A method for constructing sliding windows leak from noisy cache timing information

Rei Ueno¹ · Junko Takahashi² · Yu-ichi Hayashi³ · Naofumi Homma¹

Received: 14 November 2019 / Accepted: 2 May 2020 / Published online: 12 June 2020
© The Author(s) 2020

Abstract

This paper presents a method for constructing an operation sequence of sliding window exponentiation from the noisy cache information of RSA, which can be used for a cache attack using sliding windows leak (SWL). An SWL attack can retrieve the secret keys of RSA with non-negligible probability if the SWL is correctly captured. However, in practice, it is not always possible for an attacker to acquire a complete and correct operation sequence from cache information observation. In this paper, we first show that the capture errors in an operation sequence can be evaluated based on the Levenshtein distance between correct and estimated sequences. The dynamic time warping algorithm is employed for quantitative evaluation. Then, we present a method of accurately estimating a complete and correct operation sequence from noisy sequences obtained through multiple observations. Furthermore, we show the effectiveness of the proposed method through a set of experiments performed using RSA software in Libgcrypt. As a result, we can identify the correct operation sequence from approximately 100 observations of cache traces.

Keywords Cache attack · Sliding windows exponentiation · FLUSH + RELOAD · Libgcrypt · RSA

1 Introduction

In recent years, a new type of services that lease remote (high-end) servers to clients is being widely used, such as Amazon AWS and Google Cloud Platform. The deployment of such services makes various users share an identical computational resource (e.g., CPU and memory) as virtual machines (VMs). Consequently, cache attacks in such shared servers are attracting attention. A cache attack is a kind of side-channel attack that exploits the time differences of

cache access information to estimate secret information [1–7]. Cache attacks are possible if victim and attacker processes share a cache. These attacks can potentially occur in common cloud services where several VMs on a hypervisor share a cache. If the cryptographic implementation on VMs leaks its secret information to the attacker and the attacker can retrieve secret key via a practical time/space complexity, the attacker can, for example, eavesdrop the communication and perform a forgery in some practical scenarios even when a cryptographic protocol such as TLS is used. As cloud services are increasingly placing the VMs of multiple users in a server, the security evaluation of cryptographic software against cache attacks is highly required.

In CHES 2017, Bernstein et al. reported a cache attack on RSA software using (left-to-right) sliding window exponentiation [8]. Sliding window is one of the fastest methods for modular exponentiation [9], and it is widely employed in some open-source software (OSS) in cryptography because of its performance. Cache attacks on the sliding window method are attracting considerable attention owing to their wide applicability to OSS such as GnuPG. While sliding window is evidently a non-constant-time algorithm, it was expected to be somewhat resistant to side-channel attacks including cache attacks and simple power analysis (SPA)

This research has been supported by JSPS KAKENHI Grants No. 17H00729, No. 19H21526, and 20K19765.

✉ Rei Ueno
ueno@riec.tohoku.ac.jp
Naofumi Homma
homma@riec.tohoku.ac.jp

¹ Tohoku University, 2-1-1 Katahira, Aoba-ku, Sendai-shi 980-8577, Japan

² NTT Secure Platform Laboratories, Nippon Telegraph and Telephone Corporation, 3-9-11 Midori-cho, Musashino-shi, Tokyo 180-8535, Japan

³ Nara Institute of Science and Technology, 8916-5 Takayama-cyo, Ikoma-shi, Nara 630-0192, Japan

[10]. This is because the sequence of operations (i.e., multiplication and squaring) is less informative compared to other non-constant-time algorithms such as the left-to-right binary method. However, a cache attack utilizes the sliding windows leak (SWL) observed through cache timing information during the modular exponentiation of CRT-RSA decryption. In [8], it was shown that a cache attack can retrieve the secret key of RSA software using sliding window with non-negligible probabilities if the operation sequence of sliding window is correctly captured. A cache attack has been demonstrated through an application to CRT-RSA software in Libgcrypt, which is one of the most common cryptographic OSS [11].

However, the feasibility of Bernsteins' attack is still unclear in a practical setting because SWL (i.e., the correct operation sequence) is not always correctly captured from cache timing information. In [8], Bernstein et al. mentioned that the correct SWL should be constructed from such noisy cache timing information obtained through multiple observations. In addition, the noise (i.e., incorrectly captured part due to miscellaneous reasons on measurement and computation environment) contains capture error, which indicates the undetection and misdetection of multiplication and squaring. Undetection is the overlooking of an operation, and misdetection is the observation of an operation that is not actually performed. Therefore, the length of the captured operation sequence is different from that of the actual sequence. As we cannot determine the position of misdetection/undetection, it is quite difficult to construct the correct SWL from a set of a noisy operation sequence obtained by observing the cache timing information multiple times. Although the authors of [8] mentioned that aligning the traces and a simple majority rule is sufficient to recover the correct sequence, it is unclear how to align such operation sequences/cache traces containing misdetection/undetection of symbols to perform the majority rule with a practical complexity.

We present a method for constructing the correct SWL from noisy cache timing information (i.e., noisy operation sequence of sliding window) to evaluate feasibility of an SWL-based attack. The basic idea of the proposed method is to separate the noisy sequence of multiplication and squaring into "operation patterns." An operation pattern is defined as S^tM , which indicates that a multiplication is performed after t times of squaring. By transforming a noisy operation sequence to a pattern sequence, we can evaluate the capture errors in the operation sequence in a quantitative manner based on the dynamic time warping (DTW) algorithm [12]. Since we experimentally confirm that the most of capture errors come from the undetections of squaring operation, the transformation from an operation sequence to the operation patterns has two merits: we can consider the undetection of squaring operation as a numerical noise and we can identify the correct operation pattern by performing a majority voting on operation patterns. In this paper, we demonstrate an

experimental attack on CRT-RSA software in Libgcrypt 1.7.6 to evaluate the practicality and feasibility of an SWL-based attack using the proposed method in terms of the required number of observations and the remaining number of SWL candidates. As a result, we confirm that the proposed method can reduce the number of candidates for operation sequences to only 3 in our setup (this indicates that the number of candidates for RSA secret key can be reduced to at most 300,000) by a combination with the partial key exposure attack by Heninger and Schacham [13] and can recover the secret key of 1024-bit CRT-RSA within a day, while a straightforward estimation of location of capture errors is infeasible. (For typical example, it requires more than 2^{70} guesses of capture error locations.)

Although the basic concept and preliminary evaluation was presented in the previous version [14], in this paper, we show a further analysis to show the effectiveness of the proposed method. In particular, we show an analysis on the number of traces required for recovering correct sequence and derive the minimum number of traces for the key recovery. Moreover, we also investigate and discuss the complexity of the proposed attack, to clarify the practicality of the proposed attack. As a result, we confirm that the proposed method can recover the secret key of CRT-RSA in Libgcrypt 1.7.6 with 120 traces, and we also confirm the practicality of the proposed method.

The rest of this paper is organized as follows: Sect. 2 briefly introduces the sliding window exponentiation, how to obtain cache timing information in this paper (namely Flush+Reload), and the SWL-based attack on CRT-RSA software. Section 3 describes the proposed method based on operation patterns and DTW. Section 4 demonstrates an experimental attack to demonstrate the effectiveness of the proposed method. Section 5 contains our conclusion.

2 Preliminaries

2.1 Sliding window exponentiation

Sliding window exponentiation is one of the fastest modular exponentiation algorithms. It employs precomputation of the small odd powers of a base and windowing to reduce the number of multiplications. Algorithm 1 is the left-to-right sliding window method for modular exponentiation, which is target of this study. Here, the inputs X , N , and E correspond to the ciphertext, public key, and secret key of RSA decryption, respectively. Output R corresponds to plaintext. At Line 1, the maximum window size, w , is predetermined as a parameter. Basically, for a larger w , the computation time of the main loop (i.e., Lines 8–22) is faster, while the time for precomputation and the required memory are larger. In Lines 3–6, we precompute the odd powers of X upto X^{2^w-1} . Note here

Algorithm 1 Left-to-right sliding window exponentiation

Input: Base X , modulus N , exponent $E = (e_k e_{k-1} \dots e_1)_2$
Output: Modular exponentiation $R = X^E \bmod N$

```

1: parameter  $w$ ; ▷ Predetermined maximum window size
2: Function SLIDINGWINDOWEXPONENTIATION $_w(X, N, E)$ 
3:   int  $X_1 \leftarrow X$ ; int  $X_2 \leftarrow X^2 \bmod N$ ;
4:   for  $i$  from 1 to  $2^w - 1$  do ▷ Precomputation
5:     int  $X_{2i+1} \leftarrow X_2 X_{2i-1} \bmod N$ ;
6:   end for
7:   int  $R \leftarrow 1$ ; int  $j \leftarrow k$ ;
8:   while  $j \geq 1$  do ▷ Main loop
9:     int  $z \leftarrow \text{LeadingZerosOf}((e_j e_{j-1} \dots e_1)_2)$ ; ▷ Count leading zeros
10:     $j \leftarrow j - z$ ; ▷ Set  $j$  such that  $e_j = 1$ 
11:    int  $l \leftarrow \min(w, j)$ ;
12:    int  $s \leftarrow \text{TrailingZerosOf}((e_j e_{j-1} \dots e_{j-l+1})_2)$ ; ▷ Count trailing zeros
13:    int  $u \leftarrow (e_j e_{j-1} \dots e_{j-l+1})_2 >> s$ ; ▷ Remove trailing zeros
14:    for  $t$  from 1 to  $z + (l - s)$  do
15:       $R \leftarrow R^2 \bmod N$ ;
16:    end for
17:     $R \leftarrow R X_u \bmod N$ ;
18:     $j \leftarrow j - (l - s)$ ;
19:  end while
20:  return  $R$ ;
21: end Function

```

that $X_{2i+1} = X^{2^{i+1}}$. Lines 8–19 represent the main loop of this algorithm. In the main loop, we first count the number of leading zeros of $(e_j e_{j-1} \dots e_1)_2$ at Line 9 and set the pointer j such that $e_j = 1$ at Line 10. Line 11 is required to determine the maximum window size at the end of the main loop. At Line 12, we count the trailing zeros of the j -th- $(j-l+1)$ -th exponent bits $(e_j e_{j-1} \dots e_{j-l+1})_2$ as s , which is equivalent to $\log_2 (\text{GCD}((e_j e_{j-1} \dots e_{j-l+1})_2, 2^l))$, where GCD denotes the greatest common divisor. Here, $l-s$ is the number of exponent bits treated in the current iteration. At Line 13, we remove trailing zeros from $(e_j e_{j-1} \dots e_{j-l+1})_2$ by shifting it to the right by s and derive $u = (e_j e_{j-1} \dots e_{j-s+1})_2$. Note here that u is always odd because e_{j-s+1} should be one. After we perform the squaring operation $z + (l-s)$ times (i.e., the number of leading zeros + temporal window size) at Lines 14–16, we perform multiplication with precomputed X_u at Line 17. Finally, at Line 18, we update j for the next loop. Note here that $e_1 = 1$ because E should be odd in the case of RSA.

Similar to [8], we denote multiplication and squaring by M and S, respectively, and denote operation sequence consisting of multiplication and squaring as the string of S and M such as SSM (which indicates that a multiplication operation is performed after two squaring operations.) In Algorithm 1, for example, if the exponent E is the 16-bit value $(1000101011000011)_2$ and the window size is $w = 4$, the left-to-right operation sequence is given by SSMSSSSSSSMSSSSSSSM. In the sequence, the first, second, third, and fourth multiplications are performed with

X_1 , X_5 , X_3 , and X_3 , respectively. In this example, multiplication is performed only 4 times. In contrast, the left-to-right binary method and Montgomery ladder, which are the typical modular exponentiation algorithms, require 6 and 16 multiplications, respectively. Thus, the sliding window method is known to be one of the fastest modular exponentiation algorithms because of the precomputation, as analyzed in [15]. In addition, left-to-right sliding window exponentiation had been expected to be somewhat resistant to (cache) timing attacks and SPAs because the timing and SPA traces are less informative compared to the left-to-right binary method owing to the operand-selective multiplication (until the publication of the SWL-based attack [8]). Therefore, the left-to-right sliding window method has been widely deployed in numerous cryptographic OSS including GnuPG.

2.2 FLUSH + RELOAD

Flush+Reload is one of the most popular methods for cache attacks [2]. The use of Flush+Reload enables an attacker to know that the data in the memory of a specific location are loaded (and cached) by a victim. If the data are the code segments of multiplication and squaring for (CRT-)RSA decryption, the attacker can obtain the operation sequence of modular exponentiation via Flush+Reload. Flush+Reload is available if the attacker can run a process on a CPU core where the victim's process is running on another core, and their processes share the main memory and L3 cache. Such a scenario is practical for several contemporary cloud services that lease VMs and computational resources.

The basic idea of Flush+Reload is to repeat cache reload and flush periodically to estimate whether data are loaded and used by the victim. Flush+Reload is performed in the following three steps: (i) the attacker flushes the shared cache, (ii) the attacker waits for the victim to load target data (i.e., code segments of multiplication and squaring) from the main memory (the waiting time is referred to as slot time), and (iii) the attacker loads the target data to measure the time (i.e., clock cycles) required for the loading. For example, in x86 CPUs, loading time can be measured using the RDSTC operation. Here, if the victim loads the target data (i.e., performs multiplication and/or squaring) during the slot time, the loading time at Step (iii) should be short because the data should be cached. Otherwise, the loading time should be long because the attacker should load the data from the main memory. Thus, the attacker can estimate the timing of execution of multiplication and squaring and can obtain the operation sequence via cache information.

In this paper, to perform Flush+Reload, we employ an open-source toolkit for micro-architectural side-channel named Mastik, which is available in [16].

2.3 SWL-based attack

In [8], an SWL-based attack is performed in the following four steps: (i) First, the attacker estimates the execution timing of operations related to modular exponentiation (i.e., multiplication and squaring) through a cache attack such as Flush+Reload. (ii) Then, the estimated execution timing is translated to the operation sequence, which represents the order of execution of multiplication (M) and squaring (S). (iii) The key bits are partially estimated using an algorithm proposed in [8], which is based on the windowing rule of the left-to-right sliding window method. (iv) Finally, the entire key is recovered from the partial key by employing the algorithm developed by Heninger and Shacham [13]. It is shown that this attack can reduce the space of a secret key to less than 10^6 for all keys of 1024-bit CRT-RSA, and to 2^6 for 13% keys of 2,048-bit CRT-RSA if the attacker can obtain a complete and correct operation sequence at Step (ii). However, since operation sequences observed via Flush+Reload would usually contain a noise, it would be difficult to obtain such a correct operation sequence. This paper improves Step (ii) to perform the SWL-based attack with tolerating noise.

An SWL-based attack requires Flush+Reload to estimate the execution timing of multiplication and squaring. We describe how to apply Flush+Reload to the implementation in Libgcrypt in order to estimate the execution timing and operation sequence. As the Libgcrypt implementation employs an identical integer multiplication and modular reduction function for both multiplication and squaring¹, the execution timing of the function is not informative. Therefore, we apply probes on the pre-operations and post-operations of the function to distinguish between multiplication and squaring. The sliding window method selects a value (i.e., X_{2i+1}) from a precomputed table before performing a multiplication. In addition, there is a procedure for returning to the start of main loop after the multiplication. We also probe the execution timings of the above two operations (i.e., operand selection and jump operation) to distinguish between squaring and multiplication.

An SWL-based attack should accurately acquire the execution timing of multiplication and squaring to obtain a correct operation sequence. However, in practice, it is quite difficult to obtain such a correct operation sequence because the execution timing estimated by cache attacks, including Flush+Reload, includes noise, which is referred to as capture error. Capture error is classified into misdetection and undetection. Misdetection indicates that the attacker observes an operation via the cache information when the operation is

Fig. 1 Example of capture errors

actually not performed. Misdetection occurs when other data in the same cache line are loaded to the cache. Undetection indicates that the attacker misses the execution of an operation. Figure 1 shows two typical examples of undetection. Undetection occurs when (a) an operation is performed multiple times during a slot time or (b) the execution timing of an operation overlaps with the attacker's reload operation. Such undetection is closely related to the slot time duration, which is determined by the attacker. A shorter slot time mitigates (a), while a longer time mitigates (b). As an integer multiplication and a modular reduction in Libgcrypt (i.e., the target library in this study) require 30,000 clock cycles on average, the slot time should be shorter than 30,000 clock cycles. On the contrary, Allan et al. reported that a slot time longer than 10,000 clock cycles would be effective for preventing (b) [17].

Nevertheless, it is still difficult to obtain a correct and complete operation sequence from a single measurement of cache information for RSA decryption. This indicates that the length of the acquired operation sequence probably differs from that of the actual sequence. While the attacker would be able to observe the execution timing multiple times, there is no known method for constructing a correct operation sequence from such noisy execution timing in literature. Even though the authors of [8] stated that a correct operation sequence could be obtained by aligning numerous execution timings obtained through multiple observations and a simple majority rule, it is not clear how to align the stochastic capture errors. In other words, it is quite difficult to align the obtained noisy operations sequence and to find the correspondence of S and M among them.

It is also infeasible to estimate the location of capture errors. The length of an (actual) operation sequence depends on the window size, w , and the number of zeros counted as leading zeros. In the 1024-bit CRT-RSA in Libgcrypt, where the length of secret keys is 512 bits and w is four, the average length of an operation sequence is 614. For the simplicity, we assume that the capture errors are caused only by undetection of operation. If the operation sequence contains h undetections of operation, the number of possible locations of capture error is given by $\binom{k'}{h}$, where k' is the length of noisy operation sequence obtained from the cache

¹ In this paper, “multiplication” denotes the multiplication in the sliding window method as opposed to squaring, while “integer multiplication” indicates just a code segment which is used for both multiplication and squaring.

trace. For example, let us consider the case that the length of the correct and noisy operation sequences are 610 and 600, respectively. If we estimate the location of undetection in ascending order about h , the number of guesses until we find the correct one is given by $\sum_{h=1}^{10} \binom{600+h}{h}$ in the worst case, which is too large to estimate a correct operation sequence in an exhaustive search. A method for constructing a correct and complete operation sequence from noisy sequences is required to evaluate the feasibility of SWL-based attacks.

3 Quantitative analysis of SWL availability

We first show the characterization of the execution timing and estimated operation sequence observable via Flush + Reload. While this preliminary evaluation is performed in a similar manner to [8], to the best of the authors' knowledge, this is the first report on such a comprehensive characterization of the estimated operation sequence for evaluating SWL availability. In this experiment, we use an RSA software from Libgcrypt 1.7.6 compiled by gcc 4.8.5 with option `-O2`, employ a public toolkit for micro-architectural side-channel attack named Mastik [16], and run the RSA software on a Linux (CentOS 7.4) PC with an Intel Core i5-3470. We set the slot time of Flush + Reload to 10,000 clock cycles. In the reload step, we assume that the data should be loaded from the cache if the loading clock cycle is less than 100. Otherwise, the data are assumed to be loaded from the main memory. Similarly to [8], we employ a performance degradation attack (PDA), which increases the number of clock cycles (i.e., latency) required for executing target operations [17]. The PDA is useful for decreasing the undetection (b). Finally, as mentioned in Sect. 2.3, we apply four probes in total on the calls of integer multiplication and modular reduction (for each multiplication and squaring), operand selection, and the jump operation.

Figure 2 shows an example of a cache timing trace during CRT-RSA decryption, where the horizontal axis denotes the slot index and the vertical axis denotes the loading time (i.e., the number of clock cycles) for the attacker's reload operation. Here, the lines denoted by "Integer multiplication" and "Modular reduction" show the time required for reloading the code segments of integer multiplication and modular reduction. Note that a pair of integer multiplication and modular reduction is used per multiplication or squaring in the sliding window method. Additionally, we cannot distinguish whether multiplication or squaring is performed only from the traces of integer multiplication and modular reduction. The lines denoted by "Operand selection" and "Jump" enable us to perform this discrimination; that is, if the attacker observes either or both of them, a multiplication should be performed; otherwise, a squaring should be performed. For example, around a slot index of 1655, a squaring

Fig. 2 Example of cache timing trace

Fig. 3 Histogram of length of observed operation sequences

operation would be performed because integer multiplication and modular reduction are loaded from the cache, but operand selection and jump operation are loaded from the main memory. In contrast, multiplication would be performed around a slot index of 1665. Thus, from Fig. 2, we can estimate that the operation sequence corresponding to the cache timing trace at slot indices 1650–1665 is SSSM. (However, it may contain capture errors.)

To evaluate the influence of capture errors on the feasibility of SWL-based attacks, we perform Flush + Reload on CRT-RSA decryption with a fixed key to obtain cache timing traces. Figure 3 displays the histogram of the length of the operation sequence estimated by the obtained cache timing trace, where the length of the correct operation sequence is 617. From Fig. 3, we confirm that the lengths of the estimated operation sequence vary and most of them are shorter than the length of the correct sequence. This indicates that undetection should be dominant among capture errors.

We employ Levenshtein distance (a.k.a edit distance) to evaluate the similarity between estimated operation sequences [18]. The Levenshtein distance between two strings of characters is defined as the minimum number of edit operations, that is, insert, delete, and replace. In the context of Flush + Reload, the characters are given by M and S, and insert, delete, and replace correspond to the undetection, misdetection, and misrecognition of M and S, respectively.² Figure 4 shows the histogram of the Levenshtein distance between the correct and estimated operation sequences, and Fig. 5a, b, c shows the histograms of number of insert, delete, and replace operations, respectively. From these figures, we

² We define the direction of insertion, deletion, and replacement as from the captured sequence to the correct one.

Fig. 4 Histogram of Levenshtein distance

confirm that most capture errors are caused by the undetection of operations and the pattern of capture errors varies for each measurement. Undetection occurs in every observation, while misdetection and misrecognition occur rarely. This leads to the following important heuristic: a longer observed operation sequence contains less capture errors. However, as the locations of insert are unknown to the attacker, she cannot align noisy operation sequences by finding the correspondence between them, as mentioned in Sect. 2.3. Thus, it is impossible to obtain a correct operation sequence by averaging noisy and stochastic operation sequences. Moreover, the histograms of Levenshtein distance shown in Figs. 4 and 5 represent a complex and composite distribution and not a simple Gaussian distribution. This indicates that we require a dedicated strategy to evaluate the location and type of capture errors in a quantitative manner.

In this paper, we employ the DTW algorithm to quantitatively evaluate capture errors [12]. This algorithm is frequently used for the visualization of Levenshtein distance. First, we transform the operation sequence into another representation to which DTW can be applied. The operation sequence generated by the sliding window method should repeat one or more squaring operations followed by a multiplication. Therefore, the operation sequence of sliding window exponentiation is represented by $S^{t_1}M S^{t_2}M \dots S^{t_u}M \dots S^{t_v}M$, where t_u is an integer greater than one, where $1 \leq u \leq v$. We refer to one $S^{t_u}M$ as an “operation pattern,” and v denotes the number of operation patterns. Each operation pattern can be represented by an integer (i.e., t_u). In other words, we can transform the operation sequence to a stacked bar chart of t_u 's representing the operation pattern. Figure 6 illustrates an example of the translation of an operation sequence to pattern sequence and a stacked bar chart. We call a chain of operation patterns obtained via a cache trace “pattern-chain.”

DTW is an algorithm for measuring the similarity among multiple time series with different lengths. DTW finds the correspondence of a point of a series to other points in order to align the multiple time series. Figure 7 shows an example of alignment by DTW. Here, DTW aligns two time series, $\mathbf{A} = (a_1, a_2, \dots, a_x, \dots, a_{10}) = (2, 6, 6, 8, 4, 9, 6, 5, 3, 4)$

Fig. 5 Histogram for edit operations: **a** insert, **b** delete, and **c** replace

Operation sequence	SSSMSSSSSSSMSSSSSMSSSSSSSSSSSSSM				
Operation pattern	SSSM	SSSSSSSM	SSSSM	SM	SSSSSSSSSSSSSM
Assigned value (i.e., t_u)	3	7	4	1	15

Fig. 6 Example of transformation of operation sequence to operation pattern and t_u

and $\mathbf{B} = (b_1, b_2, \dots, b_y, \dots, b_8) = (2, 6, 5, 4, 8, 6, 5, 3)$. In the two-dimensional table, the value in the cell at (x, y) , which indicates the DTW value between a_x and b_y , is given by $(a_x - b_y)^2 + \min(p(x-1, y-1), p(x-1, y), p(x, y-1))$, where $p(x, y)$ denotes the value of the cell at (x, y) .

Note here that $p(0, 0) = 0$ and $p(0, y) = p(x, 0) = \infty$ for any x and y satisfying $x \neq y$. After calculating this table from $p(1, 1)$ to $p(10, 8)$, we determine a path from $(1, 1)$ to $(10, 8)$ such that the sum of the values in the path is minimized. The path is referred to as the warping path, and it represents correspondence (i.e., alignment). In Fig. 7, the cells highlighted in orange color show the warping path, which indicates $(a_1) \sim (b_1)$, $(a_2, a_3, a_4) \sim (b_2)$, $(a_5) \sim$

Fig. 7 Example of DTW

$(b_3, b_4), (a_6) \sim (b_5), (a_7) \sim (b_6), (a_8) \sim (b_7)$, and $(a_9, a_{10}) \sim (b_8)$, where $(A) \sim (B)$ denotes that A corresponds to B . The computational complexity of DTW is given as the product of lengths of the time series. In the case of an SWL-based attack, the complexity is given by $\mathcal{O}(v^2)$, which is sufficiently feasible.

We apply DTW to the correct operation pattern and a noisy pattern obtained by applying Flush+Reload to CRT-RSA decryption, in order to demonstrate the usefulness of DTW in the context of SWL-based Attack. Figure 8 shows the result of DTW, where the upper and lower stacked bar charts are the estimated and correct pattern-chains, respectively. Here, the lengths of correct and observed pattern-chains are 104 and 102, respectively. Furthermore, the lengths of correct and observed pattern sequences are 616 and 611, respectively. The dotted line represents the correspondence between them. The bars denoted by red and different colors represent that the lengths of the operation patterns (i.e., t_u 's) are equal and not equal to each other, respectively. The inequality of the length of a pattern-chain is basically caused by the undetection of M and misdetection of S in the observed pattern. However, DTW successfully aligns the operation patterns while considering such capture errors. Importantly, DTW appears to work well even if the number of operation patterns (i.e., v) is different from each other; that is, there is an undetection of M in the observed pattern. Thus, we can obtain the correspondence between two operation patterns and specify the location and type of capture errors for the quantitative evaluation.

Fig. 8 A result of DTW, which shows correspondence between observed and correct operation patterns

4 Constructing correct operation sequence

4.1 Proposed method

This section presents a method for constructing the correct operation sequence based on the discussion in Sect. 3. We assume that the attacker can obtain numerous noisy operation sequences for a fixed key through CRT-RSA decryption.

In Sect. 3, we confirmed the usefulness of representing the operation sequence as a pattern-chain (i.e., a stacked bar chart given by $(t_1, t_2, \dots, t_u, \dots, t_v)$), and that most lengths of patterns in a pattern-chain obtained from Flush+Reload are basically equal to the length of the corresponding correct patterns. On the other hand, it is difficult to align the pattern sequences if there are some undetections of M's. Therefore, the proposed method employs the transformation of the operation sequence to a pattern-chain, classifies the pattern sequences by their lengths, and performs operation-pattern-wise majority voting for each classes. Figure 9 shows an example of majority voting with 3 pattern-chains.

More precisely, we first transform all obtained operation sequences into the corresponding pattern sequences (and pattern-chains), and we classify the pattern-chains by their length. Here, the longer pattern sequences are meaningful for us because such longer pattern sequences would have less capture errors of misdetection of M's, as discussed in Sect. 3. Let $v_{\max}$ be the length of longest pattern-chain obtained via cache traces. In the proposed method, we first focus on the longest pattern-chains (i.e., pattern-chains with $v_{\max}$) to estimate the correct operation pattern. The u -th operation pattern of the correct pattern-chain is estimated as the majority vote of the u -th operation pattern of pattern-chains with $v_{\max}$. Then, we apply the partial key exposure attack using Heninger–Schacham algorithm and perform a brute-force of key candidates obtained by Heninger–Schacham algorithm. If the correct key is not found, we then focus on the second longest pattern-chains (i.e., series of operation patterns with $v_{\max} - 1$). As same as the above, we estimate the correct pattern-chain with an identical length by majority vote, apply Heninger–Schacham algorithm, and then perform a brute-force. At the w -th iteration ($w \geq 1$), we focus on the pattern-chains with $v_{\max} - w + 1$ for the estimation. Thus, the proposed method can reduce the key space of 1024-bit CRT-RSA to up-to $r \times 10^6$ from the noisy operation sequences

Fig. 9 Example of majority voting in proposed method

observed through cache traces, where r is the number of repetitions.

Whereas the proposed method does not directly employ techniques related to Levenshtein distance nor DTW, the proposed method can be validated from the viewpoint of DTW costs. According to the evaluation in Sect. 3, we assume that the correct pattern-chain should have the minimum sum of DTW costs for all pattern-chains obtained via cache traces, if we have the sufficient number of traces. However, the estimation of such a pattern-chain is not practical due to the difficulty in estimating/tolerating insertion and deletion. Note here that insertion(s) and deletion(s) of S are interpreted as a replacement of an operation pattern, whereas insertion and deletion of M are still equivalent of an insertion and a deletion of an operation pattern, respectively. Therefore, instead of estimating the location of insertion and deletion of operation patterns, we classify the pattern-chains by their length. Here, we assume that the pattern-chains with an identical length have the identical number of insertion/deletion, because the insertion/deletion of M is far less than that of S as evaluated in Sect. 3. This assumption also indicates that the pattern-chains with the length same as the correct one have no insertion/deletion, whereas a few pattern-chains may have pair(s) of insertion and deletion. Then, we perform the majority voting, which minimizes the sum of DTW costs between the correct pattern-chain and observed pattern-chains, in terms of replacements of operation pattern. Thus, we obtain the correct pattern-chain as one with the minimum sum of DTW costs from the observed pattern-chains.

Algorithm 2 is the algorithmic description of whole the attack. At Line 1, we first convert operation sequences to the corresponding pattern-chains. A pattern-chain is given as a list of operation patterns, which contains an integer t_u at the u -th element. Lines 4–23 are the main loop of this algorithm about w . At Line 6, we derive the pattern-chains with a length of $v = v_{\max} - w + 1$. At Lines 8–14, we obtain a candidate of correct operation patterns $c_1, c_2, \dots, c_u, \dots, c_v$ by a majority voting from 1 to v . After deriving a candidate of correct operation patterns, at Line 15, we obtain the partial key given from $c_1, c_2, \dots, c_u, \dots, c_v$ according to the reversing algorithm in [8]. At Line 16, we obtain a set of candidates for the secret key by partial key exposure attack (e.g., Heninger–Schacham algorithm for the case of CRT-RSA). Then, at

Algorithm 2 Proposed attack

Input: Noisy operation sequences obtained via cache traces

Output: Secret key d

```

1: set  $\mathcal{P} \leftarrow \text{ConvertToPatternChain}(\text{Operation sequences});$ 
2: int  $v_{\max} \leftarrow \text{Length of longest pattern-chain in } \mathcal{P};$ 
3: int  $w \leftarrow 1;$ 
4: while 1 do
5:   int  $v \leftarrow v_{\max} - w + 1;$ 
6:   set  $\mathcal{P}_v \leftarrow \text{Subset}(\mathcal{P}, v);$ 
7:   set  $\mathcal{T}_u \leftarrow \{\};$ 
8:   for  $u$  from 1 to  $v$  do
9:     for each  $T \in \mathcal{P}_v$  do
10:      int  $t_u \leftarrow T[u];$ 
11:       $\mathcal{T}_u \leftarrow \mathcal{T}_u \cup \{t_u\}$ 
12:     end for
13:     int  $c_u \leftarrow \text{MajorityVoting}(\mathcal{T}_u);$ 
14:   end for
15:   partial key  $d' \leftarrow \text{ReversingExponent}(c_1, c_2, \dots, c_v);$ 
16:   set  $\mathcal{D} \leftarrow \text{PartialKeyExposureAttack}(d');$ 
17:   for each  $d \in \mathcal{D}$  do
18:     if  $\text{IsCorrectKey}(d)$  then
19:       return  $d;$ 
20:     end if
21:   end for
22:    $w \leftarrow w + 1;$ 
23: end while

```

Line 17–21, we perform a brute-force on the secret key candidates. If the correct key is found, this algorithm returns it and stops. Whereas this algorithm is described for RSA without CRT for the simplicity, it can be easily extended to CRT-RSA by performing Lines 8–14 (i.e., majority voting) twice for d_p and d_q .

Algorithm 2 must stop assuming that the correct pattern-chain is reconstructed by the proposed method, because the correct key must be included in $\mathcal{D}$ for the correct v at Line 16. The computational time of Algorithm 2 heavily depends on the number of iterations on w . The next subsection experimentally validates that this algorithm should stop for the most cases until the fifth iteration and would be sufficiently feasible.

4.2 Experimental validation

This section demonstrates the validity of the proposed method through an experimental attack. We use the same experimental setup as that described in Sect. 3. We obtain 100,000 cache timing traces for a fixed key via Flush + Reload, in which the length of the correct pattern sequence is 104. As a result, we confirm that we successfully construct the correct operation pattern (and sequence) from noisy observed operation sequences when v is equal to the correct length.

Figure 10 shows the histogram of the length of the observed pattern sequences, where the maximum observed length is 106 and the mode is 102. Given that the correct length is 104, Fig. 10 validates our heuristic—“a longer observed pattern sequence would have less capture errors

Fig. 10 Histogram of length of operation pattern sequence obtained by Flush + Reload

and would be close to the correct pattern.” At least, the length of the correct pattern sequence should be longer than the mode. Thus, the proposed method performs the majority voting from the largest v , which would allow for efficient estimation. In this case, as the maximum length in Fig. 10 is 106, we find the correct key using the proposed method at the third execution of Heninger–Schacham algorithm following from operation-pattern-wise majority vote. We also implement Heninger–Schacham algorithm and measure the execution time. We confirm that one execution of Heninger–Schacham algorithm for 1024-bit CRT-RSA takes about 3.5h in our experiment. As a successful SWL-based attack can reduce key space to at most 10^6 as described in [8], we can retrieve the secret key with an exhaustive search of at most 3×10^6 in the above experiment, which is sufficiently feasible. Since the correct operation pattern can be found until the fifth execution of Heninger–Schacham algorithm, the proposed attack can retrieve the secret key within a day including a brute-force of key candidates. Thus, we confirm the effectiveness of the proposed method and the practicality of SWL-based attacks.

To analyze the number of traces required for recovering the correct pattern, Fig. 11 shows the number of correct operation patterns in pattern-chains constructed by the proposed method for different number of traces. Here, the DTW method is used for evaluation of the distance between correct and constructed operation patterns. In Fig. 11, the horizontal axis shows the number of traces and the vertical axis is the number of correct patterns in constructed pattern sequence. We can find that the number of correct patterns rapidly increases when the number of traces is less than 20. In addition, the operation patterns are always correctly constructed when using more than 120 traces in this experiment. This result shows that we can obtain the correct SWL from 101 noisy operation sequences. Thus, we can confirm the practicality of the proposed method.

Fig. 11 Number of correct operation patterns in pattern-chain constructed by proposed method for different numbers of traces

5 Conclusion

This paper presents a quantitative analysis of feasibility of SWL-based attacks in terms of the availability of cache timing leaks and a method for constructing a correct operation sequence from noisy and stochastic operation sequences. We employ a representation of operation sequences, named operation patterns, to which the DTW algorithm can be applied for evaluation. This leads to a heuristic that indicates that the length of the pattern-chains would have less capture errors. The proposed method for constructing correct operation patterns also employs the transformation to operation patterns. The result of an experimental attack on OSS-RSA in Libgcrypt shows that the proposed method accurately estimates the correct operation sequence from only 100 noisy observed sequences and it can retrieve the secret key of CRT-RSA within practical complexity.

A quantitative analysis of the number of cache timing traces required for a successful attack is part of future work. In addition, while DTW and Levenshtein distance are used only for evaluating the operation patterns but not used for the reconstruction method in this paper, it would be interesting to develop a method for estimating the correct operation pattern by combining all pattern sequences even with different lengths based on DTW, for improving the efficiency of estimation.

Acknowledgements We are grateful to Mr. Hayato Mori and Mr. Souki Osawa for their cooperation.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material

in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Tromer, E., Osvik, D.A., Shamir, A.: Efficient cache attacks on AES, and countermeasures. *J. Cryptol.* **23**(1), 37–71 (2010)
2. Yarom, Y., Falkner, K.: FLUSH+RELOAD: a high resolution, low noise, L3 cache side-channel attack. In: 23rd USENIX Security Symposium (2014)
3. Inci, M.S., Gülmezoglu, B., Irazoqui, G., Eisenbarth, T., Sunar, B.: Cache attacks enable bulk key recovery on the cloud. In: International Conference on Cryptographic Hardware and Embedded Systems. Volume 9813 of Lecture Notes in Computer Science, pp. 368–388. Springer (2016)
4. Yarom, Y., Genkin, D., Heninger, N.: CacheBleed: a timing attack on OpenSSL constant-time RSA. *J. Cryptogr. Eng.* **8**(1), 1–27 (2018)
5. Ge, Q., Yarom, Y., Cock, D., Heiser, G.: A survey of microarchitectural timing attacks and countermeasures on contemporary hardware. *J. Cryptogr. Eng.* **8**(1), 1–27 (2018)
6. Kocher, P., Horn, J., Fogh, A., Genkin, D., Gruss, D., Haas, W., Hamgurg, M., Lipp, M., Mangard, S., Prescher, T., Schwarz, M., Yarom, Y.: Spectre attacks: exploiting speculative execution. In: 40th IEEE Symposium on Security and Privacy (S&P'19) (2019)
7. Lipp, M., Schwarz, M., Gruss, D., Prescher, T., Haas, W., Fogh, A., Horn, J., Mangard, S., Kocher, P., Genkin, D., Yarom, Y., Hamburg, M.: Meltdown: reading kernel memory from user space. In: 27th USENIX Security Symposium (USENIX Security 18) (2018)
8. Bernstein, D.J., Breitner, J., Genkin, D., Bruinderink, L.G., Heninger, N., Lange, T., Vredendaal, C.v., Yarom, Y.: Sliding right into disaster: left-to-right sliding windows leak. In: International Conference on Cryptographic Hardware and Embedded Systems. Volume 10529 of Lecture Notes in Computer Science. Springer (2017)
9. Menezes, A.J., van Oorschot, P.C., Vanstone, S.A.: Handbook of Applied Cryptography. CRC Press, Boca Raton (1996)
10. Kocher, C.P.: Timing attacks on implementations of Diffie–Hellman, RSA, DSS, and other systems. In: Advances in Cryptology—CRYPTO 1996. Volume 1109 of Lecture Notes in Computer Science, pp. 104–113. Springer (1996)
11. GNU Privacy Guard. <https://www.gnupg.org> (2017). Accessed Oct 2017
12. Sakoe, H., Chiba, S.: Dynamic programming algorithm optimization for spoken word recognition. *IEEE Trans. Acoust., Speech Signal Process.* **26**(1), 43–49 (1978)
13. Heninger, N., Schacham, H.: Reconstructing RSA private keys from random key bits. In: Advances in Cryptology—CRYPTO 2009. Volume 5677 of Lecture Notes in Computer Science, pp. 1–17. Springer (2009)
14. Ueno, R., Junko, T., Hayashi, Y.i., Homma, N.: Constructing sliding windows leak from noisy cache timing information of OSS-RSA. In: PROOFS Workshop, pp. 64–77 (2019)
15. Koç, C.K.: Analysis of sliding window techniques for exponentiation. *Comput. Math. Appl.* **30**(10), 17–24 (1995)
16. Yarom, Y.: Mastik: a micro-architectural side-channel toolkit. <https://cs.adelaide.edu.au/~yval/Mastik/> (2018). Accessed Apr 2018
17. Allan, T., Brumley, B.B., Falkner, K., van de Pol, J., Yarom, Y.: Amplifying side channels through performance degradation. In: ACM Annual Conference on Computer Security Applications (ACSAC), pp. 422–435 (2016)
18. Gusfield, D.: Algorithms on Strings, Trees, and Sequences: Computer Science and Computational Biology. Cambridge University Press, Cambridge (1997)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.