easily pick one cookie to be the hard-coded value and one not to be, the bit thereby allowing him to easily distinguish the two.

In particular, let Π be a symmetric-key encryption scheme and let Γ be a compression scheme. Let \mathcal{CK} be the cookie space recognized by $\text{secret}[\text{A-Za-z0-9}]^\lambda$ and f be the corresponding filter, as described in Section 3.1. Recall this filter is safe for \mathcal{CK} .

Let $\Psi = \Pi \circ \text{SS}_{f, \Gamma}$ be the symmetric-key compression-encryption scheme constructed using the separating-secrets technique. Recall from Theorem 1 that Ψ is CCI-secure if Π is IND-CPA secure. By Theorem 5, Ψ is thus also RCI-secure.

Additionally, let $z \in \mathcal{CK}$, and define the point function

$$g_z(x) = \begin{cases} 1, & \text{if } z \preceq x, \\ 0, & \text{otherwise.} \end{cases}$$

We will construct Ψ' as in Figure 11 from $\Psi = \Pi \circ \text{SS}_{f, \Gamma}$ and g_z . We will show that Ψ' remains RCI-secure, but is not CCI-secure.

Claim 4. Ψ' in Figure 11 is RCI-secure, assuming Π is IND-CPA-secure and \mathcal{CK} is large.

Proof sketch of claim. Let ck_0^* and ck_1^* denote the random cookies to be distinguished.

Because $\Pi \circ \text{SS}_{f, \Gamma}$ is RCI-secure, c_1 gives the adversary no advantage in guessing the bit b . We need to assess whether c_2 helps the adversary at all in winning the RCI game.

<u>$\Psi'.\text{KeyGen}()$</u>	<u>$\Psi'.\text{Enc}_k(m)$</u>	<u>$\Psi'.\text{Dec}_k(c)$</u>
1: return $\Psi.\text{KeyGen}()$	1: $pt_s \parallel \widetilde{pt}_{ns} \leftarrow \text{SS}_{f,\Gamma}.\text{Comp}(m)$ 2: $c_1 \xleftarrow{\$} \Pi.\text{Enc}_k(pt_s \parallel \widetilde{pt}_{ns})$ 3: $c_2 \leftarrow g_z(pt_s)$ 4: return $c_1 \parallel c_2$	1: Parse $c_1 \parallel c_2 \leftarrow c$ 2: return $\Psi.\text{Dec}_k(c_1)$

Figure 11: Scheme Ψ' used in the proof of Theorem 6

The second ciphertext component c_2 is only useful to the adversary if the adversary can construct a pair (m', m'') such that c_2 is different for $m' \parallel ck_0^* \parallel m''$ versus $m' \parallel ck_1^* \parallel m''$.

Consider the construction of pt_s from m in $\Psi'.\text{Enc}_k(m)$. Since pt_s consists of a comma-separated list of cookies, and no cookie contains a comma, $g_z(pt_s) = 1$ if and only if there exists some i such that $g_z(ck_i) = 1$, where m is parsed as $m_0 \parallel ck_1 \parallel m_1 \parallel ck_2 \parallel m_2 \parallel \dots \parallel ck_n \parallel m_n$ as described above.

Now consider the handling of $m = m' \parallel ck_b^* \parallel m''$ for m', m'' of the adversary's choosing. By the argument above, $g_z(pt_s) = g_z(ck_b^*) \vee g_z(m') \vee g_z(m'')$. Moreover, $g_z(ck_0^*) \neq g_z(ck_1^*)$ if and only if one of them is equal to z but the other is not. Since ck_0^* and ck_1^* are chosen uniformly at random from \mathcal{CK} , this occurs with probability at most $2/|\mathcal{CK}|$. \square

Claim 5. Ψ' in Figure 11 is not CCI-secure.

Proof sketch of claim. The construction of a successful \mathcal{A} against the CCI security of Ψ' is straightforward. Note that \mathcal{A} knows the value z in the point function g_z . Thus, \mathcal{A} could issue a CCI challenge with $ck_0 = z$ and $ck_1 \neq z$, and then make the query $E_1(\epsilon, \epsilon)$, thereby obtaining $c_1 = \text{Enc}_k(ck_b)$ and $c_2 = g_z(ck_b)$. The last bit of such a ciphertext immediately indicates whether $ck_b = z$ or not. \square

A.5 RCI \implies CR

Theorem 7. Let Ψ be an RCI-secure symmetric-key compression-encryption scheme. Then Ψ is also a CR-secure symmetric-key compression-encryption scheme. Formally, let \mathcal{CK} be a cookie space, and let \mathcal{A} be an algorithm against the CR security of Ψ . Then, for the algorithm \mathcal{B} given in Figure 12,

$$\text{Adv}_{\Psi, \mathcal{CK}}^{\text{CR}}(\mathcal{A}) \leq \text{Adv}_{\Psi, \mathcal{CK}}^{\text{RCI}}(\mathcal{B}^{\mathcal{A}}) .$$

Proof. The proof proceeds via direct simulation. We are given an adversary \mathcal{A} against the CR security of Ψ . We must construct an adversary \mathcal{B} against the RCI security of Ψ ; note that \mathcal{B} will have access to oracles E_1 and E_2 described in the RCI security experiment. The simulator \mathcal{B} is constructed in Figure 12. Notice in particular that \mathcal{B} uses the RCI challenger's E_1 and E_2 oracles to answer \mathcal{A} 's queries.

$\mathcal{B}^{\mathcal{A}, E_1, E_2}(ck_0, ck_1)$

- 1: $ck' \xleftarrow{\$} \mathcal{A}^{E_1, E_2}()$
- 2: **if** $ck' = ck_0$ **then**
- 3: **return** 0
- 4: **else if** $ck' = ck_1$ **then**
- 5: **return** 1
- 6: **else**
- 7: **return** $b' \xleftarrow{\$} \{0, 1\}$

Figure 12: Simulator used in the proof of Theorem 7.

\mathcal{B} 's simulation of the CR experiment to \mathcal{A} is perfect: the value ck_b used by the E_1 oracle in the RCI challenger was indeed chosen at random, just as in the CR experiment, and is indeed consistent.

If \mathcal{A} 's guess ck' of the cookie is correct in the CR simulation, then it is also correct for the RCI experiment, and so \mathcal{B} 's output will be correct. If \mathcal{A} 's guess is wrong, then \mathcal{B} does as good as random guessing. \square

A.6 CR $\not\Rightarrow$ RCI

Theorem 8. *There exists a symmetric-key compression-encryption scheme that is CR-secure but not RCI-secure.*

Theorem 8 is shown using a counterexample involving a random oracle. The basic idea is that we append to each ciphertext the output of the random oracle applied to the message. For an adversary who is trying to guess an unknown cookie, this provides no information unless it queries the random oracle on the cookie itself, but a random-cookie adversary, who knows that the cookie is one of two values, could easily determine which by querying both to the random oracle.

Let Ψ be a symmetric-key compression-encryption scheme. Let $H : \{0, 1\}^* \rightarrow \{0, 1\}^\lambda$ be a random oracle. Construct Ψ' as shown in Figure 13.

<u>$\Psi'.\text{KeyGen}()$</u>	<u>$\Psi'.\text{Enc}_k(m)$</u>	<u>$\Psi'.\text{Dec}_k(c)$</u>
1: return $\Psi.\text{KeyGen}()$	1: $c_1 \xleftarrow{\$} \Psi.\text{Enc}_k(m)$ 2: $c_2 \leftarrow H(m)$ 3: return $c_1 \ c_2$	1: Parse $c_1 \ c_2 \leftarrow c$ 2: return $\Psi.\text{Dec}_k(c_1)$

Figure 13: Scheme Ψ' used in the proof of Theorem 8

Claim 6. Ψ' is CR-secure in the random oracle model, assuming Ψ is CR-secure and \mathcal{CK} is sufficiently large. Formally,

$$\text{Adv}_{\Psi', \mathcal{CK}}^{\text{CR}}(\mathcal{A}) \leq q_H \cdot \ell \cdot \text{Adv}_{\Psi, \mathcal{CK}}^{\text{CR}}(\mathcal{B}^{\mathcal{A}}) ,$$

where q_H is the number of queries that \mathcal{A} makes to the random oracle, ℓ is the maximum length of a message queried by \mathcal{A} to the random oracle, and \mathcal{B} is the algorithm given in Figure 14.

Proof. Let ck^* denote the cookie to be recovered.

The intuition of the proof is as follows. Because Ψ is CR-secure, c_1 does not help the adversary in guessing the cookie ck^* . We need to assess whether c_2 helps the adversary at all in winning the CR game. The second ciphertext component c_2 is only useful to the adversary if the adversary queries the random oracle on the plaintext of c_1 , which would mean that the adversary queries the random oracle on the target cookie ck^* . We can thus use this to win the CR experiment for Ψ' .

More precisely, the proof proceeds via direct simulation. We are given an adversary \mathcal{A} against the CR security of Ψ' . We must construct an adversary \mathcal{B} against the CR security of Ψ ; note that \mathcal{B} will have access to oracles E_1 and E_2 described in the CR security experiment for Ψ . The simulator \mathcal{B} is constructed in Figure 14.

\mathcal{B} 's simulation of the CR experiment for Ψ' to \mathcal{A} is perfect so long as E'_1 and H remain consistent. The only time an inconsistency arises is if \mathcal{A} queries H on $m' \| ck^* \| m''$ for some m', m'' that it also queries to E'_1 . With probability $1/q$, \mathcal{B} will correctly guess that the first such query to H is the i th query. Moreover, with probability at least $1/\ell$, \mathcal{B} will correctly guess which substring of that query to H is the target cookie c^* . For simplicity, we ignore the possibility of collisions on the output of the random oracle. \square

$\underline{\mathcal{B}^{\mathcal{A}, E_1, E_2}()}$ 1: $i \xleftarrow{\$} \{1, \dots, q_H\}$ 2: $j \xleftarrow{\$} \{1, \dots, \ell\}$ 3: $\text{rand} \xleftarrow{\$} \mathcal{CK}$ 4: $ck' \xleftarrow{\$} \mathcal{A}^{E_1', E_2', H}()$	$\underline{E_2'(m)}$ 1: $c_1 \xleftarrow{\$} E_2(m)$ 2: $c_2 \leftarrow H(m)$ 3: return $c_1 \ c_2$
$\underline{E_1'(m', m'')}$ 1: $c_1 \xleftarrow{\$} E_1(m', m'')$ 2: $c_2 \leftarrow H(m' \ \text{rand} \ m'')$ 3: return $c_1 \ c_2$	$\underline{H(m)}$ 1: if query $\# = i$ then 2: Parse m to identify every substring of m that is in \mathcal{CK} 3: Pick one uniformly at random 4: Output it to the CR challenger for Ψ 5: else 6: Answer $H(m)$ as normal for a random oracle

Figure 14: Simulator used in the proof of Theorem 8.

Claim 7. Ψ' is not RCI-secure.

Proof sketch of claim. The construction of a successful \mathcal{A} against the RCI security of Ψ' is straightforward. An adversary \mathcal{A} against the RCI security of Ψ' is told that the target cookie is one of two values, ck_0 and ck_1 . \mathcal{A} could make the query $E_1(\epsilon, \epsilon)$, thereby obtaining $c_1 = \text{Enc}_k(ck_b)$ and $c_2 = H(ck_b)$. \mathcal{A} could then query $H(ck_0)$ and $H(ck_1)$; one of these will equal c_2 , telling the adversary the value of b . \square

B Proof of CCI security of separating-secrets technique

Proof of Theorem 1. The proof proceeds in a sequence of games, using a hybrid approach. Each Game i proceeds as in the original CCI security experiment, except that the queries to E_1 are answered as in Figure 15. Let Adv^i denote the probability that game i outputs 1.

$\underline{E_1(m', m'')}$ 1: if query $\# \leq i$ then 2: return $\Pi.\text{Enc}_k(\text{SS}_{f, \Gamma}(m' \ ck_0 \ m''))$ 3: else if query $\# > i$ then 4: return $\Pi.\text{Enc}_k(\text{SS}_{f, \Gamma}(m' \ ck_b \ m''))$
--

Figure 15: Oracle E_1 used in Game i in proof of Theorem 1.

Game 0. This is the original CCI security game for Π . By definition,

$$\text{Adv}_{\Pi \circ \text{SS}_{f, \Gamma}, \mathcal{CK}}^{\text{CCI}}(\mathcal{A}) = \text{Adv}^0 .$$

Transition from Game $(i-1)$ to Game i , $1 \leq i \leq q$. Each hybrid transition changes how one query is answered; if the adversary's behaviour differs because of the change in answering the query, we can construct a simulator \mathcal{B}_i that wins the IND-CPA game for Ψ , as shown in Figure 16. When the IND-CPA challenger uses $b = 0$, c^* is the encryption of the separating-secrets compression of $m' \| ck_b \| m''$, so \mathcal{B}_i is playing game $(i-1)$ with \mathcal{A} . When the IND-CPA challenger uses $b = 1$, c^* is the encryption of the separating-secrets compression of $m' \| ck_0 \| m''$, so \mathcal{B}_i is playing game i with \mathcal{A} . Since f is safe for \mathcal{CK} , the separating-secrets compressions of $m' \| ck_0 \| m''$ and $m' \| ck_1 \| m''$ have the same length, and thus the pair of chosen messages given from the simulator in E_1 to the IND-CPA challenger is valid according to the IND-CPA experiment. Thus,

$$|\text{Adv}^{i-1} - \text{Adv}^i| \leq \text{Adv}_{\Psi}^{\text{IND-CPA}}(\mathcal{B}_i^{\mathcal{A}}) .$$

$\mathcal{B}_i^{A,E}()$ 1: $(ck_0, ck_1, st) \xleftarrow{\$} \mathcal{A}^{E_2}()$ s.t. $ ck_0 = ck_1 $ 2: $\hat{b} \xleftarrow{\$} \{0, 1\}$ 3: $b' \xleftarrow{\$} \mathcal{A}^{E_1, E_2}(ck_0, ck_1, st)$ 4: return b'	$E_1(m', m'')$ 1: if query $\# < i$ then 2: return $E(SS_{f,\Gamma}.Comp(m' \ ck_0 \ m''))$ 3: else if query $\# = i$ then 4: $pt \ \widetilde{pt}_{ns} \leftarrow SS_{f,\Gamma}.Comp(m' \ ck_{\hat{b}} \ m'')$ 5: $pt' \ \widetilde{pt}_{ns}' \leftarrow SS_{f,\Gamma}.Comp(m' \ ck_0 \ m'')$ 6: Give $(pt \ \widetilde{pt}_{ns}, pt' \ \widetilde{pt}_{ns}')$ to IND-CPA challenger 7: Receive c^* from IND-CPA challenger 8: return c^* 9: else if query $\# > i$ then 10: return $E(SS_{f,\Gamma}.Comp(m' \ ck_{\hat{b}} \ m''))$
$E_2(m)$ 1: return $E(SS_{f,\Gamma}.Comp(m))$	

Figure 16: Simulator \mathcal{B}_i used in the proof of Theorem 1

Analysis of Game q . Since the adversary's view is independent of b in Game q , we have

$$\text{Adv}^q = 0 .$$

Conclusion. Combining the above results, we have

$$\text{Adv}_{\Psi, \mathcal{CK}}^{\text{CCI}}(\mathcal{A}) \leq \sum_{i=1}^q \text{Adv}_{\Psi}^{\text{IND-CPA}}(\mathcal{B}_i^{\mathcal{A}}) = q \cdot \text{Adv}_{\Psi}^{\text{IND-CPA}}(\mathcal{B}^{\mathcal{A}})$$

(with a small abuse of notation in creating a single \mathcal{B} from the disparate \mathcal{B}_i). □

C Analysis of security of fixed-dictionary technique

C.1 Probability bounds, no prefix/suffix

In this section, we compute the amount of information given to the adversary from knowing the length of the compressed cookie, without any adversarially chosen prefix or suffix. This can be computed by calculating the amount of information given by knowing how many substrings of the cookie appear in the dictionary. For the analysis, we treat \mathcal{D} as a set of strings.

First we calculate the probability that a given string is a substring of a randomly chosen cookie.

Lemma 1. *Let $x \in \Omega^w$ be a word, and let $ck \xleftarrow{\$} \Omega^n = \mathcal{CK}$ be a random string of n characters. Then*

$$\Pr(x \preceq ck) \leq 1 - \left(1 - \frac{1}{|\Omega|^w}\right)^{n-w+1} .$$

Proof.

$$\begin{aligned} \Pr(x \preceq ck) &= 1 - \Pr(x \not\preceq ck) \\ &= 1 - \Pr((x \neq ck_{1:w}) \wedge (x \neq ck_{2:w}) \wedge \dots \wedge (x \neq ck_{n-w+1:w})) \\ &\leq 1 - \Pr(x \neq ck_{1:w}) \Pr(x \neq ck_{2:w}) \dots \Pr(x \neq ck_{n-w+1:w}) \\ &= 1 - \left(1 - \frac{1}{|\Omega|^w}\right)^{n-w+1} \end{aligned} \quad \square$$

We now compute that probability that one of a set of given strings is a substring of a randomly chosen cookie:

Lemma 2. Let $\mathcal{D} \subseteq \Omega^w$ with $|\mathcal{D}| = d$ be a dictionary of d words of w characters. Let $ck \stackrel{\$}{\leftarrow} \Omega^n = \mathcal{CK}$ be a random string of n characters. Then

$$\Pr(\exists x \in \mathcal{D} : x \preceq ck) \leq d \left(1 - \left(1 - \frac{1}{|\Omega|^w} \right)^{n-w+1} \right) .$$

Proof. Suppose $\mathcal{D} = \{x_1, x_2, \dots, x_d\}$.

$$\begin{aligned} \Pr(\exists x \in \mathcal{D} : x \preceq ck) &= \Pr((x_1 \preceq ck) \vee (x_2 \preceq ck) \vee \dots \vee (x_d \preceq ck)) \\ &\leq \sum_{i=1}^d \Pr(x_i \preceq ck) \\ &\leq d \left(1 - \left(1 - \frac{1}{|\Omega|^w} \right)^{n-w+1} \right) \end{aligned} \quad \square$$

Recall the definition of conditional entropy for random variables X and Y :

$$\begin{aligned} H(Y | X) &= \sum_{x \in \text{supp}(X)} \Pr(X = x) H(Y | X = x) \\ &= - \sum_{x \in \text{supp}(X)} \Pr(X = x) \\ &\quad \cdot \sum_{y \in \text{supp}(Y)} \Pr(Y = y | X = x) \log_2 \Pr(Y = y | X = x) . \end{aligned}$$

We now compute the amount of entropy about the cookie given knowledge about the number of substrings of the cookie that appear in the dictionary:

Lemma 3. Fix \mathcal{D} . Let $\#\text{SUB}(ck)$ denote the number of substrings of ck that appear in \mathcal{D} . Suppose CK is a uniform random variable on \mathcal{CK} . Then

$$\begin{aligned} H(CK | \#\text{SUB}(CK)) &\geq \left(1 - d \left(1 - \left(1 - \frac{1}{|\Omega|^w} \right)^{n-w+1} \right) \right) \\ &\quad \cdot \log_2 \left(|\mathcal{CK}| - |\mathcal{CK}| \cdot d \left(1 - \left(1 - \frac{1}{|\Omega|^w} \right)^{n-w+1} \right) \right) . \end{aligned}$$

Proof. Let $\#_s$ denote the number of cookies $ck \in \mathcal{CK}$ such that $\#\text{SUB}(ck) = s$. First note that

$$\Pr(\#\text{SUB}(CK) = s) = \frac{\#_s}{|\mathcal{CK}|} .$$

Additionally,

$$\Pr(CK = ck | \#\text{SUB}(CK) = s) = \begin{cases} \frac{1}{\#_s}, & \text{if } \#\text{SUB}(CK) = s , \\ 0, & \text{otherwise} . \end{cases}$$

Substituting into the definition of conditional entropy, $H(CK | \#\text{SUB}(CK))$

$$\begin{aligned} &= - \sum_{s \in \mathbb{N}} \Pr(\#\text{SUB}(CK) = s) \\ &\quad \cdot \sum_{ck \in \mathcal{CK}} \Pr(CK = ck | \#\text{SUB}(CK) = s) \log_2 \Pr(CK = ck | \#\text{SUB}(CK) = s) \\ &= - \sum_{s \in \mathbb{N}} \frac{\#_s}{|\mathcal{CK}|} \sum_{ck \in \mathcal{CK} : \#\text{SUB}(CK) = s} \frac{1}{\#_s} \log_2 \frac{1}{\#_s} \\ &= \frac{1}{|\mathcal{CK}|} \sum_{s \in \mathbb{N}} \#_s \log_2 \#_s . \end{aligned}$$

Let $\#_{\geq 1}$ denote the number of cookies $ck \in \mathcal{CK}$ such that $\#\text{SUB}(ck) \geq 1$. Then

$$\begin{aligned} \Pr(\#\text{SUB}(CK) \geq 1) &= \Pr(\exists x \in \mathcal{D} : x \preceq ck) = \frac{\#_{\geq 1}}{|\mathcal{CK}|} && \text{(by definition of } \#_{\geq 1}\text{)} \\ &\leq d \left(1 - \left(1 - \frac{1}{|\Omega|^w} \right)^{n-w+1} \right) && \text{(by Lemma 2)} \end{aligned}$$

Thus, the number of cookies with at least 1 substring in the dictionary is

$$\#_{\geq 1} \leq |\mathcal{CK}| \cdot d \left(1 - \left(1 - \frac{1}{|\Omega|^w} \right)^{n-w+1} \right) .$$

Consequently, the number of cookies with no substring in the dictionary is

$$\#_0 = |\mathcal{CK}| - \#_{\geq 1} \geq |\mathcal{CK}| - |\mathcal{CK}|d \left(1 - \left(1 - \frac{1}{|\Omega|^w} \right)^{n-w+1} \right) .$$

Finally,

$$\begin{aligned} H(CK \mid \#\text{SUB}(CK)) &= \frac{1}{|\mathcal{CK}|} \sum_{s \in \mathbb{N}} \#_s \log_2 \#_s \geq \frac{1}{|\mathcal{CK}|} \#_0 \log_2 \#_0 \\ &\geq \left(1 - d \left(1 - \left(1 - \frac{1}{|\Omega|^w} \right)^{n-w+1} \right) \right) \\ &\quad \cdot \log_2 \left(|\mathcal{CK}| - |\mathcal{CK}| \cdot d \left(1 - \left(1 - \frac{1}{|\Omega|^w} \right)^{n-w+1} \right) \right) \quad \square \end{aligned}$$

For example, if we have 16-byte cookies ($\mathcal{CK} = \{0\mathbf{x}00, \dots, 0\mathbf{x}\mathbf{FF}\}^{16}$), and the dictionary \mathcal{D} is a set of $d = 4096$ words of length $w = 4$ bytes, then

$$H(CK \mid \#\text{SUB}(CK)) \geq 127.998395 .$$

Concluding our analysis of the information learned given to the adversary without any adversarially chosen prefix or suffix, we give a bound on the amount of entropy about the cookie given the length of the compressed cookie:

Lemma 4. *Fix \mathcal{D} with d words of length w over character set Ω . Denote the length of a cookie ck compressed with dictionary \mathcal{D} by $\text{COMPLEN}(ck) = |\text{FD}_{\mathcal{D},w,\ell} \cdot \text{Comp}(ck)|$. Suppose CK is a uniform random variable on \mathcal{CK} . Then*

$$\begin{aligned} H(CK \mid \text{COMPLEN}(CK)) &\geq H(CK \mid \#\text{SUB}(CK)) \\ &\geq \left(1 - d \left(1 - \left(1 - \frac{1}{|\Omega|^w} \right)^{n-w+1} \right) \right) \\ &\quad \cdot \log_2 \left(|\mathcal{CK}| - |\mathcal{CK}| \cdot d \left(1 - \left(1 - \frac{1}{|\Omega|^w} \right)^{n-w+1} \right) \right) . \end{aligned}$$

Lemma 4 follows from the data processing inequality and Lemma 3.

C.2 Probability bounds, prefix/suffix

Suppose CK is a uniform random variable on $\mathcal{CK} = \Omega^n$. We know that $H(CK) = n \log_2(|\Omega|)$. Trivially, $H(CK | CK_1) = (n-1) \log_2(|\Omega|)$, where CK_1 is the first character of CK . Similarly, $H(CK | CK_{1:a}) = (n-a) \log_2(|\Omega|)$ and finally $H(CK | CK_{1:a}, CK_{n-b:b}) = (n-a-b) \log_2(|\Omega|)$.

Consider the following CRIME-like attack on the beginning of the cookie. Let \mathcal{D} be a dictionary with d words of length w over character set Ω . Let $ck \in \Omega^n$. Let $O(\cdot)$ be an oracle that, upon input a of length $w-m$, with $1 \leq m \leq w-1$, returns 1 if and only if $a||ck_{1:m} \in \mathcal{D}$.

The CRIME-like attack works as follows:

1. For each $x \in \mathcal{D}$, query $x_{1:w-1}$ to the oracle. If a query for $x_{1:w-1}$ returns 1, then it is known that $ck_{1:1} \in Z_1 = \{z : x_{1:w-1}||z \in \mathcal{D}\}$. If no query returns 1, then return \emptyset .
2. For $m = 2, \dots, w-1$: For each $x \in \mathcal{D}$ such that $x_{w-m} \in Z_{m-1}$, query $x_{1:w-m}$ to the oracle. If a query for $x_{1:w-m}$ returns 1, then it is known that $ck_{1:m} \in Z_m = \{z_1 z_2 \dots z_m : x_{1:w-m}||z_1 z_2 \dots z_m \in \mathcal{D}\}$. If no query returns 1, then return Z_1, \dots, Z_{m-1} .
3. Return Z_1, \dots, Z_{w-1} .

A corresponding attack on the suffix is obvious.

Let $\text{CRIMEpre}(ck)$ denote the output obtained from running the above prefix CRIME attacks on ck , $\text{CRIMEsuf}(ck)$ denote the output from the corresponding suffix attack. Let $\text{CRIME}(ck) = (\text{CRIMEpre}(ck), \text{CRIMEsuf}(ck))$.

Noting that in the best case the CRIME attack allows the attacker to learn the first $w-1$ and the last $w-1$ characters of the cookie, some trivial lower bounds are:

$$\begin{aligned} H(CK_{1:w-1} | \text{CRIME}(CK)) &\geq 0 \\ H(CK_{n-w+1:w-1} | \text{CRIME}(CK)) &\geq 0 \end{aligned}$$

However, the CRIME attack provides no information about the remaining characters, so $I(CK_{1:w-1}, CK_{w:n-w+1}) = 0$ and $I(CK_{1:n-w+1}, CK_{n-w+1:w-1}) = 0$, and thus $H(CK_{w:n-w+2} | \text{CRIME}(CK), \text{COMPLEN}(CK)) = H(CK_{w:n-w+2} | \text{COMPLEN}(CK))$.

Finally, we have that

$$\begin{aligned} &H(CK | \text{CRIME}(CK), \text{COMPLEN}(CK)) \\ &\geq H(CK_{1:w-1} | \text{CRIMEpre}(CK)) + H(CK_{w:n-w+2} | \text{COMPLEN}(CK)) \\ &\quad + H(CK_{n-w+1:w-1} | \text{CRIMEsuf}(CK)) \\ &\geq 0 + H(CK_{w:n-w+2} | \text{COMPLEN}(CK)) + 0 \end{aligned}$$

and we can obtain a lower bound on $H(CK_{w:n-w} | \text{COMPLEN}(CK))$ using Lemma 4.