

# On the Security of RC4 in TLS and WPA\*

Nadhem J. AlFardan<sup>1</sup>

Daniel J. Bernstein<sup>2</sup>

Kenneth G. Paterson<sup>1</sup>

Bertram Poettering<sup>1</sup>

Jacob C.N. Schuldt<sup>1</sup>

<sup>1</sup> Information Security Group  
Royal Holloway, University of London

<sup>2</sup> University of Illinois at Chicago and Technische Universiteit Eindhoven

July 8th 2013

## Abstract

The Transport Layer Security (TLS) protocol aims to provide confidentiality and integrity of data in transit across untrusted networks. TLS has become the *de facto* protocol standard for secured Internet and mobile applications. TLS supports several symmetric encryption options, including a scheme based on the RC4 stream cipher. In this paper, we present ciphertext-only plaintext recovery attacks against TLS when RC4 is selected for encryption. Variants of these attacks also apply to WPA, a prominent IEEE standard for wireless network encryption. Our attacks build on recent advances in the statistical analysis of RC4, and on new findings announced in this paper. Our results are supported by an experimental evaluation of the feasibility of the attacks. We also discuss countermeasures.

## 1 Introduction

TLS is arguably the most widely used secure communications protocol on the Internet today. Starting life as SSL, the protocol was adopted by the IETF and specified as an RFC standard under the name of TLS 1.0 [7]. It has since evolved through TLS 1.1 [8] to the current version TLS 1.2 [9]. Various other RFCs define additional TLS cryptographic algorithms and extensions. TLS is now used for securing a wide variety of application-level traffic: It serves, for example, as the basis of the HTTPS protocol for encrypted web browsing, it is used in conjunction with IMAP or SMTP to cryptographically protect email traffic, and it is a popular tool to secure communication with embedded systems, mobile devices, and in payment systems.

Technically speaking, TLS sessions consist of two consecutive phases: the execution of the TLS Handshake Protocol which typically deploys asymmetric techniques to establish a secure session key, followed by the execution of the TLS Record Protocol which uses symmetric key cryptography (block ciphers, the RC4 stream cipher, MAC algorithms) in combination with the established session key and sequence numbers to build a secure channel for transporting application-layer data. In the Record Protocol, there are mainly three encryption options:

---

\*The research of the third, fourth and fifth authors was supported by an EPSRC Leadership Fellowship, EP/H005455/1. The research of the second author was supported by the National Science Foundation under grant 1018836 and by the Netherlands Organisation for Scientific Research (NWO) under grant 639.073.005.

- HMAC followed by CBC-mode encryption using a block cipher,
- HMAC followed by encryption using the RC4 stream cipher, or
- authenticated encryption using GCM or CCM mode of operation of a block cipher.

The third of these three options is only available with TLS 1.2 [24, 19], which is yet to see widespread adoption.<sup>1</sup> The first option has seen significant cryptanalysis (padding oracle attacks [6], BEAST [10], Lucky 13 [3]). While countermeasures to the attacks on CBC-mode in TLS exist, many commentators now recommend, and many servers now offer, RC4-based encryption options ahead of CBC-mode.<sup>2</sup> Indeed, the ICSI Certificate Notary<sup>3</sup> recently performed an analysis of 16 billion TLS connections and found that around 50% of the traffic was protected using RC4 ciphersuites [4].

This makes it timely to examine the security of RC4 in TLS. While the RC4 algorithm is known to have a variety of cryptographic weaknesses (see [26] for an excellent survey), it has not been previously explored how these weaknesses can be exploited in the context of TLS. Here we show that new and recently discovered biases in the RC4 keystream do create serious vulnerabilities in TLS when using RC4 as its encryption algorithm.

While the main focus of this paper lies on the security of RC4 in TLS, our attacks (or variants thereof) might also be applicable to other protocols where RC4 is meant to ensure data confidentiality. Indeed, the WPA protocol used for encrypting wireless network traffic also utilizes the RC4 stream cipher in a way that allows (partial) plaintext recovery in specific settings — using basically the same attack strategies as in the TLS case.

We hope that this work will help spur the adoption of TLS 1.2 and its authenticated encryption algorithms, as well as the transition from WPA to (the hopefully more secure) WPA2.

## 1.1 Overview of Results

We present two plaintext recovery attacks on RC4 that are exploitable in specific but realistic circumstances when this cipher is used for encryption in TLS. Both attacks require a fixed plaintext to be RC4-encrypted and transmitted many times in succession (in the same, or in multiple independent RC4 keystreams). Interesting candidates for such plaintexts include passwords and, in the setting of secure web browsing, HTTP cookies.

A statistical analysis of ciphertexts forms the core of our attacks. We stress that the attacks are ciphertext-only: no sophisticated timing measurement is needed on the part of the adversary, the attacker does not need to be located close to the server, and no packet injection capability is required (all premises for Lucky 13). Instead, it suffices for the adversary to record encrypted traffic for later offline analysis. Provoking the required repeated encryption and transmission of the target plaintext, however, might require more explicit action: e.g., resetting TCP connections or guiding the victim to a website with specially prepared JavaScript (see examples below).

Since both our attacks require large amounts of ciphertext, their practical relevance could be questioned. However, they do show that the strength of RC4 in TLS is much lower than the employed 128-bit key would suggest. We freely admit that our attacks are not particularly deep, nor sophisticated: they only require an understanding of how TLS uses RC4, solid statistics on

---

<sup>1</sup>SSL Pulse (<https://www.trustworthyinternet.org/ssl-pulse/>) reported in June 2013 that only 15.1% of 170,000 websites surveyed support TLS 1.2; most major browsers currently do not support TLS 1.2.

<sup>2</sup>For examples of RC4 being recommended in the face of CBC attacks, see advice at Qualys' website <https://community.qualys.com/blogs/securitylabs/2011/10/17/mitigating-the-beast-attack-on-tls>, Ivan Ristic's personal blog <http://blog.ivanristic.com/2009/08/is-rc4-safe-for-use-in-ssl.html>, PhoneFactor's blog <http://blog.phonefactor.com/2011/09/23/slaying-beast-mitigating-the-latest-ssltls-vulnerability>, and F5's suggested Lucky 13 mitigation at <http://support.f5.com/kb/en-us/solutions/public/14000/100/sol14190.html>. Other examples abound on discussion forums and vendor websites.

<sup>3</sup><http://notary.icsi.berkeley.edu>

the biases in RC4 keystreams, and some experience of how modern browsers handle cookies. We consider it both surprising and alarming that such simple attacks are possible for such an important and heavily-studied protocol as TLS. We further discuss the implications of our attack in Section 7.

### 1.1.1 Our single-byte bias attack

Our first attack targets the initial 256 bytes of RC4 ciphertext. It is fixed-plaintext and multi-session, meaning that it requires a fixed sequence of plaintext bytes to be independently encrypted under a large number of (random) keys. This setting corresponds to what is called a “broadcast attack” in [18, 16, 26]. As we argue below, such attacks are a realistic attack vector in TLS. Observe that, in TLS, the first 36 bytes of the RC4 keystream are used to encrypt a TLS Handshake `Finished` message. This message is not fixed across TLS sessions. As a consequence, our methods can be applied only to recover up to 220 bytes of the TLS application plaintext.

Our attack exploits statistical biases occurring in the first 256 bytes of RC4 keystream. Such biases, i.e., deviations from uniform in the distributions of the keystream bytes at certain positions, have been reported and theoretically analyzed by [18], [16], and [26]. The corresponding authors also propose algorithms to exploit such biases for plaintext recovery. In this paper, we discuss shortcomings of their algorithms, empirically obtain a *complete view* of all single-byte biases occurring in the first 256 keystream positions, and propose a generalized algorithm that fully exploits all these biases for advanced plaintext recovery. As a side result of our research, in Section 3.1 we report on significant biases in the RC4 keystream that seemingly follow specific patterns and that have not been identified or analysed previously.

For concreteness, we describe how our single-byte bias attack could be applied to recover cookies in HTTPS traffic. Crucial here is to find an automated mechanism for efficiently generating a large number of encryptions of the target cookie. In line with the scenario employed by the BEAST and Lucky 13 attacks against CBC-mode encryption in TLS [3, 10], a candidate mechanism is for JavaScript malware downloaded from an attacker-controlled website and running in the victim’s browser to repeatedly send HTTPS requests to a remote server. The corresponding cookies are automatically included in each of these requests in a predictable location, and can thus be targeted in our attack. If client and server are configured to use TLS session resumption, the renewal of RC4 keys could be arranged to happen with particularly high frequency — as required for our attack to be successful.<sup>4</sup> Alternatively, the attacker can cause the TLS session to be terminated after the target encrypted cookie is sent; the browser will automatically establish a new TLS session when the next HTTPS request is sent.

As a second example, consider the case where IMAP passwords<sup>5</sup> are attacked. In a setup where an email client regularly connects to an IMAP server for (password-authenticated) mail retrieval, let the adversary reset the TCP connection between client and server immediately after the encrypted password is transmitted. In some client configurations this might trigger an automatic resumption of the session, including a retransmission of the (encrypted) password. If this is the case, the adversary is in the position to harvest a large set of independently encrypted copies of the password —one per reset— precisely fulfilling the precondition of our attack.

Our single-byte bias attack is on the verge of practicality. In our experiments, the first 40 bytes of TLS application data after the `Finished` message were recovered with a success rate of over 50% per byte, using  $2^{26}$  sessions. With  $2^{32}$  sessions, the per-byte success rate is more than 96% for the first 220 bytes (and is 100% for all but 12 of these bytes). If, for example, a target plaintext byte is known to be a character from a set of cardinality 16 (e.g., in a 4-bits-

---

<sup>4</sup>Unfortunately, we do not currently know of a way to trigger TLS session resumption from JavaScript running in a browser.

<sup>5</sup>The Internet Message Access Protocol (IMAP) is a popular protocol for email retrieval.

per-byte-encoded HTTP cookie), our algorithm recovers the first 112 bytes of plaintext with a success rate of more than 50% per byte, using  $2^{26}$  sessions. For further details, see Section 5.

### 1.1.2 Our double-byte bias attack

As we have seen, our single-byte bias attack on RC4 is quite effective in recovering ‘early’ plaintext bytes in the fixed-plaintext multi-session setting. It has, however, a couple of limitations when it comes to attacking practical systems that employ TLS. Focussing on the recovery of cookies in HTTPS-secured web sessions, we note that modern web browsers typically send a large number of HTTP headers before any cookies (these headers carry information about the particular client or server software, accepted MIME types, compression options, etc.). In practice, cookie data appears only at positions that come after the attackable initial 220 bytes of the ciphertext<sup>6</sup>. Independently of this issue, in the attack scenarios proposed above, a large number of HTTPS sessions would have to be established and torn down again, inducing non-negligible computing and bandwidth overheads via the TLS Handshake. Lastly, it has been proposed to routinely drop the first few hundred keystream bytes of RC4 before starting encryption in order to avoid the relatively strong early keystream biases [20] — if this were to be implemented in TLS, our single-byte bias attack would effectively be defeated.

Complementary to our single-byte bias attack, we present a second fixed-plaintext ciphertext-only attack on RC4. It exploits biases that appear in the entire keystream (and not just in the first 256 positions) and does not assume, but tolerates, frequent changes of the encryption key. Our second attack hence covers some scenarios where our single-byte bias attack does not seem to be applicable; it would, for example, be able to recover cookies from (long-persisting) HTTPS sessions. It would also be applicable if the initial keystream bytes were to be discarded.

In contrast to our first attack, our second attack exploits certain biases in consecutive pairs of bytes in the RC4 keystream that were first reported by Fluhrer and McGrew [12]. We empirically evaluate the probability of occurrence for each possible pair of bytes beginning at each position (modulo 256), obtaining a complete view of the distributions of pairs of bytes in positions  $(i, i + 1)$  (modulo 256). Our analysis strongly suggests that there are no further biases in consecutive positions of the same strength as the Fluhrer-McGrew biases. We use the obtained results in a specially designed attack algorithm to recover repeatedly encrypted plaintexts.

Our double-byte bias attack is again close to being practical. In our experiments, we focus on our attack’s ability to correctly recover 16 consecutive bytes of plaintext, roughly equating to an HTTP cookie. With  $13 \cdot 2^{30}$  encryptions of the plaintext, we achieve a success rate of 100% in recovering all 16 bytes. We obtain better success rates for restricted plaintexts, as in the single-byte case. For further details, see Section 5.

## 1.2 Related Work

In independent and concurrent work, Isobe *et al.* [14] have considered the security of RC4 against broadcast attacks. They present attacks based on both single-byte and multi-byte biases. They identify three biases in the first output bytes  $Z_r$  of RC4 that we also identify (specifically, the biases towards  $Z_3 = 0x83$ ,  $Z_r = r$ , and  $Z_r = -r$  when  $r$  is a multiple of 16) as well as a new conditional bias  $Z_1 = 0 | Z_2 = 0$ .

The single-byte bias attack in [14] only considers the *strongest* bias at each position, whereas our single-byte bias attack simultaneously exploits *all* biases in each keystream position. Specifically, we use Bayes’s law to compute the *a posteriori* plaintext distribution from the *a priori* plaintext distribution and the precomputed distributions of the  $Z_r$ . This explains why our single-byte attack out-performs that of [14]. For example, we achieve reliable plaintext recovery

---

<sup>6</sup>Note that when attacking secret URL parameters from HTTPS connections or passwords from IMAP sessions such limitations do not arise.

in the first 256 positions with  $2^{32}$  ciphertexts, while Isobe *et al.* [14] require  $2^{34}$  ciphertexts. We also achieve uniformly higher success rates for lower numbers of sessions. Previous authors exploring broadcast attacks on RC4 also only used single biases, leading to attacks that simply do not work [16, 26] or which have inferior performance to ours [25].

The multi-byte bias attack in [14] exploits the positive bias towards the pattern *ABSAB* that was identified by Mantin [17]. Here *A* and *B* are keystream bytes and *S* is a short string consisting of any keystream bytes (possibly of length 0). The attack in [14] assumes that 3-out-of-4 bytes in particular positions are known and uses the Mantin bias to recover the fourth. A limited experimental evaluation of the attack is reported in [14]: the attack is applied only to recovery of plaintext bytes 258-261, assuming all previous plaintext bytes have been successfully recovered, with success rates of 1 (for each of the 4 targeted bytes) using  $2^{34}$  ciphertexts. As explained in [14], this multi-byte attack would fail if the initial bytes of RC4 output were to be discarded. By contrast, our double-byte bias attack, which exploits the Fluhrer-McGrew biases, recovers more bytes with comparable success rate using slightly fewer ciphertexts and is resilient to initial byte discarding. It is an interesting open problem to determine whether the Mantin *ABSAB* bias can be combined with the Fluhrer-McGrew biases to gain enhanced attack performance.

A further point of comparison between our work and that of [14] concerns practical implementation. We have extensively explored the applicability of our attacks to RC4 as used in TLS and WPA, while [14] makes only brief mention of TLS in its concluding section and gives no mechanisms for generating the large numbers of ciphertexts needed for the attacks.

Finally, the authors of [14] claim in their abstract that their methods “*can recover the first  $2^{50}$  bytes  $\approx 1000$  T bytes of the plaintext, with probability close to 1, from only  $2^{34}$  ciphertexts*”. We point out that their methods would only recover  $2^{16}$  *distinct* bytes of output, rather than the advertised  $2^{50}$  bytes, since their attacks require the *same* plaintext to be encrypted  $2^{34}$  times. Furthermore, their multi-byte bias attack is not resilient to errors occurring in the recovery of early plaintext bytes (whereas ours is), so this claim would only be true if their multi-byte bias attack does not fail at any stage, and this is as yet untested.

### 1.3 Paper Organisation

Section 2 provides further background on the RC4 stream cipher and the TLS Record Protocol. Section 3 summarises weaknesses in RC4 that we exploit in our attacks. Section 4 describes our two plaintext recovery attacks on RC4. We evaluate the attacks in Section 5, with our main focus there being on TLS. The applicability of our attacks to WPA is analysed in Section 6. Finally, Section 7 discusses countermeasures to our attacks, and concludes with a recap of the main issues raised by our work.

## 2 Further Background

### 2.1 The RC4 Stream Cipher

The stream cipher RC4, originally designed by Ron Rivest, became public in 1994 and found application in a wide variety of cryptosystems; well-known examples include SSL/TLS, WEP [1], WPA [2], and some Kerberos-related encryption modes [15]. RC4 has a remarkably short description and is extremely fast when implemented in software. However, these advantages come at the price of lowered security: several weaknesses have been identified in RC4 [12, 11, 18, 17, 16, 26, 28, 27, 33], some of them being confirmed and exploited in the current paper.

Technically, RC4 consists of two algorithms: a *key scheduling algorithm* (KSA) and a *pseudo-random generation algorithm* (PRGA), which are specified in Algorithms 1 and 2. The KSA takes as input a key *K*, typically a byte-array of length between 5 and 32 (i.e., 40 to 256 bits), and produces the initial internal state  $st_0 = (i, j, S)$ , where *S* is the canonical representation

Algorithm 1: RC4 key scheduling (KSA)	Algorithm 2: RC4 keystream generator (PRGA)
<b>input</b> : key $K$ of $l$ bytes <b>output</b> : initial internal state $st_0$ <b>begin</b> <b>for</b> $i = 0$ <b>to</b> 255 <b>do</b> $\mathcal{S}[i] \leftarrow i$ $j \leftarrow 0$ <b>for</b> $i = 0$ <b>to</b> 255 <b>do</b> $j \leftarrow j + \mathcal{S}[i] + K[i \bmod l]$ $\text{swap}(\mathcal{S}[i], \mathcal{S}[j])$ $i, j \leftarrow 0$ $st_0 \leftarrow (i, j, \mathcal{S})$ <b>return</b> $st_0$	<b>input</b> : internal state $st_r$ <b>output</b> : keystream byte $Z_{r+1}$ updated internal state $st_{r+1}$ <b>begin</b> $\text{parse}(i, j, \mathcal{S}) \leftarrow st_r$ $i \leftarrow i + 1$ $j \leftarrow j + \mathcal{S}[i]$ $\text{swap}(\mathcal{S}[i], \mathcal{S}[j])$ $Z_{r+1} \leftarrow \mathcal{S}[\mathcal{S}[i] + \mathcal{S}[j]]$ $st_{r+1} \leftarrow (i, j, \mathcal{S})$ <b>return</b> $(Z_{r+1}, st_{r+1})$

Figure 1: Algorithms implementing the RC4 stream cipher. All additions are performed modulo 256.

of a permutation on the set  $[0, 255]$  as an array of bytes, and  $i, j$  are indices into this array. The PRGA will, given an internal state  $st_r$ , output ‘the next’ keystream byte  $Z_{r+1}$ , together with the updated internal state  $st_{r+1}$ . Particularly interesting to note is the fact that updated index  $j$  is computed in dependence on current  $i, j$ , and  $\mathcal{S}$ , while  $i$  is just a counter (modulo 256).

## 2.2 The TLS Record Protocol

We describe in detail the cryptographic operation of the TLS Record Protocol in the case that RC4 is selected as the encryption method.

Data to be protected by TLS is received from the application and may be fragmented and compressed before further processing. An individual record  $R$  (viewed as a sequence of bytes) is then processed as follows. The sender maintains an 8-byte sequence number  $\text{SQN}$  which is incremented for each record sent, and forms a 5-byte field  $\text{HDR}$  consisting of a 2-byte version field, a 1-byte type field, and a 2-byte length field. It then calculates an HMAC over the string  $\text{HDR}||\text{SQN}||R$ ; let  $T$  denote the resulting tag.

For RC4 encryption, record and tag are concatenated to create the plaintext  $P = R||T$ . This plaintext is then xored in a byte-by-byte fashion using the RC4 keystream, i.e., the ciphertext bytes are computed as

$$C_r = P_r \oplus Z_r \quad \text{for } r = 1, 2, 3, \dots,$$

where  $P_r$  are the individual bytes of  $P$ , and  $Z_r$  are the RC4 keystream bytes. The data transmitted over the wire then has the form

$$\text{HDR}||C,$$

where  $C$  is the concatenation of the bytes  $C_r$ .

The RC4 algorithm itself is initialised at the start of each TLS connection, using a 128-bit encryption key  $K$ . This key  $K$  is computed with a hash-function-based key derivation function from the TLS master secret that is established during the TLS Handshake Protocol. In more detail, the key  $K$  may be established either via a full TLS Handshake or via TLS session resumption. In a full TLS Handshake, a total of 4 communication round-trips are needed, and usually some public key cryptographic operations are required of both client and server. A full TLS Handshake run establishes a new TLS *session* and a new TLS master secret from which all other keys, including RC4 key  $K$ , are derived. TLS session resumption involves a lightweight version of the TLS Handshake Protocol being run to establish a new *connection* within an existing session: essentially, an exchange of nonces takes place, followed by an exchange of

**Finished** messages; no public key cryptographic operations are involved. The keys for the new connection, including  $K$ , are derived from the existing master secret and the new nonces. Given the design of the key derivation process, it is reasonable to model  $K$  as being uniformly random in the different sessions/connections.

The initialisation of RC4 in TLS is the standard one for this algorithm. Notably, none of the initial keystream bytes is discarded when RC4 is used in TLS, despite these bytes having known weaknesses. Note also that the first record sent under the protection of RC4 for each session or connection will be a **Finished** message, typically of length 36 bytes, consisting of a Handshake Protocol header, a PRF output, and a MAC on that output. This is typically 36 bytes in size (4 bytes of Handshake Protocol header, 12 bytes of PRF output, and 20 bytes of MAC in the case of HMAC-SHA-1). This record will not be targeted in our attacks, since it is not constant across multiple sessions.

The decryption process reverses this sequence of steps, but its details are not germane to our attacks. For TLS, any error arising during decryption should be treated as fatal, meaning an (encrypted) error message is sent to the sender and the session terminated with all keys and other cryptographic material being disposed of. This gives an attacker a convenient method to cause a session to be terminated and force new encryption and MAC keys to be set up. Another method is to somehow induce the client or server to initiate session resumption.

### 3 Biases in the RC4 Keystream

In this section, we summarise known biases in the RC4 keystream, and report new biases that we have observed experimentally.

#### 3.1 Single-byte Biases

The first significant bias in the RC4 keystream was observed by Mantin and Shamir in [18]. Their main result can be stated as:

**Result 1.** [18, Thm 1] *The probability that  $Z_2$ , the second byte of keystream output by RC4, is equal to 0x00 is approximately  $1/128$  (where the probability is taken over the random choice of the key).*

Since this result concerns only the second byte of the keystream, and this byte is always used to encrypt a **Finished** message in TLS, we are unable to exploit it in our attacks. More recently, the following result was obtained by Sen Gupta *et al.* in [26] as a refinement of an earlier result of Maitra *et al.* [16]:

**Result 2.** [26, Thm 14 and Cor 3] *For  $3 \leq r \leq 255$ , the probability that  $Z_r$ , the  $r$ -th byte of keystream output by RC4, is equal to 0x00 is*

$$\Pr(Z_r = 0x00) = \frac{1}{256} + \frac{c_r}{256^2} ,$$

where the probability is taken over the random choice of the key,  $c_3 = 0.351089$ , and  $c_4, c_5 \dots, c_{255}$  is a decreasing sequence with terms that are bounded as follows:

$$0.242811 \leq c_r \leq 1.337057.$$

In other words, bytes 3 to 255 of the keystream have a bias towards 0x00 of approximately  $1/2^{16}$ . This result was experimentally verified in [26] and found to be highly accurate (see Figure 11 of that paper). The biases here are substantially smaller than those observed in Result 1.

Additionally, Sen Gupta *et al.* [26] have identified a key-length-dependent bias in RC4 keystreams. Specifically, [26, Theorem 5] shows that when the key-length is  $\ell$  bytes, then byte  $Z_\ell$  is biased towards value  $256 - \ell$ , with the bias always being greater than  $1/2^{16}$ . For RC4 in TLS, we have  $\ell = 16$ .

Experimentally, we have observed additional biases in the RC4 keystream that do not yet have a theoretical explanation. As an example, Figure 2 shows the empirical distribution for the RC4 keystream bytes  $Z_1$ ,  $Z_{16}$ ,  $Z_{32}$  and  $Z_{50}$ , calculated over  $2^{44}$  independent, random 128-bit keys. For the first keystream byte,  $Z_1$ , we generally confirm the ‘sine-like’ behaviour also recently reported by Sen Gupta *et al.* [26]. We note, however, that the single strong bias *away* from value 0x81 (decimal 129) is not consistent with [26] which provides a theoretical treatment of the distribution of  $Z_1$ . The disparity likely arises because Sen Gupta *et al.* work with 256-byte keys, while our work is exclusively concerned with 128-bit (16-byte) keys as used in TLS; in other words, our observed bias of  $Z_1$  away from 0x81 seems to be key-length-dependent. For  $Z_{16}$ , we have three main biases: the bias towards 0x00, the very dominant key-length-dependent bias towards 0xF0 (decimal 240) from [26], and a new bias towards 0x10 (decimal 16). For  $Z_{32}$ , we also have three main biases: the bias towards 0x00, a large, new bias towards 0xE0 (decimal 224), and a new bias towards 0x20 (decimal 32). For  $Z_{50}$ , there are significant biases towards byte values 0x00 and 0x32 (decimal 50), as well as an upward trend in probability as the byte value increases.

Individual inspection of ciphertext distributions at all positions  $1 \leq r \leq 256$  reveals two new significant biases that occur with specific regularities: a bias towards value  $r$  for all  $r$ , and a bias towards value  $256 - r$  at positions  $r$  that are *multiples* of (key-length) 16; note that the latter finding both confirms and extends the results from [26]. Both of these new biases were also observed by Isobe *et al.* [14] in concurrent work, with a theoretical explanation being given for the bias towards  $r$ . Figure 3 shows the estimated strength of these biases in comparison with the strength of the bias towards 0x00 for the keystream bytes  $Z_1, \dots, Z_{256}$ . The estimates are based on the empirical distribution of the RC4 keystream bytes, calculated over  $2^{44}$  random 128-bit RC4 keys. We note that the key-length dependent bias dominates the other two biases until position  $Z_{112}$ , and that the bias of  $Z_r$  towards  $r$  dominates the bias towards 0x00 observed by [16] between positions  $Z_5$  and  $Z_{31}$ , except for byte  $Z_{16}$  where the bias towards 0x00 is slightly stronger. Finally, our computations have revealed a number of other, smaller biases in the initial bytes of the RC4 keystream. <sup>7</sup>

### 3.2 Multi-byte Biases

Besides the single-byte biases highlighted above, several multi-byte biases have been identified in the RC4 keystream. In contrast to the single-byte biases, most of the identified multi-byte biases are “long term” biases which appear periodically at regular intervals in the keystream.

The most extensive set of multi-byte biases was identified by Fluhrer and McGrew [12] who analyzed the distribution of pairs of byte values for consecutive keystream positions  $(Z_r, Z_{r+1})$ ,  $r \geq 1$ . More precisely, they estimated the distribution of consecutive keystream bytes for scaled-down<sup>8</sup> versions of RC4 by assuming an idealized internal state of RC4 in which the permutation  $\mathcal{S}$  and the internal variable  $j$  are random (see Figure 1), and then extrapolated the results to standard RC4.

The reported biases for standard RC4 are listed in Table 1. Note that all biases are dependent on the internal variable  $i$  which is incremented (modulo 256) for each keystream byte generated. It should also be noted that, due to the assumption that  $\mathcal{S}$  and  $j$  are random, the biases cannot be expected to hold for the initial keystream bytes. However, this idealization becomes a

<sup>7</sup>These can be seen pictorially at <http://www.isg.rhul.ac.uk/tls/biases.pdf>.

<sup>8</sup> In detail, instead of an internal permutation  $\mathcal{S}$  of 8-bit values, Fluhrer and McGrew consider variants of RC4 based on permutations of 3-bit, 4-bit, and 5-bit values, respectively. Note that in these versions of RC4, the internal variables  $i$  and  $j$ , as well as the output  $Z_r$ , will also be 3-bit, 4-bit and 5-bit values, respectively.

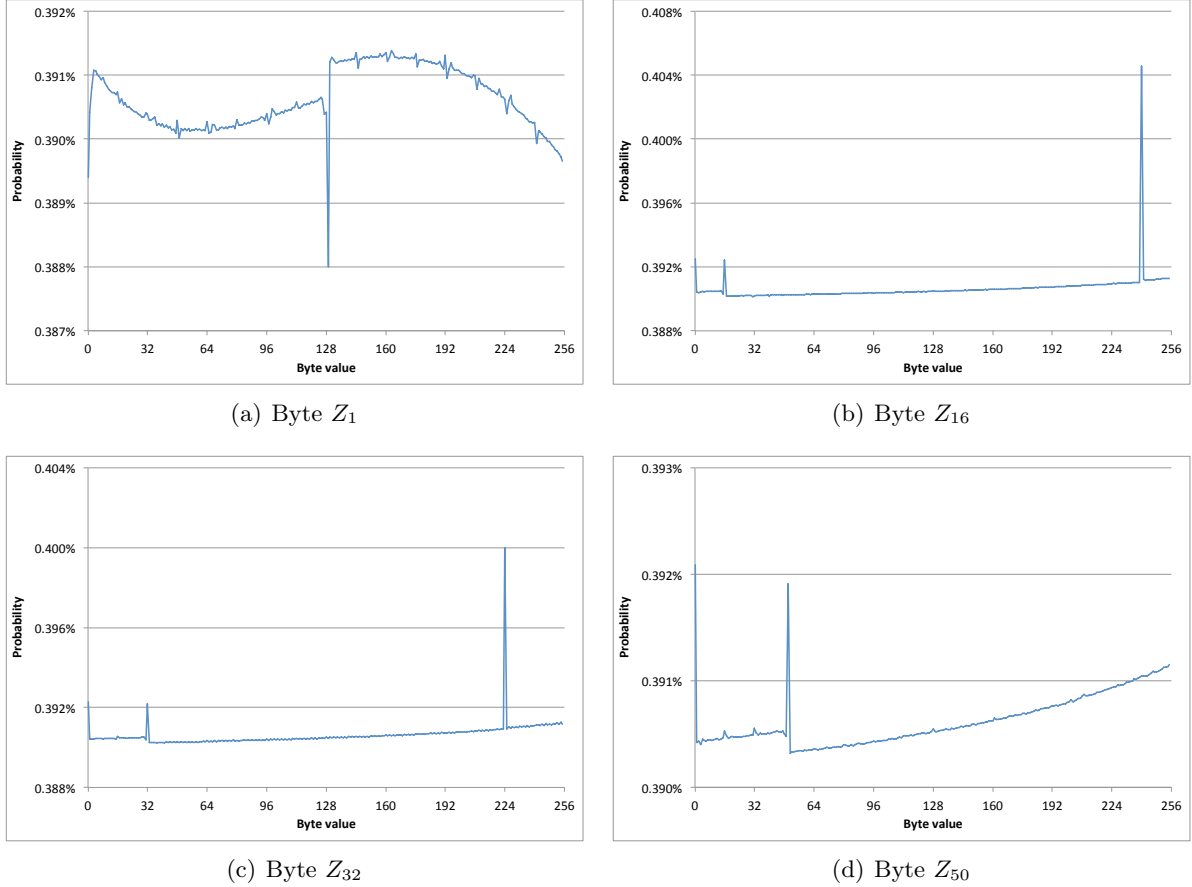


Figure 2: Measured distributions of RC4 keystream bytes  $Z_1$ ,  $Z_{16}$ ,  $Z_{32}$ , and  $Z_{50}$ .

close approximation to the internal state of RC4 after a few invocations of the RC4 keystream generator, [12].

We experimentally verified the Fluhrer-McGrew biases by analysing the output of  $2^{10}$  RC4 instances using 128-bit keys and generating  $2^{40}$  keystream bytes each. For each keystream, the initial 1024 bytes were dropped. Based on this data, we found the biases from [12] to be accurate, also for 128-bit keys. This is in-line with the experiments and observations reported in [12]. Furthermore, we did not identify any additional significant long term biases for consecutive keystream bytes which are repeated with a periodicity that is a proper divisor of 256. Hence, for the purpose of implementing the attack presented in Section 4.2, we assume that the biases identified in [12] are the only existing long term biases for consecutive keystream bytes, and that all other pairs of byte-values are uniformly distributed.

Independently of [12], Mantin [17] identified a positive bias towards the pattern  $ABSAB$ , where  $A$  and  $B$  represent byte values and  $S$  is a short string of bytes (possibly of length 0). The shorter the string  $S$  is, the more significant is the bias. Additionally, Sen Gupta *et al.* [26] identified a bias towards the byte values  $(0, 0)$  for keystream positions  $(Z_r, Z_{r+2})$ , separated by any single keystream byte for  $r \geq 1$ . However, we do not make use of these biases in the attacks presented in this paper.

## 4 Plaintext Recovery Attacks

For the purpose of exposition, we first explain how the broadcast attack by Maitra *et al.* [16] and Sen Gupta *et al.* [26] is meant to work. Suppose byte  $Z_r$  of the RC4 keystream has a dominant bias towards value  $0x00$ . As RC4 encryption is defined as  $C_r = P_r \oplus Z_r$ , the corresponding

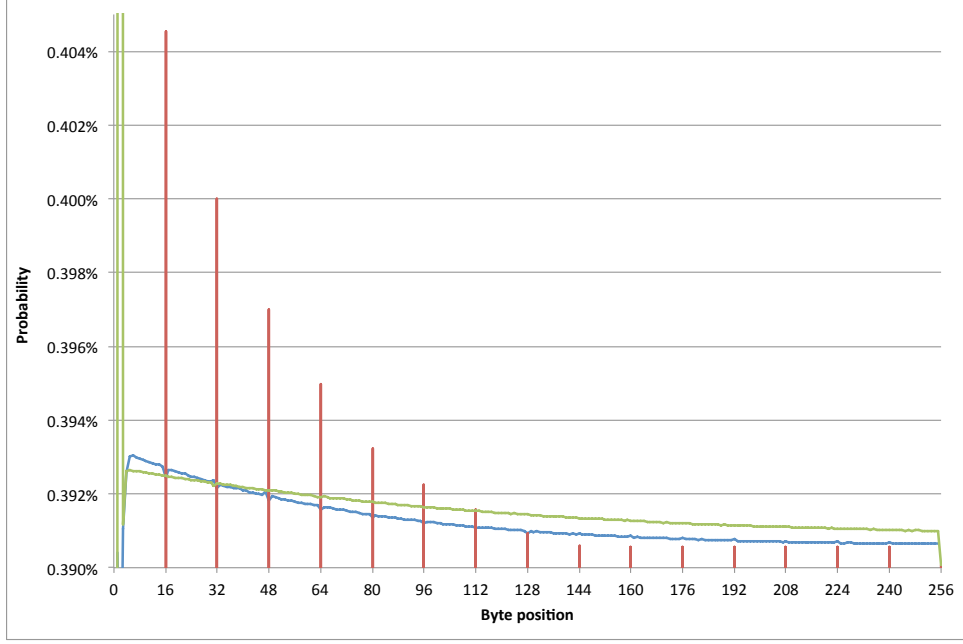


Figure 3: Measured strength of the bias towards 0x00 (green), the bias towards value  $r$  in  $Z_r$  (blue), and the key-length dependent bias towards byte value  $256 - r$  (red) for keystream bytes  $Z_1, \dots, Z_{256}$ , based on keystreams generated by  $2^{44}$  independent random keys. Note that the large peak for the 0x00 bias in  $Z_2$  extends beyond the bounds of the graph and is not fully shown for illustrative purposes.

---

**Algorithm 3:** Basic plaintext recovery attack

---

**input** :  $S$  independent encryptions  $(C_j)_{1 \leq j \leq S}$  of fixed plaintext  $P$ , position  $r$

**output**: estimate  $P_r^*$  for plaintext byte  $P_r$

**begin**

$N_{0x00} \leftarrow 0, \dots, N_{0xFF} \leftarrow 0$

**for**  $j = 1$  **to**  $S$  **do**

$N_{C_{j,r}} \leftarrow N_{C_{j,r}} + 1$

$P_r^* \leftarrow \arg \max_{\mu \in \{0x00, \dots, 0xFF\}} N_\mu$

---

ciphertext byte  $C_r$  has a bias towards plaintext byte  $P_r$ . Thus, obtaining sufficiently many ciphertext samples  $C_r$  for a fixed plaintext  $P_r$  allows inference of  $P_r$  by a majority vote:  $P_r$  is equal to the value of  $C_r$  that occurs most often. This is the core idea of Algorithm 3 that we reproduce from [16, 26]. Let  $S$  denote the number of ciphertexts available to the attacker and, for all  $1 \leq j \leq S$ , let  $C_{j,r}$  denote the  $r$ -th byte of ciphertext  $C_j$ . For a fixed position  $r$ , Algorithm 3 runs through all  $j$ , and in each iteration increments one out of 256 counters, namely the one that corresponds to value  $C_{j,r}$ . After processing all ciphertexts, the character corresponding to the largest counter in the obtained histogram is the output of the algorithm.

The algorithm is tailor-made for plaintext recovery in the case described by Result 2: it assumes that the largest bias in the RC4 keystream is towards 0x00. However, it is highly likely to fail to reliably suggest the correct plaintext byte  $P_r$  if the RC4 keystream has, in position  $r$ , additional biases of approximately the same size (or larger) as the bias towards 0x00. Such additional biases would simply be misinterpreted as the bias towards 0x00 and hence falsify the result. As we observed in Section 3.1 (and Figure 3), several other quite strong biases in the RC4 keystream do indeed exist. This clearly invalidates Algorithm 3 for practical use.

Byte pair	Condition on $i$	Probability
(0, 0)	$i = 1$	$2^{-16}(1 + 2^{-9})$
(0, 0)	$i \neq 1, 255$	$2^{-16}(1 + 2^{-8})$
(0, 1)	$i \neq 0, 1$	$2^{-16}(1 + 2^{-8})$
( $i + 1, 255$ )	$i \neq 254$	$2^{-16}(1 + 2^{-8})$
(255, $i + 1$ )	$i \neq 1, 254$	$2^{-16}(1 + 2^{-8})$
(255, $i + 2$ )	$i \neq 0, 253, 254, 255$	$2^{-16}(1 + 2^{-8})$
(255, 0)	$i = 254$	$2^{-16}(1 + 2^{-8})$
(255, 1)	$i = 255$	$2^{-16}(1 + 2^{-8})$
(255, 2)	$i = 0, 1$	$2^{-16}(1 + 2^{-8})$
(129, 129)	$i = 2$	$2^{-16}(1 + 2^{-8})$
(255, 255)	$i \neq 254$	$2^{-16}(1 - 2^{-8})$
(0, $i + 1$ )	$i \neq 0, 255$	$2^{-16}(1 - 2^{-8})$

Table 1: Fluhrer-McGrew biases for consecutive pairs of byte values. In the table,  $i$  is the internal variable of the RC4 keystream generation algorithm (see Figure 1).

#### 4.1 Our Single-byte Bias Attack

We propose a plaintext-recovery algorithm that takes into account *all* possible single-byte RC4 biases at the same time, along with their strengths. The idea is to first obtain a detailed picture of the distributions of RC4 keystream bytes  $Z_r$ , for all positions  $r$ , by gathering statistics from keystreams generated using a large number of independent keys ( $2^{44}$  in our case). That is, for all  $r$ , we (empirically) estimate

$$p_{r,k} := \Pr(Z_r = k), \quad k = 0x00, \dots, 0xFF,$$

where the probability is taken over the random choice of the RC4 encryption key (i.e., 128-bit keys in the TLS case). Using these biases  $p_{r,k}$ , in a second step, plaintext can be recovered with optimal accuracy using a maximum-likelihood approach, as follows.

Suppose we have  $S$  ciphertexts  $C_1, \dots, C_S$  available for our attack. For any fixed position  $r$  and any candidate plaintext byte  $\mu$  for that position, vector  $(N_{0x00}^{(\mu)}, \dots, N_{0xFF}^{(\mu)})$  with

$$N_k^{(\mu)} = |\{j \mid C_{j,r} = k \oplus \mu\}_{1 \leq j \leq S}| \quad (0x00 \leq k \leq 0xFF)$$

represents the distribution on  $Z_r$  required to obtain the observed ciphertexts  $\{C_{j,r}\}_{1 \leq j \leq S}$  by encrypting  $\mu$ . We compare these *induced* distributions (one for each possible  $\mu$ ) with the accurate distribution  $p_{r,0x00}, \dots, p_{r,0xFF}$  and interpret a close match as an indication for the corresponding plaintext candidate  $\mu$  being the correct one, i.e.,  $P_r = \mu$ . More formally, we observe that the probability  $\lambda_\mu$  that plaintext byte  $\mu$  is encrypted to ciphertext bytes  $\{C_{j,r}\}_{1 \leq j \leq S}$  follows a multinomial distribution and can be precisely calculated as

$$\lambda_\mu = \frac{S!}{N_{0x00}^{(\mu)}! \dots N_{0xFF}^{(\mu)}!} \prod_{k \in \{0x00, \dots, 0xFF\}} p_{r,k}^{N_k^{(\mu)}}. \quad (1)$$

By computing  $\lambda_\mu$  for all  $0x00 \leq \mu \leq 0xFF$  and identifying  $\mu$  such that  $\lambda_\mu$  is largest, we determine the (optimal) maximum-likelihood plaintext byte value. Algorithm 4 specifies the details of the described single-byte bias attack, including the optimizations discussed next.

Observe that, for each fixed position  $r$  and set of ciphertexts  $\{C_{j,r}\}_{1 \leq j \leq S}$ , values  $N_k^{(\mu)}$  can be computed from values  $N_k^{(\mu')}$  by equation  $N_k^{(\mu)} = N_{k \oplus \mu' \oplus \mu}^{(\mu')}$ , for all  $k$ . In other words, vectors  $(N_{0x00}^{(\mu)}, \dots, N_{0xFF}^{(\mu)})$  and  $(N_{0x00}^{(\mu')}, \dots, N_{0xFF}^{(\mu')})$  are permutations of each other; by consequence, term  $S!/(N_{0x00}^{(\mu)}! \dots N_{0xFF}^{(\mu)}!)$  in equation (1) can safely be ignored when determining the largest  $\lambda_\mu$ . Furthermore, computing and comparing  $\log(\lambda_\mu)$  instead of  $\lambda_\mu$  makes the computation slightly more efficient.

---

**Algorithm 4:** Single-byte bias attack

---

**input** :  $\{C_j\}_{1 \leq j \leq S}$  –  $S$  independent encryptions of fixed plaintext  $P$   
           $r$  – byte position  
           $(p_{r,k})_{0x00 \leq k \leq 0xFF}$  – keystream distribution at position  $r$   
**output**:  $P_r^*$  – estimate for plaintext byte  $P_r$   
**begin**  
     $N_{0x00} \leftarrow 0, \dots, N_{0xFF} \leftarrow 0$   
    **for**  $j = 1$  **to**  $S$  **do**  
         $N_{C_{j,r}} \leftarrow N_{C_{j,r}} + 1$   
    **for**  $\mu = 0x00$  **to**  $0xFF$  **do**  
        **for**  $k = 0x00$  **to**  $0xFF$  **do**  
             $N_k^{(\mu)} \leftarrow N_{k \oplus \mu}$   
             $\lambda_\mu \leftarrow \sum_{k=0x00}^{0xFF} N_k^{(\mu)} \log p_{r,k}$   
     $P_r^* \leftarrow \arg \max_{\mu \in \{0x00, \dots, 0xFF\}} \lambda_\mu$   
    **return**  $P_r^*$ 

---

## 4.2 Our Double-byte Bias Attack

As we have seen, Algorithm 4 allows the recovery of the initial 256 bytes of plaintext when multiple encryptions under different keys are observed by the attacker. In the following, we describe an algorithm which allows the recovery of plaintext bytes at *any* position in the plaintext. Furthermore, the algorithm does not require the plaintext to be encrypted under many different keys but works equally well for plaintexts repeatedly encrypted under a *single* key.

Our algorithm is based on biases in the distribution of consecutive bytes  $(Z_r, Z_{r+1})$  of the RC4 keystream that occur as long term biases, i.e., that appear periodically at regular intervals in the keystream. As described in Section 3, we empirically measured the biases which are repeated with a period of 256 bytes. However, in  $2^{50}$  experimentally generated keystream bytes we observed no significant new biases besides those already identified by Fluhrer and McGrew [12]; for the purpose of constructing our algorithm, we hence use the biases described in Table 1 and assume that all other consecutive byte pairs are equally likely to appear in the keystream. In other words, we assume that we have accurate estimates  $p_{r,k_1,k_2}$  for  $1 \leq r \leq 256$ ,  $0x00 \leq k_1, k_2 \leq 0xFF$ , such that

$$\begin{aligned} p_{r,k_1,k_2} &= \Pr[(Z_{256q+r}, Z_{256q+r+1}) = (k_1, k_2)] \\ &\approx \Pr[(\text{RC4}(r, j, \mathcal{S}), \text{RC4}^{+1}(r, j, \mathcal{S})) = (k_1, k_2)] , \end{aligned}$$

where the first probability is taken over all possible RC4 keys and choices of  $q$  from a large subset of  $\mathbb{N}$ , the second probability is taken over all configurations of the internal state  $\mathcal{S}$  and the index  $j$  of the RC4 keystream generation algorithm, and  $\text{RC4}(i, j, \mathcal{S})$  and  $\text{RC4}^{+1}(i, j, \mathcal{S})$  denote the next and second next output byte of RC4's PRGA on input  $st = (i, j, \mathcal{S})$  (see Algorithm 2). By the results of Fluhrer and McGrew, these probabilities are close<sup>9</sup>.

Let  $L$  be an integer multiple of 256. In the following description of our plaintext recovery algorithm, we assume that a fixed  $L$ -byte plaintext  $P = P_1 || \dots || P_L$  is encrypted repeatedly under a single key, i.e., we consider a ciphertext  $C$  obtained by encrypting  $P || \dots || P$ . (In fact, it is sufficient for our attack that the target plaintext bytes form a subsequence of consecutive bytes

---

<sup>9</sup>Note that the internal state  $\mathcal{S}$ , which corresponds to a permutation over byte values, will not be distributed as a random permutation immediately after the key scheduling algorithm is run, even if the used key is picked uniformly at random. Furthermore,  $j$  will not be random, but initialized to 0. However, random  $\mathcal{S}$  and  $j$  will be a close approximation after keystream bytes have been generated a short period of time (see [12] for further discussion of this property).

that are constant across blocks of  $L$  bytes.) Let  $C_j$  denote the substring of  $C$  corresponding to the encryption of the  $j$ -th copy of  $P$ , and let  $C_{j,r}$  denote the  $r$ -th byte of  $C_j$  (i.e.,  $C_{j,r}$  corresponds to byte  $(j-1) \cdot L + r$  of  $C$ ).

Given this setting, it seems reasonable to take an approach towards plaintext recovery similar to that of Algorithm 4: for each position  $r$ , the most likely plaintext pair  $(\mu_r, \mu_{r+1})$  could be computed from the ciphertext bytes  $\{(C_{j,r}, C_{j,r+1})\}_{1 \leq j \leq S}$  and the probability estimates  $\{p_{r,k_1,k_2}\}_{0x00 \leq k_1, k_2 \leq 0xFF}$ . In other words, a plaintext candidate would be obtained by splitting ciphertexts  $C$  into byte pairs and individually computing the most likely corresponding plaintext pairs.

However, by considering overlapping byte pairs, it is possible to construct a more accurate estimate of the likelihood of a plaintext candidate being correct than by just considering the likelihood of individual byte-pairs. More specifically, for any plaintext candidate  $P' = \mu_1 || \dots || \mu_L$  we compute an estimated likelihood  $\lambda_{P'} = \lambda_{\mu_1 || \dots || \mu_L}$  for  $P'$  being correct via the recursion

$$\lambda_{\mu_1 || \dots || \mu_{\ell-1} || \mu_\ell} = \delta_{\mu_\ell | \mu_{\ell-1}} \cdot \lambda_{\mu_1 || \dots || \mu_{\ell-1}} \quad (\ell \leq L), \quad (2)$$

where  $\delta_{\mu_\ell | \mu_{\ell-1}}$  denotes the probability that  $P_\ell = \mu_\ell$  assuming  $P_{\ell-1} = \mu_{\ell-1}$ , and  $\lambda_{\mu_1 || \dots || \mu_{\ell-1}}$  is the estimated likelihood of  $\mu_1 || \dots || \mu_{\ell-1}$  being the correct  $(\ell-1)$ -length prefix of  $P$ . Implicitly, this expression assumes that  $\mu_\ell$  may depend on  $\mu_{\ell-1}$  but not on any earlier plaintext bytes (conditioned on a fixed ciphertext  $C$ ). We show below how values  $\delta_{\mu_\ell | \mu_{\ell-1}}$  can be computed given the ciphertext bytes  $\{(C_{j,\ell-1}, C_{j,\ell})\}_{1 \leq j \leq S}$  and the probability estimates  $\{p_{\ell-1,k_1,k_2}\}_{0x00 \leq k_1, k_2 \leq 0xFF}$ . Note that, by rewriting equation (2) and assuming that  $\lambda_{\mu_1} = \Pr[P_1 = \mu_1]$  is accurately known, we obtain likelihood estimate  $\lambda_{P'} = \Pr[P_1 = \mu_1] \prod_{\ell=2}^L \delta_{\mu_\ell | \mu_{\ell-1}}$ .

Our algorithm computes the plaintext candidate  $P^* = \mu_1 || \dots || \mu_L$  which maximizes the estimated likelihood  $\lambda_{P^*}$ . This is done by exploiting the following easy-to-see optimality-preserving property of  $\lambda_{P^*}$  when defined in this way: for all prefixes  $\mu_1 || \dots || \mu_{\ell-1}$  of  $P^*$ ,  $\ell \leq L$ , we have that  $\lambda_{\mu_1 || \dots || \mu_{\ell-1}}$  takes its largest value amongst all  $(\ell-1)$ -length plaintext candidates having  $\mu_{\ell-1}$  as the last byte.

The basic idea of our algorithm is to iteratively construct  $P^*$  by considering the prefixes of  $P^*$  with increasing length. As just argued, these correspond to the (partial) plaintext candidates with the highest likelihood and a specific choice of the last byte value. However, when computing a candidate for a length  $\ell \leq L$ , it is not known in advance what the specific value of the last byte  $\mu_\ell$  should be. Our algorithm hence computes the most likely partial plaintext candidates for *all* possible values of  $\mu_\ell$ . More specifically, for each  $(\ell-1)$ -length partial candidate  $\mu_1 || \dots || \mu_{\ell-1}$  and any value  $\mu_\ell$ , we compute the likelihood of the  $\ell$ -length plaintext candidate  $\mu_1 || \dots || \mu_{\ell-1} || \mu_\ell$  via equation (2) as  $\lambda_{\mu_1 || \dots || \mu_\ell} = \delta_{\mu_\ell | \mu_{\ell-1}} \cdot \lambda_{\mu_1 || \dots || \mu_{\ell-1}}$ . Due to the optimality-preserving property, the string  $\mu_1 || \dots || \mu_\ell$  with the highest estimated likelihood will correspond to the most likely plaintext candidates of length  $\ell$  with the last byte  $\mu_\ell$ . This guarantees that the  $\ell$ -length prefix of (optimal)  $P^*$  will be among the computed  $\ell$ -length candidates and, furthermore, when the length of  $P^*$  is reached, that  $P^*$  itself will be obtained.

To initialize the above process, the algorithm assumes that the first plaintext byte  $\mu_1$  of  $P$  is known with certainty, i.e.,  $\lambda_{\mu_1} = 1$  (this can, for example, be assumed if the attack is used to recover HTTP cookies from HTTPS traffic, see Section 4.3.1). Likewise, the algorithm assumes that the last byte  $\mu_L$  of  $P$  is known, i.e.,  $\lambda_{\mu_L} = 1$  (this is also the case when recovering HTTP cookies). This leads to a single  $\mu_L$  being used in the last iteration of the above process which will then return the most likely plaintext candidate  $P^*$ . (See Remark 1 for how the algorithm can be modified to work without these assumptions.)

It remains to show the details of how to compute  $\delta_{\mu_{i+1} | \mu_i}$ . This is done similarly to the maximum-likelihood computation of the probability estimate used in Algorithm 4. More precisely, each combination of index  $i$ , pair  $(\mu_i, \mu_{i+1})$ , and ciphertext bytes  $\{(C_{j,i}, C_{j,i+1})\}_{1 \leq j \leq S}$  induces a distribution on the keystream bytes  $\{(Z_{(j-1)L+i}, Z_{(j-1)L+i+1})\}_{1 \leq j \leq S}$ . The latter can

be represented as a vector  $(N_{i,0x00,0x00}, \dots, N_{i,0xFF,0xFF})$ , where

$$N_{i,k_1,k_2} = |\{1 \leq j \leq S \mid (C_{j,i}, C_{j,i+1}) = (k_1 \oplus \mu_i, k_2 \oplus \mu_{i+1})\}|.$$

As in Section 4.1, we see that this vector follows a multinomial distribution, and that the probability that  $(N_{i,0x00,0x00}, \dots, N_{i,0xFF,0xFF})$  will arise (i.e., the probability that  $(\mu_i, \mu_{i+1})$  corresponds to the  $i$ -th and the  $(i+1)$ -th plaintext bytes) is given by

$$\Pr[P_i = \mu_i \wedge P_{i+1} = \mu_{i+1} \mid C] = \frac{S!}{N_{i,0x00,0x00}! \cdots N_{i,0xFF,0xFF}!} \prod_{k_1, k_2 \in \{0x00, \dots, 0xFF\}} p_{i,k_1,k_2}^{N_{i,k_1,k_2}}. \quad (3)$$

We can now compute  $\delta_{\mu_{i+1}|\mu_i}$  as

$$\begin{aligned} \delta_{\mu_{i+1}|\mu_i} &= \Pr[P_{i+1} = \mu_{i+1} \mid P_i = \mu_i \wedge C] \\ &= \frac{\Pr[P_i = \mu_i \wedge P_{i+1} = \mu_{i+1} \mid C]}{\Pr[P_i = \mu_i \mid C]}. \end{aligned} \quad (4)$$

We assume that no significant single-byte biases are present in the keystream, i.e., that  $\Pr[P_i = \mu_i \mid C]$  is uniform over the possible plaintext values  $\mu_i$ . Under this condition, since the term will stay invariant for all plaintext candidates, we can ignore the contribution of factor  $1/\Pr[P_i = \mu_i \mid C]$  in (4), when comparing probability estimates. This is likewise the case for the terms  $S!/(N_{i,0x00,0x00}!, \dots, N_{i,0xFF,0xFF}!)$  in (3), due to similar observations as made for Algorithm 4.

We combine the results of the discussion from the preceding paragraphs, including the proposed optimizations, to obtain our double-byte bias attack in Algorithm 5.

*Remark 1.* The above assumption, that the first and last byte of the plaintext  $P$  is known, can easily be avoided. Specifically, if the first byte is unknown, Algorithm 5 can be initialized by computing, for each possible value  $\mu_2$ , the most likely pairs  $(\mu_1, \mu_2)$ . This can be done based on the ciphertext bytes  $\{(C_{j,1}, C_{j,2})\}_{1 \leq j \leq S}$  and the probability estimates  $\{p_{1,k_1,k_2}\}_{0x00 \leq k_1, k_2 \leq 0xFF}$ . Likewise, if the last byte is unknown, the algorithm will identify  $P^*$  as the plaintext candidate with the highest likelihood estimate among the computed plaintext candidates of length  $L$ . Note, however, that knowing the first and last plaintext byte will lead to a more accurate likelihood estimate and will thereby increase the success rate of the algorithm.

### 4.3 Optimizations

In specific settings where the attacker has *a priori* information about the encrypted plaintext the recovery performance of Algorithms 4 and 5 can be further improved. In the following, we consider such settings and discuss corresponding ways to exploit the additional information.

#### 4.3.1 Partially-known plaintext

In the TLS setting, if the communication protected by TLS is web traffic, all messages will be structured as specified in the HTTP and HTML standards, and will contain parts which are known to or easily guessable by the adversary. Hence, the adversary might know both location and the surrounding plaintext of the information he is trying to recover. For instance, the HTTP cookie header field is guaranteed to start with the characters ‘Cookie:’ and will be terminated with a newline character. Furthermore, knowing a user’s browser and the site the user connects to will most likely reveal to the adversary where the cookie header field is located in the HTTP stream.

An example where this kind of information is exploited by an adversary is Algorithm 5: as described in Section 4.2, knowing the first and last byte surrounding a targeted plaintext allows a more accurate plaintext estimation (cf. Remark 1). Note that partial knowledge of the plaintext will not improve the success rate of Algorithm 4.

---

**Algorithm 5:** Double-byte bias attack

---

**input** :  $C$  – encryption of  $S$  copies of fixed plaintext  $P$   
( $C_{j,r}$  denotes the  $r$ -th byte of the substring of  $C$  encrypting the  $j$ -th copy of  $P$ )  
 $L$  – length of  $P$  in bytes (must be a multiple of 256)  
 $\mu_1$  and  $\mu_L$  – the first and last byte of  $P$   
 $\{p_{r,k_1,k_2}\}_{1 \leq r \leq L-1, 0x00 \leq k_1, k_2 \leq 0xFF}$  – keystream distribution

**output**: estimate  $P^*$  for plaintext  $P$

**notation**: let  $\max_2(Q)$  denote  $(P, \lambda) \in Q$  such that  $\lambda \geq \lambda' \forall (P', \lambda') \in Q$

**begin**

```
 $N_{(r, k_1, k_2)} \leftarrow 0$  for all  $1 \leq r < L, 0x00 \leq k_1, k_2 \leq 0xFF$ 
for  $j = 1$  to  $S$  do
  for  $r = 1$  to  $L - 1$  do
     $N_{(r, C_{j,r}, C_{j,r+1})} \leftarrow N_{(r, C_{j,r}, C_{j,r+1})} + 1$ 
 $Q \leftarrow \{(\mu_1, 0)\}$ 
for  $r = 1$  to  $L - 2$  do
   $Q_{ext} \leftarrow \{\}$  // List of plaintext candidates of length  $r + 1$ 
  for  $\mu_{r+1} = 0x00$  to  $0xFF$  do
     $Q_{\mu_{r+1}} \leftarrow \{\}$  // List of plaintext candidates ending with  $\mu_{r+1}$ 
    for each  $(P', \lambda_{P'}) \in Q$  do
       $P' \rightarrow \mu_1 || \dots || \mu_r$ 
       $\lambda_{P' || \mu_{r+1}} \leftarrow \lambda_{P'} + \sum_{k_1=0x00}^{0xFF} \sum_{k_2=0x00}^{0xFF} N_{(r, k_1 \oplus \mu_r, k_2 \oplus \mu_{r+1})} \cdot \log p(r, k_1, k_2)$ 
       $Q_{\mu_{r+1}} \leftarrow Q_{\mu_{r+1}} \cup \{(P' || \mu_{r+1}, \lambda_{P' || \mu_{r+1}})\}$ 
     $Q_{ext} \leftarrow Q_{ext} \cup \{\max_2(Q_{\mu_{r+1}})\}$ 
   $Q \leftarrow Q_{ext}$ 
 $Q_{\mu_L} \leftarrow \{\}$  // List of plaintext candidates ending with  $\mu_L$ 
for each  $(P', \lambda_{P'}) \in Q$  do
   $P' \rightarrow \mu_1 || \dots || \mu_{L-1}$ 
   $\lambda_{P' || \mu_L} \leftarrow \lambda_{P'} + \sum_{k_1=0x00}^{0xFF} \sum_{k_2=0x00}^{0xFF} N_{(r, k_1 \oplus \mu_{L-1}, k_2 \oplus \mu_L)} \cdot \log p(r, k_1, k_2)$ 
   $Q_{\mu_L} \leftarrow Q_{\mu_L} \cup \{(P' || \mu_L, \lambda_{P' || \mu_L})\}$ 
 $(P^*, \lambda_{P^*}) \leftarrow \max_2(Q_{\mu_L})$ 
return  $P^*$ 
```

---

#### 4.3.2 Restricted character set

In some cases an adversary might know that the encrypted plaintext characters he is trying to recover belong to a small set (at least at some positions). A prime example is encodings of session IDs used by current frameworks for dynamic web pages. For instance, in the popular PHP server-side scripting language, the encoding of such session IDs, which are stored and transmitted in HTTP cookies, can be limited to a representation with 4 bits per character [23]. Other cases of restricted plaintext character sets include base64-encoded transmissions (6 bits per character), or HTML markup.

Algorithms 4 and 5 are readily adjusted to exploit cases where the set of possible plaintext characters is restricted: it suffices to correspondingly restrict the set of candidate plaintexts  $\mu$  considered in the algorithms. Note, however, that the accuracy of the adapted algorithms is expected to depend on the specific character set in use: different sets might correspond to different recovery rates — even if the sets have the same size.

### 4.3.3 Multiple plaintext candidates

Algorithm 4 currently returns the (single) most likely plaintext candidate for each considered position; however, it is easily modified to output the list of plaintext candidates in order of decreasing likelihood  $\lambda_\mu$ , or all plaintext candidates with likelihood above some threshold. (Note that Algorithm 5 does not seem to have this property.) Such extended results might yield a more accurate plaintext recovery if combined with additional knowledge about the encrypted plaintext.

For example, the plaintext might be known to have a certain structure, like English text or HTML markup, which implies that some combinations of characters are more likely to occur than others. This knowledge can assist an adversary in choosing the correct plaintext byte among the candidates with the highest likelihood for a given position. More specifically, by taking into account not only the likely plaintext candidates at each position, but also the probability that the combined string of characters will appear as plaintext, an adversary can obtain a more accurate estimate of the probability that a given set of ciphertexts correspond to a given plaintext. In general, any statistical language model that can estimate the probability of a given sequence of characters appearing as the plaintext can be used in combination with the modified algorithm, and appropriately combining the probability estimates from the language model and the modified algorithm will yield a more accurate plaintext recovery.

In addition, even if the plaintext has no known structure, the likelihood of the plaintext candidates for each position might still provide useful information. For example, if an adversary attempts to recover an HTTP cookie which will give him access to a certain website, a variant of Algorithm 4 will enable him to generate a list of cookie candidates in order of decreasing likelihood, and then successively attempt to connect to the website using the cookie candidates from the list until access is granted.

Furthermore, specific properties of the RC4 keystream can cause a significant increase of the attack’s success probability in the multiple plaintext candidate setting. Indeed, the RC4 keystream has two almost equally strong biases in many of the initial positions, which will cause two different plaintext candidates to have a high likelihood estimate when running Algorithm 4 on a reasonable number of ciphertexts. For these positions, the correct plaintext will be among the two most likely plaintext candidates with high probability, but the probability that it corresponds to the most likely plaintext candidate could be significantly lower.

## 5 Experimental Results

Through simulation, we measured the performance of the single-byte and double-byte bias attacks. We furthermore validated our algorithms in real attack scenarios.

### 5.1 Simulation of Single-byte Bias Attack

We simulated the first plaintext recovery attack described in Section 4. We used RC4 keystreams for  $2^{44}$  random keys to estimate the per-output-byte probabilities  $\{p_{r,k}\}_{1 \leq r \leq 256, 0x00 \leq k \leq 0xFF}$ . We then ran the attack in Algorithm 4 256 times for each of  $S = 2^{24}, 2^{25}, \dots, 2^{32}$  sessions to estimate the attack’s success rate. The results for  $S = 2^{24}, 2^{26}, \dots, 2^{32}$  are shown in Figures 4 and 5(a)–5(d). In each figure, we show the success rate in recovering the correct plaintext byte versus the position  $r$  of the byte in the output stream. Some notable features of these figures are:

- Even with as few as  $2^{24}$  sessions, some positions of the plaintext are correctly recovered with high probability. The ones with highest probability seem to arise because of the key-length-dependent biases that we observed in positions that are multiples of 16. These

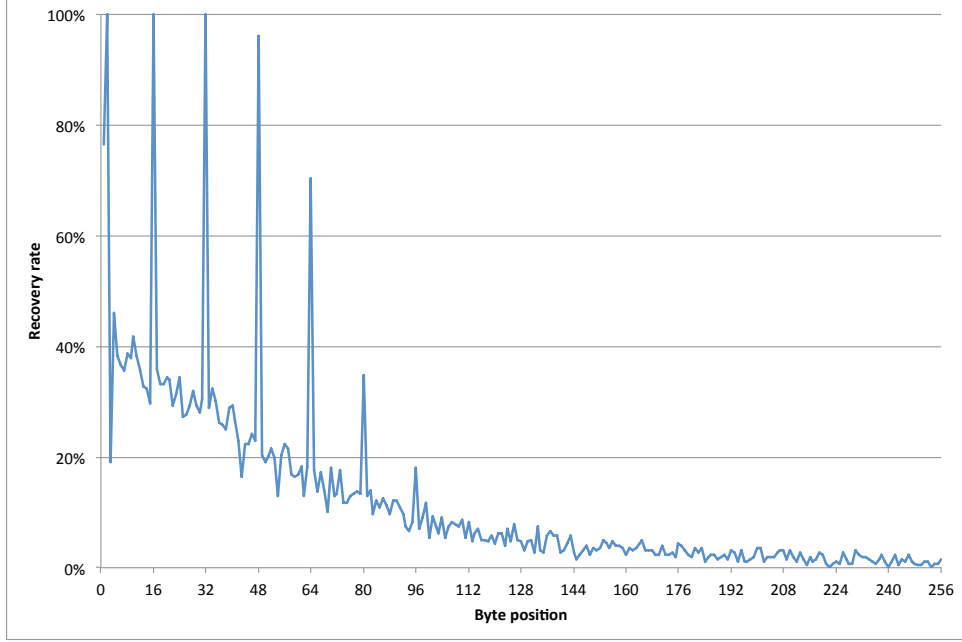


Figure 4: Recovery rate of the single-byte bias attack for  $S = 2^{24}$  sessions for the first 256 bytes of plaintext (based on 256 experiments).

large biases make it easier to recover the correct plaintext bytes when compared to other ciphertext positions.

- With  $S = 2^{26}$  sessions, the first 80 plaintext bytes are recovered with rate at least 50% per byte (but recall that in TLS, the first 36 bytes are not constant across sessions, and so would not be recoverable using our algorithm).
- With  $S = 2^{32}$  sessions, all of the first 256 bytes of output are recovered with rate close to 100%: the rate is at least 96% in all positions, and is 100% for all but 12 positions.
- The rate at which bytes are correctly recovered increases steadily as the number of sessions  $S$  is increased.

Secondly, we executed the recovery attack in a setting where plaintexts are encoded with a 4-bits-per-byte encoding scheme using the characters ‘0’ to ‘9’ and ‘a’ to ‘f’ (see note on restricted character sets in Section 4.3.2). We reused the probability estimates  $\{p_{r,k}\}_{1 \leq r \leq 256, 0x00 \leq k \leq 0xFF}$  for the RC4 keystream bytes generated for the simulation above, and ran a modified version of Algorithm 4 which takes into account the restricted plaintext space. The modified algorithm was run 256 times for each of  $S = 2^{24}, 2^{25}, \dots, 2^{32}$  sessions. The results for  $S = 2^{24}, 2^{26}, 2^{28}$  and  $S = 2^{30}$  are shown in Figures 6(a)–6(d). For comparison, the figures include the success rate of the original attack for an unrestricted plaintext space. We note:

- With  $S = 2^{26}$  sessions, the first 112 plaintext bytes are recovered with rate at least 50% per byte. This represents a marked improvement over the case of an unrestricted plaintext space, where only the first 72 bytes were recovered with rate at least 50% per byte.
- With  $S = 2^{24}, \dots, 2^{30}$  sessions, the recovery attack for the restricted plaintext space has a better success rate than the recovery attack for the unrestricted plaintext space with twice the number of sessions (i.e.  $S = 2^{25}, \dots, 2^{29}$ ) for almost all positions.

Lastly, we simulated the plaintext recovery attack for a version of Algorithm 4 which is modified to return multiple plaintext candidates as described in Section 4.3.3. Specifically, we

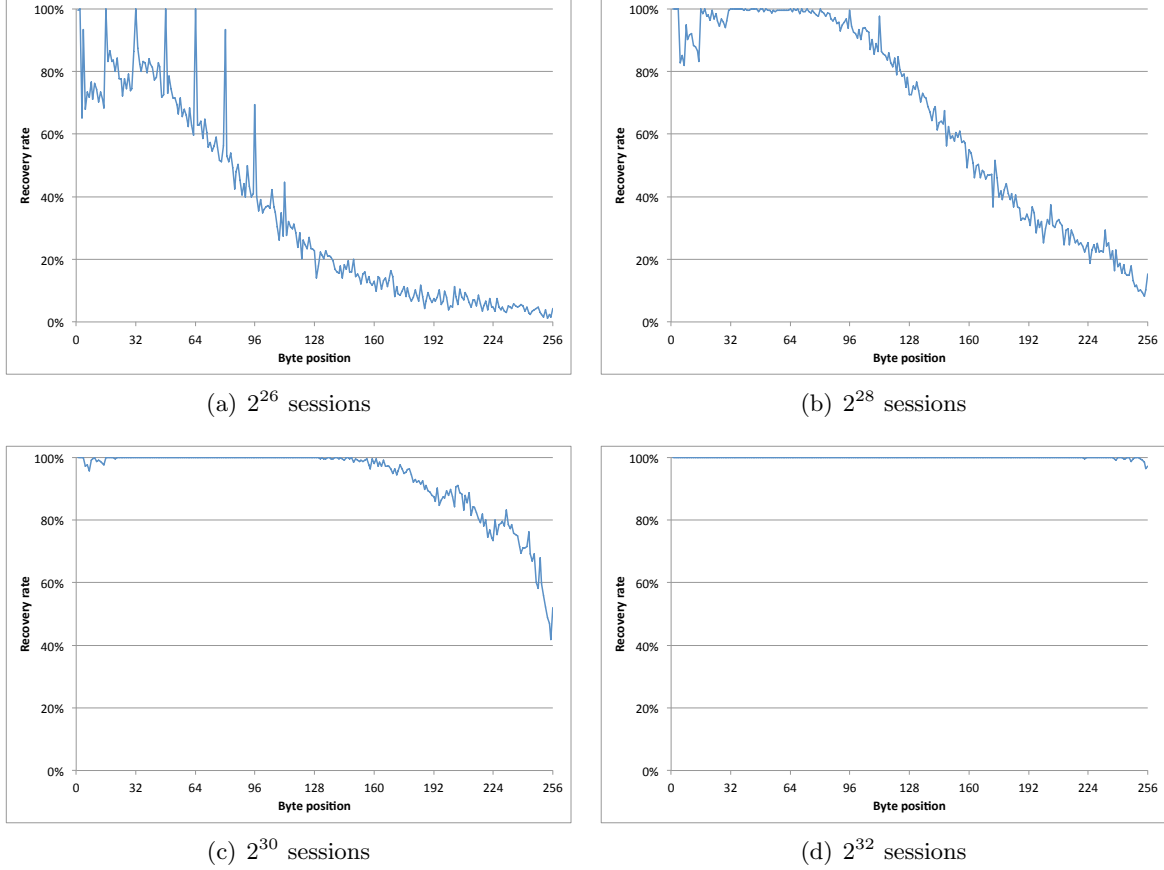


Figure 5: Recovery rate of the single-byte bias attack for  $S = 2^{26}$ ,  $2^{28}$ ,  $2^{30}$  and  $2^{32}$  sessions for the first 256 bytes of plaintext (based on 256 experiments).

modified the algorithm such that instead of returning just a single plaintext candidate byte for each position, the algorithm returns the two most likely candidates (according to the values  $\lambda_\mu$ ). If one of the two plaintext candidate bytes returned by the algorithm corresponded to the correct plaintext byte, we considered the algorithm to be successful for this position (note that for this type of attack, it is assumed that the correct plaintext byte among the two candidates can be determined by other means — see Section 4.3.3 for a discussion of this). The setup was similar to the above: the probability estimates  $\{p_{r,k}\}_{1 \leq r \leq 256, 0x00 \leq k \leq 0xFF}$  for the RC4 keystream bytes were reused from the simulation of the original attack, and the modified algorithm was run 256 times for each of  $S = 2^{24}, 2^{25}, \dots, 2^{32}$  sessions. The results for  $S = 2^{24}, 2^{26}, 2^{28}$ , and  $S = 2^{30}$  are shown in Figures 7(a)–7(d). As above, the success rate of the original plaintext recovery attack is shown in the figures for comparison. We note that:

- With  $S = 2^{26}$  sessions, the first 96 plaintext bytes are recovered with rate at least 50% per byte. (Again, however, recall that the first 36 bytes of plaintext would not be recoverable in our attack for TLS.)
- In general, the modified algorithm has a noticeably better success rate of recovering the plaintext bytes for the initial positions compared to the unmodified algorithm. However, this advantage is not really relevant in attacks against TLS due to the first 36 bytes being encryptions of **Finished** messages.

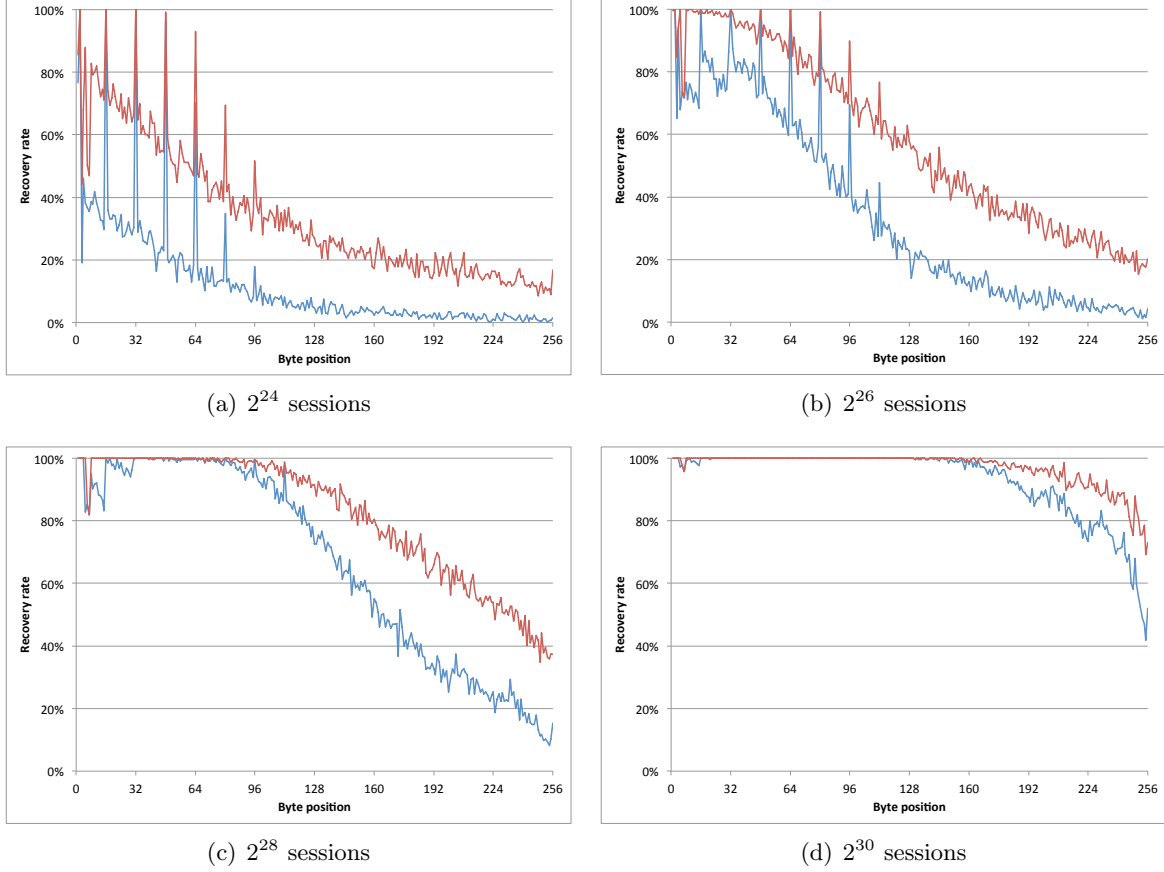


Figure 6: Recovery rates for the restricted plaintext space (red) and the original single-byte bias attack (blue) for  $S = 2^{24}, 2^{26}, 2^{28}$  and  $2^{30}$  sessions (based on 256 experiments).

## 5.2 Simulation of Double-byte Bias Attack

We simulated the second plaintext recovery attack based on Algorithm 5. In the simulation, we encrypted  $S = 1 \cdot 2^{30}, \dots, 13 \cdot 2^{30}$  copies of the same 256-byte plaintext and attempted to recover 16 bytes located at a fixed position in the plaintext. More precisely, we simulated an attack in which we assume the first byte of the plaintext is known, the following 16 bytes are the unknown bytes targeted by the attack, and the byte immediately following these is known. The remaining bytes are assumed not to be of interest in the attack. This attack scenario is very similar to the case in which an adversary attempts to recover a cookie value from an HTTP request (see Section 4.3.1 for a discussion of this). Depending on the number of plaintext copies, we used between one and five 128-bit RC4 keys for the encryption<sup>10</sup>. As highlighted in Section 4.2, we used the biases described by Fluhrer-McGrew [12] to compute the probability estimates  $\{p_{r,k_1,k_2}\}_{1 \leq r \leq 255, 0x00 \leq k_1, k_2 \leq 0xFF}$  required by Algorithm 5.

The attack was run 128 times for each of  $S = 1 \cdot 2^{30}, \dots, 13 \cdot 2^{30}$  encrypted copies of the plaintext to estimate the success rate of the attack. The results are shown in Figure 8(a): the dashed line shows the average fraction of successfully recovered plaintext bytes versus the number of encrypted plaintexts, whereas the solid line shows the success rate of recovering the full 16-byte plaintext versus the number of encrypted plaintexts. We note:

- With  $S = 6 \cdot 2^{30}$  encrypted copies of the plaintext, more than 50% of the plaintext is

<sup>10</sup>Our experiments showed that there is no significant difference in the recovery rate when running the attack on encryptions of the plaintext generated by a single key and encryptions generated by a small number of different keys.

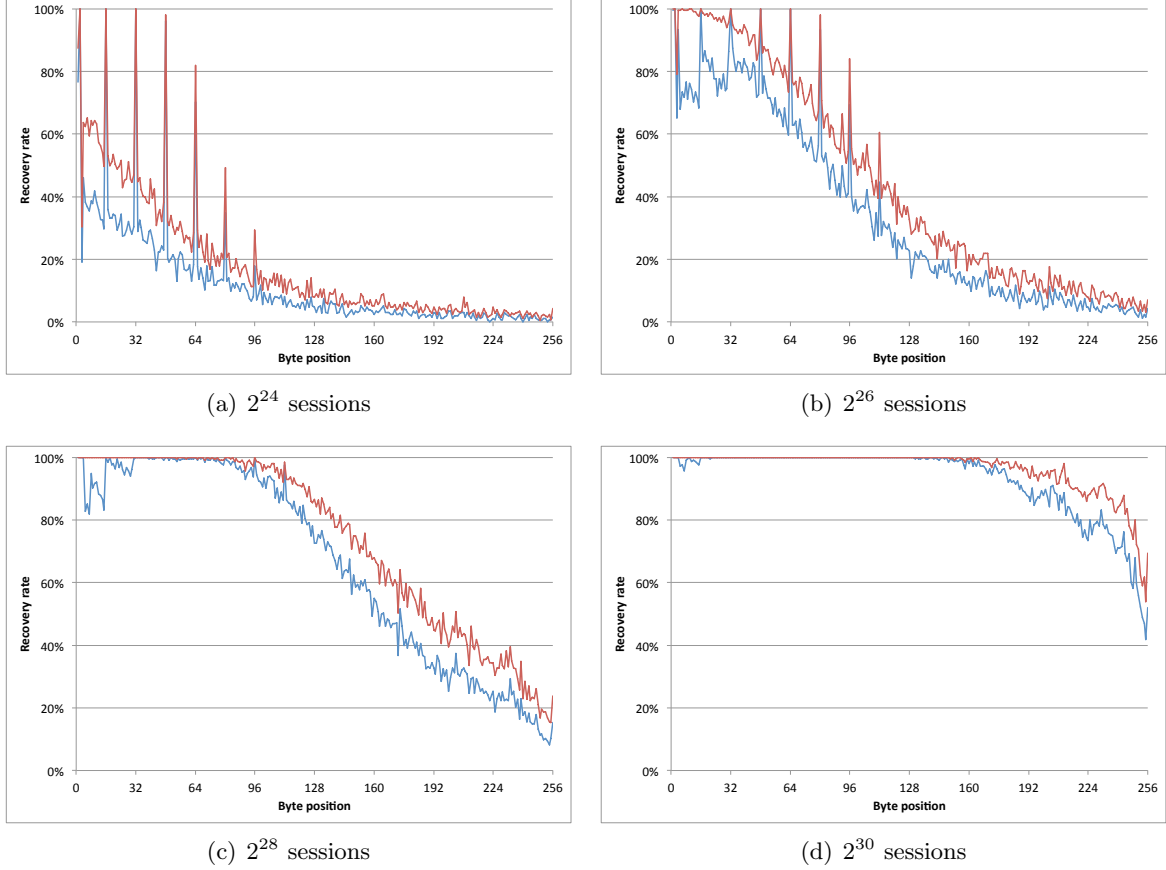


Figure 7: Recovery rates for the modified attack returning two plaintext candidates per position (red) and the original attack (blue) for  $S = 2^{24}$ ,  $2^{26}$ ,  $2^{28}$  and  $2^{30}$  sessions (based on 256 experiments).

correctly recovered on average. Furthermore, in 19% of the 128 trials, the full 16-byte plaintext was recovered.

- With  $S = 8 \cdot 2^{30}$  encrypted copies of the plaintext, the full plaintext is correctly recovered in significantly more than 50% of the 128 trials (more precisely, the full plaintext was recovered in 72% of the trials).
- With  $S = 13 \cdot 2^{30}$  encrypted copies of the plaintext, the full plaintext was recovered in all trials.
- The rate at which the full plaintext is correctly recovered increases fairly rapidly after  $S = 5 \cdot 2^{30}$  copies of the plaintext are encrypted, and with  $S = 11 \cdot 2^{30}$ , the full plaintext is correctly recovered in nearly all trials (99%).

In addition, similarly to Section 5.1, we simulated the attack for plaintexts encoded with a 6-bits-per-byte (base64) and a 4-bits-per-byte encoding scheme. Specifically, we firstly ran a modified version of Algorithm 5 which takes into account the restricted plaintext space by only considering candidate plaintext bytes which correspond to byte-values used in a base64 encoding. Furthermore, we used a plaintext where the 16 bytes targeted by the attack consisted of bytes with a byte-value corresponding to the character ‘b’, which is a valid base64 encoded message. As in the attack above for a non-restricted plaintext space, the probability estimates  $\{p_{r,k_1,k_2}\}_{1 \leq r \leq 255, 0x00 \leq k_1, k_2 \leq 0xFF}$  were based on the biases from [12]. The attack was run 128 times for each of  $S = 1 \cdot 2^{30}, \dots, 12 \cdot 2^{30}$  encrypted copies of the plaintext, and the results are shown in Figure 8(b). We note:

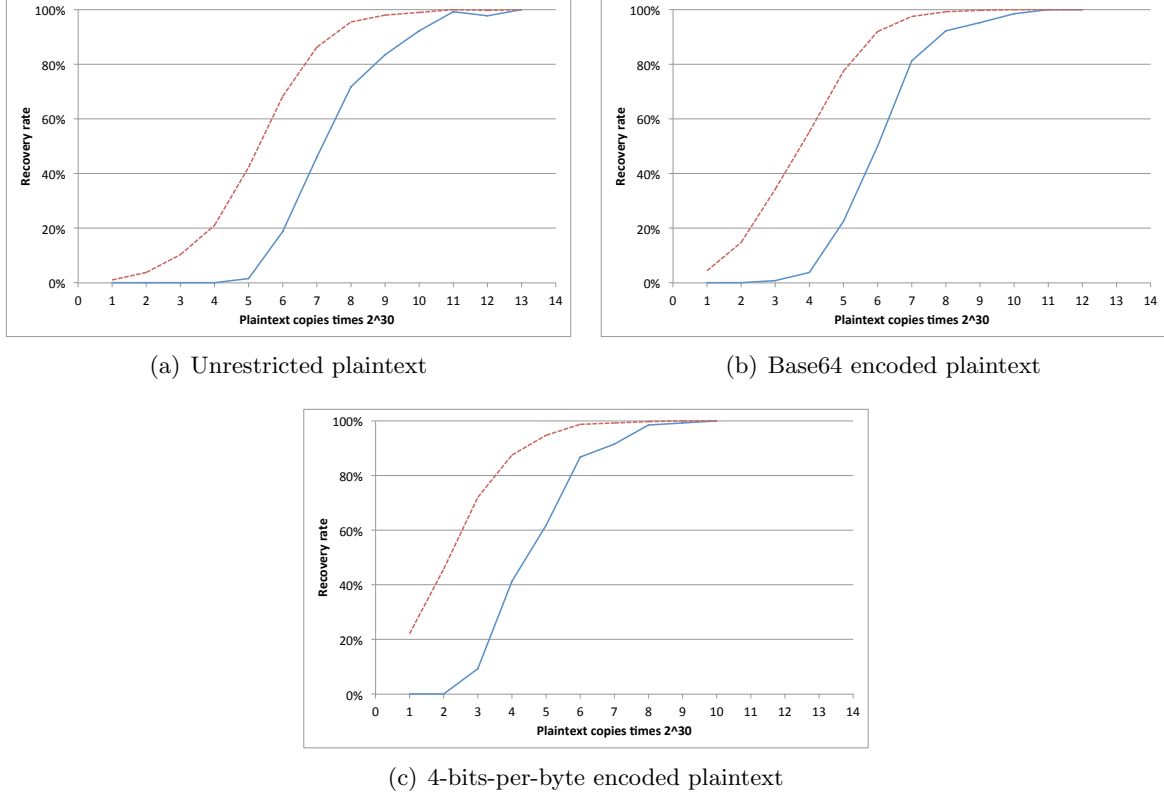


Figure 8: Average fraction of successfully recovered plaintext bytes (red dashed line), and success rate for recovering the full 16-byte plaintext (blue solid line) of the double-byte bias attack based on 128 experiments. Note that the unit of the  $x$ -axis is  $2^{30}$  encrypted copies of the plaintext.

- With  $S = 4 \cdot 2^{30}$  encrypted copies of the plaintext, more than 50% of the plaintext is correctly recovered on average. Furthermore, in 4% of the 128 trials, the full 16-byte plaintext is recovered.
- With  $S = 6 \cdot 2^{30}$  encrypted copies of the plaintext, the full plaintext is correctly recovered in 50% of the 128 trials.
- With  $S = 10 \cdot 2^{30}$  encrypted copies of the plaintext, the full plaintext is correctly recovered in nearly all trials (98%).

Regarding the 4-bit-per-byte encoding scheme, we again assumed a plaintext character set consisting of ‘0’ to ‘9’ and ‘a’ to ‘f’. The setup was similar to the above experiment for base64 encoded messages: we ran a modified version of Algorithm 5 which takes into account the restricted plaintext space, the probability estimates  $\{p_{r,k_1,k_2}\}_{1 \leq r \leq 255, 0x00 \leq k_1, k_2 \leq 0xFF}$  was based on the biases from [12], and we used a plaintext consisting of bytes with a byte-value corresponding to the character ‘b’. The attack was run 128 times for each of  $S = 1 \cdot 2^{30}, \dots, 10 \cdot 2^{30}$  encrypted copies of the plaintext, and the results can be seen in Figure 8(c). We note:

- With  $S = 3 \cdot 2^{30}$  encrypted copies of the plaintext, significantly more than 50% of the plaintext is correctly recovered on average (more precisely, 72% is recovered correctly on average).
- With  $S = 5 \cdot 2^{30}$  encrypted copies of the plaintext, the full plaintext is recovered in more than 50% of the 128 trials.

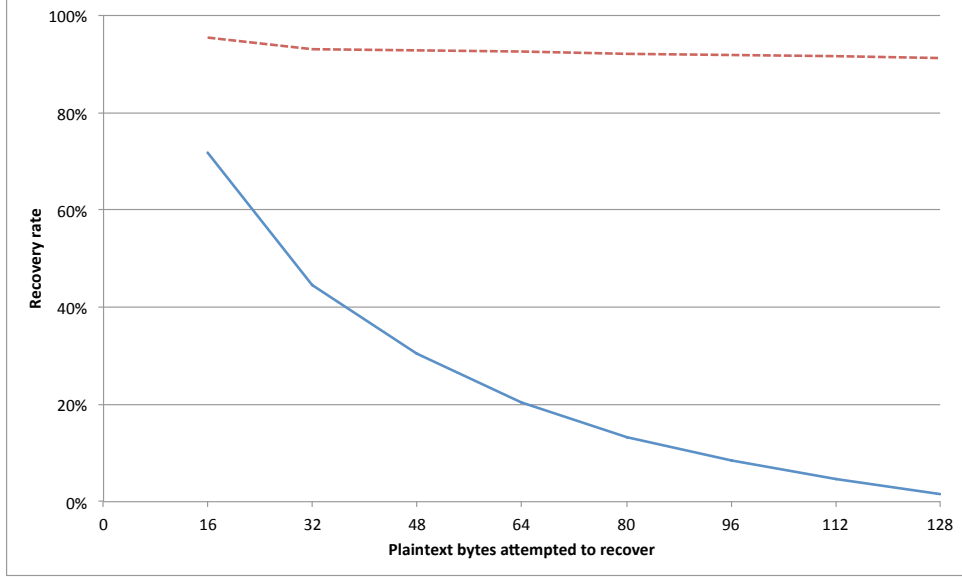


Figure 9: Average fraction of successfully recovered plaintext bytes (dashed line), and success rate for recovering all plaintext bytes the algorithm was setup to recover (solid line) of the double-byte bias attack for  $8 \cdot 2^{30}$  encrypted copies of the plaintext (based on 128 experiments).

- With  $S = 8 \cdot 2^{30}$  encrypted copies of the plaintext, the full plaintext is recovered in nearly all trials (98%).

Lastly, we simulated the unmodified Algorithm 5 for  $8 \cdot 2^{30}$  encrypted copies of a plaintext of length 256 bytes, but attempted to recover an increasing number of plaintext bytes starting with 16 bytes as used in the above experiments. More precisely, we set up the algorithm to recover 16, 32,  $\dots$ , 128 consecutive bytes of the plaintext, and ran the algorithm 128 times for each of these. As in the previous experiments, the probability estimates  $\{p_{r,k_1,k_2}\}_{1 \leq r \leq 255, 0x00 \leq k_1, k_2 \leq 0xFF}$  based on the biases from [12] were used. The result can be seen in Figure 9. As above, the dashed line represents the average fraction of correctly recovered plaintext bytes, whereas the solid line represents the success rate of correctly recovering all plaintext bytes the algorithm was set up to recover. We note that while the former remains almost constant, the latter falls gradually (as expected) as the number of plaintext bytes whose recovery is attempted increases.

### 5.3 Practical Validation

We tested the success rates of our plaintext recovery algorithms in realistic attack settings involving web servers and browsers that are connected through TLS-secured network links. Here, we report on the results.

#### 5.3.1 Validating the operation of RC4 in TLS

We first experimentally verified that the OpenSSL implementation of TLS does indeed use RC4 in the way explained in Section 2.2, in particular without discarding any initial keystream bytes. We did this by setting up an OpenSSL version 1.0.1c client and server running in a virtualised environment, making use of `s_client` and `s_server`, generic tools that are available as part of the OpenSSL distribution package. The two virtual machines were running Ubuntu 12.10 and kernel version 3.5.0-17.

### 5.3.2 Validating the single-byte bias attack

Recall that our single-byte bias attack targets the first 256 bytes of plaintext across multiple TLS sessions or connections with random keys. In order to efficiently generate the large number of ciphertexts needed to test our attack, we again used the `s_client` and `s_server` tools, this time modifying the `s_client` source code to force a session resumption for each TLS packet sent.

Using this approach, we were able to generate around  $2^{21}$  encryptions of a fixed plaintext per hour; with  $2^{25}$  recorded ciphertexts, we obtained results comparable to the simulation of our single-byte bias attack reported in Section 5.1 above. A second possible approach to ensure frequently enough rekeying is to actively interfere with the TLS session after each ciphertext is sent, causing it to fail and be restarted, by injecting a bad TLS packet or by resetting the corresponding TCP connection.

We admit that we do not currently have an automated mechanism for forcing session resumption, e.g., from JavaScript. However, JavaScript running in the browser can trigger the browser to establish a fresh TLS session (with a fresh, random key) after each HTTP connection torn down by the attacker. We estimate that this second approach would be significantly slower than using session resumption because of the additional overhead of running the full TLS Handshake. Thus, even though our double-byte bias attack has higher complexity in terms of its ciphertext requirements than our single-byte bias attack, in practice it could be the more efficient attack in terms of total running time, because it can be executed in a single session (or a small number of sessions).

Furthermore, while the single-byte bias attack successfully recovered fixed plaintext bytes in the initial 256 bytes of the TLS ciphertexts, our subsequent experimentation with modern web browsers revealed that these bytes consisted mostly of less interesting HTTP headers rather than cookies. For this reason, after this basic validation, we switched our experimental focus to the double-byte bias attack.

### 5.3.3 Validating the double-byte bias attack

The double-byte bias attack does not rely on session resumption or session renegotiation and is hence easier to implement in practice. As our experimental setup for this attack, we used a network comprising three (non-virtualized) nodes: a legitimate web server (`www.abc.com`) that serves 16-byte secure cookies over HTTPS, a malicious web server (`www.evil.com`) serving a malicious JavaScript, and a client running a web browser representing a user. The legitimate and malicious web servers run Apache and PHP. For the client, we experimented with various browsers, including Firefox, Opera and Chrome. The nodes were connected through a 100 Mbps Ethernet link; they were equipped with Intel Core i7 processors with 2.3 GHz cores and 16 GB of RAM. None of our experiments used all available CPU resources, nor saturated the network bandwidth.

In this setup, we let the client visit `https://www.abc.com`. This will result in the legitimate web server sending the client a secure cookie which will be stored by the client's browser. This cookie will be the target of the attack. We then let the client visit `http://www.evil.com` and run the malicious JavaScript served by the malicious web server. Note that the same-origin policy (SOP) implemented by the client's browser will prevent the JavaScript from directly accessing the secure cookie. However, the JavaScript will direct repeated HTTP requests to the legitimate server over TLS (i.e. using HTTPS)<sup>11</sup>. The client's browser will then automatically attach the cookie to each request and thereby repeatedly encrypt the target cookie as required in our attack.

---

<sup>11</sup>This is made possible by Cross-Origin Resource Sharing (CORS), a mechanism developed to allow JavaScript to make requests to another domain than the domain the script originates from.

The JavaScript uses `XMLHttpRequest` objects<sup>12</sup> to send the requests. We tested `GET`, `POST`, and `HEAD` requests, but found that `POST` requests gave the best performance (using Firefox). Furthermore, we found that the requests needed to be sent in blocks to ensure that the browser stayed responsive and didn't become overloaded.

For all the browsers we tested (Firefox, Chrome, and Opera), we found that the requests generated by the JavaScript resulted in TLS messages containing more than 256 bytes of ciphertext. To keep the target cookie in a fixed position in the TLS message (modulo 256) as needed for the double-byte bias attack, we therefore added padding by manipulating the HTTP headers in the request to bring the encrypted `POST` requests up to exactly 512 bytes. This padding introduces some overhead to the attack. The exact amount and location of padding needed is browser-dependent, since different browsers behave differently in terms of the content and order of HTTP headers included in `POST` requests. In practice, then, the attacker's JavaScript would need to perform some browser fingerprinting before carrying out its attack.

As an alternative method for generating request to the legitimate web server, we tried replacing the JavaScript code with basic HTML code, using HTML tags such as `img`, pointing to `https://www.abc.com`. The target cookie was still sent in every request, but we found this approach to be less effective (i.e. slower) than using JavaScript.

For Firefox with 512-byte ciphertexts encrypting padded `XMLHttpRequest` `POST` requests, we were able to generate 6 million ciphertexts per hour on our network, with each request containing the target cookie in the same position (modulo 256) in the corresponding plaintext. Given that our attack needs on the order of  $13 \cdot 2^{30}$  encryptions to recover a 16-byte plaintext with high success probability, we estimate that the running time for the whole attack would be on the order of 2000 hours using our experimental setup. The attack generates large volumes of network traffic over long periods of time, and so should not be considered a practical threat. Nevertheless, it demonstrates that our double-byte bias attack does work in principle.

## 6 Application to WPA

The cryptographic mechanisms that aim at protecting transmitted data in modern wireless computer networks have seen an ongoing evolution. Most prominent are the results of the IEEE 802.11 standardization effort; amongst others, IEEE introduced *Wired Equivalent Privacy* (WEP) in 1999, *Wi-Fi Protected Access* (WPA) in 2003, and WPA2 in 2004.

In a nutshell, the WEP protocol [1] works as follows: when a message  $m$  is to be transmitted, a CRC32 checksum  $C(m)$  is appended to it and the resulting string encrypted using RC4; the corresponding packet-specific RC4 key consists of the concatenation of a monotonically increasing sequence number and a shared secret. Practical attacks against integrity, authenticity, and secrecy of transmitted data were reported soon after the publication of WEP, exploiting a wide range of shortcomings of the protocol (including short sequence numbers, lack of randomization, linearity of both RC4 encryption and CRC32, and others) [5]. Refined versions of these attacks [11, 30] rely on advanced cryptanalysis of RC4 and are based on the fact that the (public) packet sequence number is part of the encryption key. Today, WEP is considered fully broken, and its usage is discouraged even in the IEEE 802.11 standard itself.

To counter these attacks, IEEE decided to redesign the cryptographic components of their wireless standards from scratch. Indeed, the WPA2 standard builds on the AES block cipher run in the CCM mode of operation instead of on RC4 [2]. However, as most wireless devices implement their core cryptographic routines directly in silicon, switching from WEP to WPA2 also requires the replacement of hardware in all involved network computers and access points. In order to mitigate these costs, IEEE additionally proposed the *Temporal Key Integrity Protocol* (TKIP) under the name of WPA as an intermediate solution. The design of WPA/TKIP is

---

<sup>12</sup><http://www.w3.org/TR/XMLHttpRequest/>

by intention quite close to that of the original WEP, so that all required modifications can be implemented using only firmware updates. Indeed, WPA/TKIP also encrypts packets using RC4, but with (supposedly) better per-packet keys. In the following we analyze the applicability of our single-byte bias attack on TKIP encryption. Previous security analyses of WPA can be found in [21, 29, 13, 22, 28, 31, 32]<sup>13</sup>.

We reproduce some details of the TKIP key schedule from [2, Sect 11.4.2]. In TKIP it is generally assumed that all communicating parties share two symmetric keys: the *temporal encryption key* TK (128 bits), and a key MK for a MAC (called ‘MIC’ in the context of IEEE 802.11). Additional inputs to the encryption process are the plaintext packet (including sender address, destination address, a priority flag, and the message itself), the TKIP sequence counter TSC (48 bits), and the transmitter address TA (48 bits). Each 16-byte RC4 key K is computed via a so-called ‘key mixing’ procedure, i.e.,  $K \leftarrow \text{KM}(\text{TA}, \text{TK}, \text{TSC})$ , where TSC is different for each processed packet. Internally, KM implements key derivation by mixing together its inputs using a custom 8-round Feistel cipher, where the round function relies on the AES S-boxes (for details see Figures 11-14 and 11-15 in [2]). Key K is derived from the output of this routine, with some structure added to ‘preclude the use of known RC4 weak keys’ [2]. More precisely, we have

$$K_0 = (\text{TSC} \gg 8) \& 0\text{xff} \quad K_1 = ((\text{TSC} \gg 8) | 0\text{x20}) \& 0\text{x7f} \quad K_2 = \text{TSC} \& 0\text{xff} \quad (5)$$

and  $K_3, \dots, K_{15}$  being assigned from the output of the Feistel cipher.

In the context of our analysis, we need to assess the strength of potential biases in the RC4 output streams for the keys K output by KM (observe that these biases might differ from the ones coming from uniform 128-bit keys as specific ‘weak keys’ are intentionally avoided by TKIP). If strong biases exist, then an attack on WPA is likely to be feasible using the algorithms presented in Section 4, without modification. That is, in a setting where the same plaintext message is repeatedly transmitted in a WPA-protected wireless network, one can expect that this plaintext is (at least partially) recoverable. Examples of constant plaintext bytes that are repeatedly transmitted over WPA include IP header fields, ARP packets, and SMB packets.

We experimentally determined single-byte keystream biases in WPA. Precisely, we implemented the KM key derivation function, verified it against the test vectors from [2, Annex M.1.2], and computed keys

$$\text{KM}(\text{TA}, \text{TK}, \text{TSC}), \text{KM}(\text{TA}, \text{TK}, \text{TSC} + 1), \dots, \text{KM}(\text{TA}, \text{TK}, \text{TSC} + (2^{22} - 1))$$

for  $2^{19}$  random assignments of variables TA, TK, TSC, aiming at modelling a realistic application of TKIP. Considering all resulting  $2^{41}$  RC4 keys, we measured the distribution of keystream bytes at positions 1–256. Independently, from the set of all keys K consistent with (5), we picked  $2^{41}$  keys at random (i.e., with random TSC, but without using the Feistel cipher to generate  $K_3, \dots, K_{15}$ ) and identified the corresponding keystream distributions at the same positions. We observed that the difference between these two sets of distributions is negligible, allowing us to make the assumption that the action of the Feistel cipher does not affect the output distribution of RC4. We hence base all of the following observations on the statistics obtained from the  $2^{41}$  random keys conforming with (5).

In contrast to the internal Feistel cipher of KM, the structure on RC4 keys implied by (5) has a significant influence on the biases in the resulting RC4 keystreams.<sup>14</sup> For instance, at positions 17, 33, 49, 65, 81, and 97 (i.e.,  $16k + 1$  for small  $k$ ), new peaks in the distribution show up that do not appear in RC4 with random 128-bit keys. For the distributions at positions 17, 33, and 49 see Figures 10(b), 10(c), and 10(d). Even more extreme is the difference at position 1, shown in Figure 10(a). It is also interesting to observe how the bias towards 0x00 behaves in

<sup>13</sup>The Wi-Fi Alliance acknowledges some of the vulnerabilities reported in these papers at <http://www.wi-fi.org/knowledge-center/faq/please-explain-various-security-standards-algorithms>.

<sup>14</sup>These biases can be seen pictorially at [http://www.isg.rhul.ac.uk/tls/kip\\_biases.pdf](http://www.isg.rhul.ac.uk/tls/kip_biases.pdf).

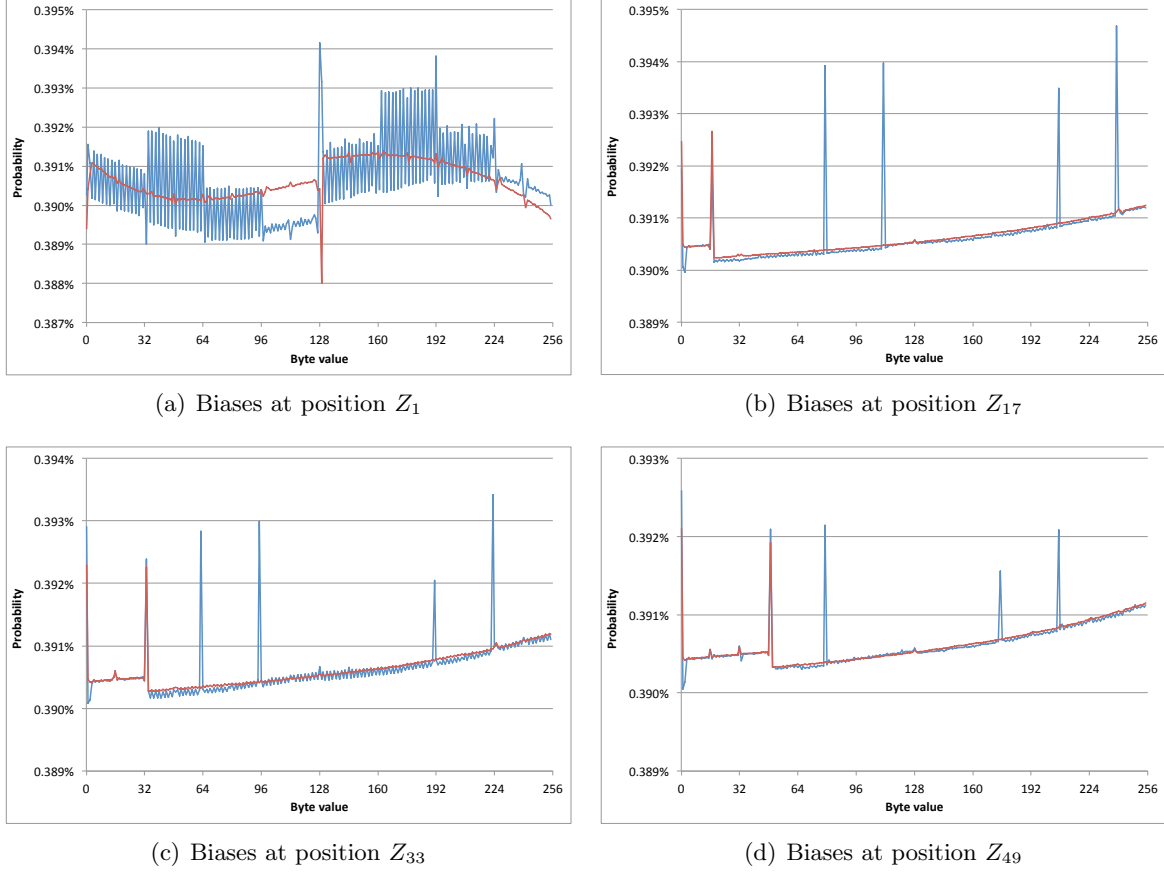


Figure 10: Measured distribution of the RC4/TKIP keystream at positions  $Z_1$ ,  $Z_{17}$ ,  $Z_{33}$ , and  $Z_{49}$  (blue). These estimates were obtained by considering  $2^{41}$  KM-generated keys. For reference, we overlay the biases of the RC4 keystream with uniform 128-bit keys (red).

the TKIP case: it is persistently less pronounced than in the case of random 128-bit keys at positions 2–32 and 128–160, whereas for the other positions its strength alternates from byte to byte, between being significantly stronger and being significantly weaker than the corresponding bias for uniform keys. This is illustrated in Figure 11, which compares the strength of the bias towards 0x00 for TKIP and for random 128-bit keys.

Akin to our approach from Section 5.1, we simulated the single-byte bias plaintext recovery attack from Section 4.1 on WPA to determine its performance. We first used reference RC4 keystreams from close to  $2^{41}$  keys (generated using KM) to estimate the per-output-byte probabilities  $\{p_{r,k}\}_{1 \leq r \leq 256, 0x00 \leq k \leq 0xFF}$ . We then ran the attack in Algorithm 4 256 times for each of  $S = 2^{24}, 2^{26}, 2^{28}, 2^{30}$  simulated frames to estimate the attack’s success rate. The results are shown in Figures 12(a)–12(d), which display the success rate of recovering the correct plaintext byte versus the byte position  $r$  in the keystream. Some notable features of these figures are:

- With  $2^{26}$  frames, the first 55 plaintext bytes are recovered with rate at least 50% per byte. While the corresponding result for random 128-bit keys are higher, comparing Figure 5(a) and Figure 12(b) reveals that many plaintext bytes are recovered with a significantly higher rate in the TKIP case, leading to a higher average recovery rate.
- With  $S = 2^{30}$  frames, the first 130 plaintext bytes are recovered with rate close to 100%; the first 211 bytes are recovered with rate at least 50%. Note again that, while the corresponding results for random 128-bit keys are better in this respect, many plaintext bytes are recovered with significantly higher probability in the TKIP case; see Figures 5(c) and 12(d).

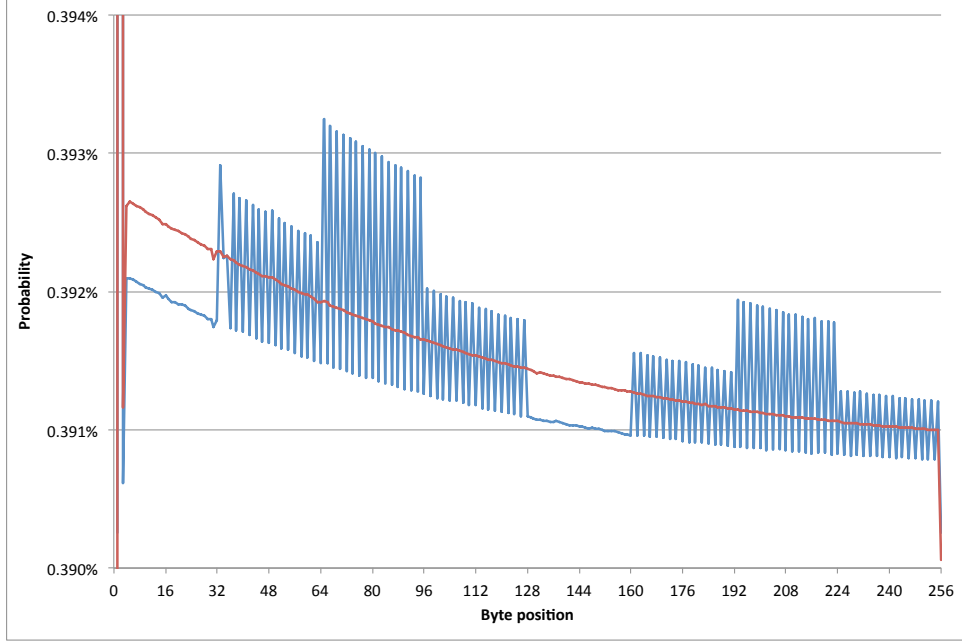


Figure 11: Strength of the bias towards 0x00 of RC4/TKIP keystream bytes at positions 1–256 (blue). The red line corresponds to the biases of RC4 with uniform 128-bit keys and is reproduced from Figure 3 for comparison.

- Independently of the number  $S$  of considered frames, the recovery rate is highly correlated with the strength of the bias towards 0x00 at the same position: to see this, compare Figures 12(a)–12(d) with Figure 11.

By comparing Figure 12 with Figure 5, we observe that the structure on keys enforced by (5), which was aiming to ‘preclude the use of known RC4 weak keys’ [2] (to prevent WEP key recovery), effectively allows easier recovery of plaintext bytes than with uniform keys, at least in some positions.

We do not analyze whether WEP is also amenable to broadcast attacks, simply for the reason that the attacks on WEP mentioned above [11, 30] are already highly powerful and practical. Moreover, WEP is deprecated and is gradually becoming less common.

## 7 Discussion and Conclusions

We have shown that plaintext recovery for RC4 in TLS is possible for the first 200 or so bytes of the plaintext stream (after the `Finished` message), provided sufficiently many independent encryptions of the same plaintext are available. The number of encryptions required (around  $2^{28}$  to  $2^{32}$  for reliable recovery) is large, but not completely infeasible. We have also shown that plaintext recovery for RC4 is possible from arbitrary positions in the plaintext, given enough encryptions of the same plaintext bytes. Here, the number of encryptions required is rather higher (around  $13 \cdot 2^{30}$ ), but the attack is more flexible and more efficient in practice because it avoids rerunning the TLS Handshake. We have explained how the attacks also apply directly to WPA. Certainly, the security level provided by RC4 is *far* below the strength implied by the 128-bit key in both TLS and WPA.

This said, it would be incorrect to describe the attacks as being a practical threat to TLS and WPA today. However, our attacks are open to further enhancement, using, for example, the ability of our algorithms to output likelihoods for candidate plaintext bytes coupled with more sophisticated plaintext models. It may also be possible to enhance the rate of ciphertext

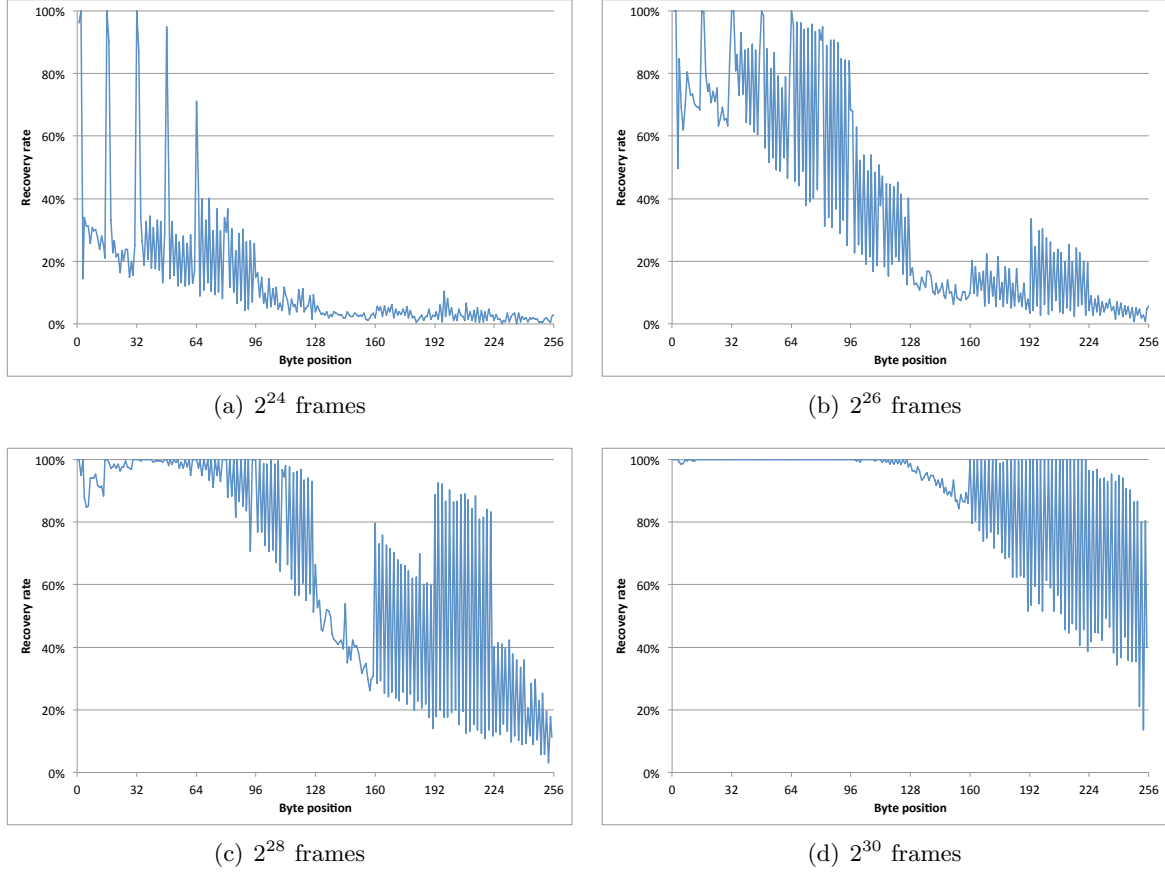


Figure 12: TKIP success rate for  $2^{24}$ ,  $2^{26}$ ,  $2^{28}$  and  $2^{30}$  frames (4, 16, 64, and 256 keys generating  $2^{22}$  frames each)

generation in browsers using methods beyond our knowledge. It would seem dangerous to assume that the attacks will not be improved by other researchers in future.

There are countermeasures to the attacks. We discussed these countermeasures extensively with vendors during the disclosure process that we followed prior to making our attacks public. They include:

1. discarding the initial keystream bytes output by RC4, as recommended in [20];
2. fragmenting the initial HTTP requests at the browser so that the initial keystream bytes are mostly (or entirely) used to encrypt TLS's MAC fields;
3. adding random padding to HTTP requests; and
4. limiting the lifetime of cookies or the number of times they can be sent from the browser.

The first countermeasure cannot easily be implemented in TLS (or WPA) because it would require mass coordination between the many different implementations. The first two countermeasures are not effective against our double-byte bias attack; the second countermeasure applies only for HTTPS. The third countermeasure (which applies only for HTTPS) can be relatively easily implemented in browsers but increases the complexity of our attacks rather than defeating them completely. The fourth countermeasure (which applies only for HTTPS) is currently effective, but not immune to further improvements of our attacks.

Some vendors (e.g. Opera<sup>15</sup>) have implemented a combination of these (and other) countermeasures; others (e.g. Google in Chrome) are focussing on implementing TLS 1.2 and AES-

<sup>15</sup><http://my.opera.com/securitygroup/blog/2013/03/20/on-the-preciousness-of-rc4>

GCM; yet others (e.g. Microsoft in Windows 8.1 Preview<sup>16</sup>) have modified their code so that RC4 is no longer enabled by default for TLS.

We recognise that, with around 50% of TLS traffic currently using RC4, recommending that it be avoided completely in TLS is not a suggestion to be made lightly. Nevertheless, given the rather small security margin provided by RC4 against our attacks, our recommendation is that RC4 should henceforth be avoided in TLS, and deprecated as soon as possible. We also make the same recommendation for WPA; fortunately, WPA seems destined to become a legacy protocol in the future.

## Acknowledgements

We thank David McGrew for raising the question of the security of RC4 in TLS.

## References

- [1] Wireless LAN medium access control (MAC) and physical layer (PHY) specification, 1997.
- [2] Wireless LAN medium access control (MAC) and physical layer (PHY) specification: Amendment 6: Medium access control (MAC) security enhancements, 2004.
- [3] N. AlFardan and K. G. Paterson. Lucky 13: Breaking the TLS and DTLS record protocols. In *IEEE Symposium on Security and Privacy*, 2013. URL <http://www.isg.rhul.ac.uk/tls/Lucky13.html>.
- [4] B. Amann. Personal communication, February 2013.
- [5] N. Borisov, I. Goldberg, and D. Wagner. Intercepting mobile communications: The insecurity of 802.11. In C. Rose, editor, *MOBICOM*, pages 180–189. ACM, 2001.
- [6] B. Canvel, A. Hiltgen, S. Vaudenay, and M. Vuagnoux. Password interception in a SSL/TLS channel. *Advances in Cryptology-CRYPTO 2003*, pages 583–599, 2003.
- [7] T. Dierks and C. Allen. The TLS Protocol Version 1.0. RFC 2246, Internet Engineering Task Force, Jan. 1999. URL <http://www.rfc-editor.org/rfc/rfc2246.txt>.
- [8] T. Dierks and E. Rescorla. The Transport Layer Security (TLS) Protocol Version 1.1. RFC 4346, Internet Engineering Task Force, Apr. 2006. URL <http://www.rfc-editor.org/rfc/rfc4346.txt>.
- [9] T. Dierks and E. Rescorla. The Transport Layer Security (TLS) Protocol Version 1.2. RFC 5246, Internet Engineering Task Force, Aug. 2008. URL <http://www.rfc-editor.org/rfc/rfc5246.txt>.
- [10] T. Duong and J. Rizzo. Here come the  $\oplus$  Ninjas. Unpublished manuscript, 2011.
- [11] S. R. Fluhrer, I. Mantin, and A. Shamir. Weaknesses in the key scheduling algorithm of RC4. In S. Vaudenay and A. M. Youssef, editors, *Selected Areas in Cryptography*, volume 2259 of *Lecture Notes in Computer Science*, pages 1–24. Springer, 2001.
- [12] S. R. Fluhrer and D. McGrew. Statistical analysis of the alleged RC4 keystream generator. In B. Schneier, editor, *FSE*, volume 1978 of *Lecture Notes in Computer Science*, pages 19–30. Springer, 2000.

---

<sup>16</sup><http://technet.microsoft.com/en-us/library/dn303404.aspx>

- [13] F. M. Halvorsen, O. Haugen, M. Eian, and S. F. Mjølsnes. An improved attack on TKIP. In A. Jøsang, T. Maseng, and S. J. Knapskog, editors, *NordSec*, volume 5838 of *Lecture Notes in Computer Science*, pages 120–132. Springer, 2009. ISBN 978-3-642-04765-7.
- [14] T. Isobe, T. Ohigashi, Y. Watanabe, and M. Morii. Full plaintext recovery attack on broadcast RC4. In *Preproceedings of FSE*, 2013.
- [15] K. Jaganathan, L. Zhu, and J. Brezak. The RC4-HMAC Kerberos Encryption Types Used by Microsoft Windows. RFC 4757 (Informational), Dec. 2006. URL <http://www.ietf.org/rfc/rfc4757.txt>.
- [16] S. Maitra, G. Paul, and S. Sengupta. Attack on broadcast RC4 revisited. In A. Joux, editor, *FSE*, volume 6733 of *Lecture Notes in Computer Science*, pages 199–217. Springer, 2011.
- [17] I. Mantin. Predicting and distinguishing attacks on rc4 keystream generator. In R. Cramer, editor, *EUROCRYPT*, volume 3494 of *Lecture Notes in Computer Science*, pages 491–506. Springer, 2005.
- [18] I. Mantin and A. Shamir. A practical attack on broadcast RC4. In M. Matsui, editor, *FSE*, volume 2355 of *Lecture Notes in Computer Science*, pages 152–164. Springer, 2001.
- [19] D. McGrew and D. Bailey. AES-CCM Cipher Suites for Transport Layer Security (TLS). RFC 6655 (Proposed Standard), 2012.
- [20] I. Mironov. (Not so) random shuffles of RC4. In M. Yung, editor, *CRYPTO*, volume 2442 of *Lecture Notes in Computer Science*, pages 304–319. Springer, 2002.
- [21] V. Moen, H. Raddum, and K. J. Hole. Weaknesses in the temporal key hash of WPA. *Mobile Computing and Communications Review*, 8(2):76–83, 2004.
- [22] M. Morii and Y. Todo. Cryptanalysis for RC4 and breaking WEP/WPA-TKIP. *IEICE Transactions*, 94-D(11):2087–2094, 2011.
- [23] PHP Documentation Group. PHP manual, Feb 2013. <http://www.php.net/manual/en/session.configuration.php#ini.session.hash-bits-per-character>.
- [24] J. Salowey, A. Choudhury, and D. McGrew. AES Galois Counter Mode (GCM) Cipher Suites for TLS. RFC 5288 (Proposed Standard), Aug. 2008. URL <http://www.ietf.org/rfc/rfc5288.txt>.
- [25] S. Sen Gupta, S. Maitra, G. Paul, and S. Sarkar. Proof of empirical RC4 biases and new key correlations. In *Selected Areas in Cryptography*, pages 151–168, 2011.
- [26] S. Sen Gupta, S. Maitra, G. Paul, and S. Sarkar. (Non-) random sequences from (non-) random permutations – analysis of RC4 stream cipher. *Journal of Cryptology*, pages 1–42, 2012.
- [27] P. Sepehrdad, S. Vaudenay, and M. Vuagnoux. Discovery and exploitation of new biases in RC4. In A. Biryukov, G. Gong, and D. R. Stinson, editors, *Selected Areas in Cryptography*, volume 6544 of *Lecture Notes in Computer Science*, pages 74–91. Springer, 2010.
- [28] P. Sepehrdad, S. Vaudenay, and M. Vuagnoux. Statistical attack on RC4 – distinguishing WPA. In K. G. Paterson, editor, *EUROCRYPT*, volume 6632 of *Lecture Notes in Computer Science*, pages 343–363. Springer, 2011.
- [29] E. Tews and M. Beck. Practical attacks against WEP and WPA. In D. A. Basin, S. Capkun, and W. Lee, editors, *WISEC*, pages 79–86. ACM, 2009. ISBN 978-1-60558-460-7.

- [30] E. Tews, R.-P. Weinmann, and A. Pyshkin. Breaking 104 bit wep in less than 60 seconds. In S. Kim, M. Yung, and H.-W. Lee, editors, *WISA*, volume 4867 of *Lecture Notes in Computer Science*, pages 188–202. Springer, 2007.
- [31] Y. Todo, Y. Ozawa, T. Ohigashi, and M. Morii. Falsification attacks against WPA-TKIP in a realistic environment. *IEICE Transactions*, 95-D(2):588–595, 2012.
- [32] M. Vanhoef and F. Piessens. Practical verification of WPA-TKIP vulnerabilities. In K. Chen, Q. Xie, W. Qiu, N. Li, and W.-G. Tzeng, editors, *ASIACCS*, pages 427–436. ACM, 2013. ISBN 978-1-4503-1767-2.
- [33] S. Vaudenay and M. Vuagnoux. Passive-only key recovery attacks on RC4. In C. M. Adams, A. Miri, and M. J. Wiener, editors, *Selected Areas in Cryptography*, volume 4876 of *Lecture Notes in Computer Science*, pages 344–359. Springer, 2007.