

Cracking the Channel Hopping Sequences in IEEE 802.15.4e-Based Industrial TSCH Networks

Xia Cheng

Department of Computer Science
State University of New York at
Binghamton
xcheng12@binghamton.edu

Junyang Shi

Department of Computer Science
State University of New York at
Binghamton
jshi28@binghamton.edu

Mo Sha

Department of Computer Science
State University of New York at
Binghamton
msha@binghamton.edu

ABSTRACT

Industrial networks typically connect hundreds or thousands of sensors and actuators in industrial facilities, such as manufacturing plants, steel mills, and oil refineries. Although the typical industrial applications operate at low data rates, they pose unique challenges because of their critical demands for reliable and real-time communication in harsh industrial environments. IEEE 802.15.4 based Wireless Sensor-Actuator Networks (WSANs) technology is appealing for use to construct industrial networks because it does not require wired infrastructure and can be manufactured inexpensively. Battery-powered wireless modules easily and inexpensively retrofit existing sensors and actuators in industrial facilities without running cables for communication and power. To address the stringent real-time and reliability requirements, WSANs made a set of unique design choices such as employing the Time-Synchronized Channel Hopping (TSCH) technology that distinguish themselves from traditional wireless sensor networks that require only best effort services. The function-based channel hopping used in TSCH simplifies the network operations at the cost of security. Our study shows that an attacker can reverse engineer the channel hopping sequences by silently observing the channel activities and put the network in danger of selective jamming attacks. To our knowledge, this paper represents the first systematic study that investigates the security vulnerability of TSCH channel hopping in IEEE 802.15.4e under realistic traffic. In this paper, we demonstrate the process of cracking the TSCH channel sequences, present two case studies using publicly accessible TSCH implementations (developed for Orchestra and WirelessHART), and provide a set of insights.

CCS CONCEPTS

• Security and privacy → Mobile and wireless security; • Networks → Link-layer protocols;

KEYWORDS

Time-Synchronized Channel Hopping, Selective Jamming Attack, IEEE 802.15.4e, Industrial Wireless Sensor-Actuator Networks

ACM Reference Format:

Xia Cheng, Junyang Shi, and Mo Sha. 2019. Cracking the Channel Hopping Sequences in IEEE 802.15.4e-Based Industrial TSCH Networks. In *IoTDI '19: Internet of Things Design and Implementation, April 15–18, 2019, Montreal, QC, Canada*. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3302505.3310075>

1 INTRODUCTION

The Internet of Things (IoT) refers to a broad vision whereby things, such as everyday objects, places, and environments, are interconnected with one another via the Internet [31]. Until recently, most of the IoT infrastructures and applications developed by businesses have focused on smart homes and wearables. However, it is the “production and manufacturing” IoT, which underlies the Fourth Industrial Revolution (or Industry 4.0), that promises to be one of the largest potential economic effects of IoT [43] – up to \$47 trillion in added value globally by 2025, according to McKinsey’s report on future disruptive technologies [24].

Industrial networks, the underlying support of industrial IoT, typically connect hundreds or thousands of sensors and actuators in industrial facilities, such as manufacturing plants, steel mills, oil refineries, and infrastructures implementing complex monitoring and control processes. Although the typical industrial applications operate at low data rates, they pose unique challenges because of their critical demands for *reliable* and *real-time* communication in harsh industrial environments. Failure to achieve such performance can lead to production inefficiency, safety threats, and financial loss. These requirements have been traditionally met by specifically chosen wired solutions, e.g., the Highway Addressable Remote Transducer (HART) communication protocol [13], where cables connect sensors and forward sensor readings to a control room where a controller collects sensor readings and sends commands to actuators. However, wired networks are often costly to deploy and maintain in industrial environments and difficult to reconfigure to accommodate new production requirements. IEEE 802.15.4-based Wireless Sensor-Actuator Networks (WSANs) technology is appealing for use in industrial applications because it does not require wired infrastructure and can be manufactured inexpensively. Battery-powered wireless modules easily and inexpensively retrofit existing sensors and actuators in industrial facilities without running cables for communication and power.

There have been two major technology breakthroughs in industrial WSANs. An initial breakthrough came in 1997 from the UC Berkeley’s Smart Dust project [28], which demonstrated tiny, low-power motes could sense, compute, and communicate through wireless mesh networks. A second breakthrough came in 2006 with the time-synchronized mesh protocol (TSMP) [29] with a core

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

IoTDI '19, April 15–18, 2019, Montreal, QC, Canada

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-6283-2/19/04...\$15.00

<https://doi.org/10.1145/3302505.3310075>

technology of Time-Synchronized Channel Hopping (TSCH): All devices in a network are time synchronized and hop channels to exploit frequency diversity. The TSCH technology was adopted by the leading industrial WSN standards (WirelessHART [49] and ISA100 [15]) and the one being standardized by IETF (6TiSCH [14]). A decade of real-world deployments of WirelessHART and ISA100 have demonstrated the feasibility of using TSCH-based WSNs to achieve reliable low-power wireless communication in industrial facilities. Therefore, TSCH was amended into the IEEE 802.15.4e standard in 2012 [1] as a mode to support industrial or embedded applications with critical performance requirements.

To address the stringent real-time and reliability requirements, TSCH made a set of unique design choices that distinguish itself from traditional Medium Access Control (MAC) protocols designed for the Wireless Sensor Networks (WSNs) that require only best effort services [22]. Specifically, TSCH divides time into slices of fixed length that are grouped in a slotframe. Nodes are synchronized and share the notion of a slotframe that repeats over time. Frequency diversity is used to mitigate effects of multipath fading and to improve the robustness and the network capacity. Channel hopping is achieved by sending successive packets on different frequencies. All devices in the network compute the channel hopping sequences by following a function. TSCH's function-based channel hopping simplifies the network operations at the cost of security. Our study shows that an attacker can reverse engineer the channel hopping sequences by silently observing the channel activities and put the network in danger of selective jamming attacks, where the attacker jams only the transmission of interest on its specific communication channel in its specific time slot, which makes the attacks energy efficient and hardly detectable. The selective jamming attacks are more severe threats in WSNs compared to the simple jamming attacks, because jamming a channel or the whole spectrum continuously can be easily detected and located by a wireless intrusion prevention system (WIPS) [34, 50, 51, 54]. Many countermeasures have been developed in the WSN literature to minimize the damage, such as adjusting routing [7, 16, 52]. However, the existing approaches may fail to detect more stealthy attacks such as selective jamming attacks, because the transmission failures caused by the attacks only happen occasionally and are buried in the normal fluctuations of low-power links.

To our knowledge, this paper represents the first systematic study that investigates the security vulnerability of TSCH channel hopping in IEEE 802.15.4e under realistic settings. The contributions of this work are three-fold:

- We present the security vulnerability of TSCH channel hopping in IEEE 802.15.4e by demonstrating the process of cracking the channel hopping sequences;
- We perform two case studies using publicly accessible TSCH implementations¹;
- We provide a set of insights distilled from our analysis and case studies to secure the network by increasing the cracking difficulty.

¹To avoid bias in our attack design and experiments, we use the implementations provided by the third party researchers in our case studies and have different authors to design the attacking program and configure the networks to collect data traces.

The remainder of the paper is organized as follows. Section 2 introduces the background of TSCH channel hopping. Section 3 presents the security vulnerability by demonstrating the cracking process. Section 4 and Section 5 describe two case studies. Section 6 presents our lessons learned. Section 7 reviews related work. Section 8 concludes the paper.

2 BACKGROUND ON TSCH CHANNEL HOPPING

To provide time-deterministic packet deliveries and combat narrow-band interference and multi-path fading, TSCH combines time-slotted MAC access, multi-channel communication, and channel hopping. TSCH divides time into slices of fixed length that are grouped in a slotframe. Each time slot is long enough to deliver a data packet and an acknowledgement between a pair of devices. Nodes are synchronized and share the notion of a slotframe that repeats over time. Channel hopping is achieved by sending successive packets on different frequencies. TSCH uses the channel hopping sequences, which are fixed and known by all devices in the network, instead of using the pseudo-random ones to minimize the channel synchronization overhead. Initially, 16 different channels are available for communication. Each channel is identified by *ChannelOffset*. However, some of these frequencies could be blacklisted due to low-quality communication and, hence, the total number of channels $N_{channel}$ used for channel hopping may be fewer than 16. In TSCH, a link is defined as the pairwise assignment of directed communication between two nodes in a given time slot on a given channel offset. Hence, a link between two communicating nodes can be represented by a pair of values that specifies the time slot in the slotframe and the channel offset used by the nodes in that time slot. Let $[n, ChannelOffset]$ denote a link between two nodes. Then the communication frequency (channel) to be used for communication in time slot n of the slotframe is derived as

$$f = F[(ASN + ChannelOffset) \% S_{length}] \quad (1)$$

where ASN is the Absolute Slot Number, defined as the total number of time slots elapsed since the start of the network, and “ $\%$ ” is the modulo operator. F is a lookup table that contains a sequence of available physical channels. S_{length} is the length of the sequence. Please note that S_{length} may be larger than $N_{channel}$, implying that some channels appear multiple times in the table F . The first device in the network sets ASN to 1 and the newcomers learn ASN from the existing devices. Each device in the network increments ASN at every time slot and uses it as a global time slot counter. Eq. 1 defines the TSCH channel hopping mechanism by returning a different channel for the same link (*ChannelOffset*) at different time slots. Many links can transmit simultaneously in the same time slot, provided that they are assigned with different *ChannelOffset*. Figure 1 shows an example where the network consists of four links and uses three channels. Each link has been assigned with a *ChannelOffset* (0, 1, or 2) that represents channel 24, 25, or 26 in F and each node keeps tracking ASN . In each time slot, the sender and receiver of an active link use Eq. 1 to compute their communication channel. The table in Figure 1 lists the channel hopping sequence for each link. For example, node b and node a use the channel 25 ($(ChannelOffset + ASN) \% 3 = 1$) to transmit




Figure 1: Example channel hopping sequences with 3 channels used. The numbers beside each link indicate the $ChannelOffset$ assigned to that link.

and receive data in time slot 1 if link $b \rightarrow a$ is active in that slot. The transmission scheduler that runs on top of the MAC layer is responsible for deciding which set of links should be active in each time slot. The underlined numbers in Figure 1 describe an example schedule of active links, which allows node a to collect readings from the rest of the nodes in every four time slots. The IEEE 802.15.4 standard specifies neither any scheduling algorithm nor the way of generating physical channel sequence, but it defines the abovementioned mechanism to execute a schedule provided by the scheduler in the upper layer.

The function-based channel hopping used in TSCH simplifies the network operations because there is no need for the network device to synchronize the channel. In a conventional TSCH network, each node learns the current ASN and the channels used in the network from its neighbors upon joining the network, and then uses those information to generate a channel hopping sequence, typically until it leaves the network. However, the channel sequences generated by TSCH present a strong pattern, which introduces a security vulnerability to the network. We will present how an attacker derives the channel hopping sequences without knowing any parameter of Eq. 1 in the following section.

3 VULNERABILITY ANALYSIS

We let $ASN = N_s * N + X$ and rewrite Eq. 1 as

$$f = F[(N_s * N + X + ChannelOffset) \% S_{length}] \quad (2)$$

where N_s is the number of time slots in the combined slotframe², N is the number of slotframes elapsed since the start of the network, and X is the time slot offset in the combined slotframe. In this section, we demonstrate how an attacker, without any prior knowledge on the operating network (any parameter of Eq. 2), cracks the channel hopping sequences by silently observing the channel activities.

The attacker is assumed to be a device that is capable of monitoring all transmission activities on all 16 channels in 2.4 GHz ISM band in the network and has moderate computational capability (e.g., a Raspberry Pi 3 Model B [26] integrating with a Wi-Spy USB Spectrum Analyzer [47]). Today, many TSCH networks are deployed in open fields to support wireless monitor and control applications (e.g., in oil drilling plants). The attacker may be deployed or airdropped into the field and powered by batteries or energy harvesting. The intention of the attacker is to reverse engineer the channel hopping sequences by silently observing the

channel activities and then launch selective jamming attacks, where the attacker jams only the transmission of interest on its specific communication channel in its specific time slot. Please note that the IEEE 802.15.4 standard leaves the upper layer protocol to decide whether to encrypt the parameters in Eq. 1 during transmissions. However, leaving it unprotected makes the problem trivial and the networks vulnerable to attacks. Therefore, in this paper, we assume that the parameters in Eq. 2 are encrypted and the attacker can only get information from the unencrypted MAC message header³ and cannot understand the encrypted payload that stores ASN, $ChannelOffset$, and S_{length} .

TSCH allows a transmitter to skip some scheduled transmissions at runtime (e.g., skipping the retransmission if the first attempt succeeds, skipping the routing traffic if no update is needed), significantly increasing the difficulty of cracking. We start our analysis from a basic case where transmitters transmit packets in all scheduled slots (*ideal case*) and then extend our analysis to a realistic case where transmitters may skip some time slots with scheduled transmissions (*realistic traffic*).

3.1 Cracking the Channel Hopping Sequences in the Ideal Case

In this case, we assume that all devices transmit packets in their scheduled time slots and the attacker begins to eavesdrop on the channels in the time slot S_1 . The encrypted ASN of slot S_1 is unknown to the attacker. Here are the key steps on how an attacker cracks the channel hopping sequences, identifies the slots with scheduled transmissions in the slotframe, and predicts the future channel usage:

- (1) **Grouping the eavesdropped packets:** The attacker snoops the channels and groups the eavesdropped packets based on their source and destination addresses stored in the unencrypted MAC message headers. The attacker then identifies the channel usage sequence of each network device.
- (2) **Identifying the least common multiple of N_s and S_{length} (denoted as $LCM(N_s, S_{length})$):** According to Eq. 2, each network device must use the same channel in the time slot S_1 and $S_1 + LCM(N_s, S_{length})$, S_2 and $S_2 + LCM(N_s, S_{length})$, \dots . In other words, the channel hopping sequence used by each device repeats in every $LCM(N_s, S_{length})$ time slots. Based on each network device's channel usage sequence, the attacker identifies its usage repetition cycle and measures its time duration $T_{repetition}$. The attacker can also derive the length of a time slot T_{slot} by measuring the minimum time duration between the start of two consecutive transmissions. Finally, the attacker gets $LCM(N_s, S_{length}) = T_{repetition} / T_{slot}$. Please note that the measured $T_{repetition}$ of some network devices may be less than $LCM(N_s, S_{length}) * T_{slot}$, thus the attacker should use the largest value among all measured $T_{repetition}$.
- (3) **Identifying the time slots with scheduled transmissions:** From the eavesdropped transmission activities, the attacker identifies the time slots that are scheduled for transmissions in the slotframe.

²TSCH allows the upper layer protocol to define more than one type of slotframes. All slotframes are merged into a single combined slotframe for execution at runtime.

³Due to the overhead concern, IEEE 802.15.4e does not require any encryption to the MAC message header.




Figure 2: Section 3.1 example: four slots in the slotframe, three channels (24, 25, and 26) used in the network, and four transmissions scheduled in each slotframe.

Table 1: Channel Offset Table (P increases by 1 every 12 slots).

Link	Slot No.	Channel	Link	Slot No.	Channel
b->a	$12 \cdot P + 3$	26	c->a	$12 \cdot P + 2$	26
b->a	$12 \cdot P + 7$	24	c->a	$12 \cdot P + 6$	24
b->a	$12 \cdot P + 11$	25	c->a	$12 \cdot P + 10$	25
d->c	$12 \cdot P + 4$	26	e->c	$12 \cdot P + 1$	24
d->c	$12 \cdot P + 8$	24	e->c	$12 \cdot P + 5$	25
d->c	$12 \cdot P + 12$	25	e->c	$12 \cdot P + 9$	26

- (4) **Creating a channel offset table:** The goal of cracking the channel hopping sequences by an attacker is to predict the future channel usage and then perform selective jamming attacks. Thus, there is no need for an attacker to obtain the actual values of F , S_{length} , N , X , and $ChannelOffset$ in Eq. 2. The attacker can assume the time slot S_1 is the first slot in the slotframe and set $N = 0$, and then create a table that pairs each time slot with scheduled transmission (between slot S_1 and $S_{LCM(N_s, S_{length})}$) to a channel for each link.

After deriving $LCM(N_s, S_{length})$, the time slots with scheduled transmissions in the slotframe, and the channel offset table, the attacker knows the exact channel hopping sequence of each link in future, and thus can perform precise strikes to any transmission of interest. The channel hopping sequences can be cracked without error within the bounded time $2 * LCM(N_s, S_{length}) * T_{slot}$ ⁴.

We use an example to illustrate the cracking process. We assume that an attacker is placed in the network presented in Figure 1 and begins to snoop the channels when $ASN = 3$. The attacker does not know ASN and assumes $P = 0$. After snooping for a while, the attacker observes some activities on three channels (channel 24, 25, and 26) and finds that the channel usage repeats in every 12 time slots, as Figure 2 shows. The attacker then derives $LCM(N_s, S_{length}) = 12$ and finds out that all four slots are scheduled with transmissions. Finally, the attacker generates Table 1 and uses it to predict the channel usage for each link in future slots.

3.2 Cracking the Channel Hopping Sequences under Realistic Traffic

In this case, the attacker cannot easily derive $LCM(N_s, S_{length})$ by identifying the channel usage repetition, because a transmitter may skip some scheduled time slots at runtime, which breaks the repetition pattern. The realistic traffic poses a significant challenge

⁴We assume that the attacker needs to snoop $2 * LCM(N_s, S_{length})$ slots to confirm the channel usage repetition cycle.

Algorithm 1: N_s Identification Algorithm

Input : $TSUR[]$
Output: N_s

- 1 Initialize $Density[]$ to 0 and $Position[][]$ to 0;
- 2 **for** $i = 1; i \leq N_u; i++$ **do**
- 3 **for** $j = 1; j \leq N_r; j++$ **do**
- 4 **if** $Position[i][TSUR[j] \% i] == 0$ **then**
- 5 $Position[i][TSUR[j] \% i] = 1;$
- 6 $Density[i] = Density[i] + 1/i;$
- 7 **end**
- 8 **end**
- 9 **end**
- 10 Output k ($Density[k]$ is the smallest value in $Density[]$);




Figure 3: Algorithm execution example ($i = 4$ and $TSUR[] = \{1, 5, 10, 14, 17, \dots\}$).

for the attacker to derive N_s , because it cannot easily pinpoint the beginning and the end of a channel usage repetition cycle. However, we find that the attacker is able to accomplish the cracking by employing a “trial-and-error” learning method. Algorithm 1 shows the method that derives N_s . Please note that the attacker does not know ASN because it does not know when the network starts. Instead of using ASN , the attacker defines OSN (Observation Slot Number) and sets the first observed time slot S_1 with $OSN = 1$. If the attacker observes the actual transmission in a time slot, it adds its OSN value into the array, named $TSUR$ (Time Slot Usage Record). The $TSUR$ array is the input of Algorithm 1. N_r denotes the number of usage records that the attacker collects (number of elements in $TSUR[]$). Algorithm 1 first defines a one-dimensional array $Density$ and a two-dimensional array $Position$ with initial values (line 1). Each element $Density[i]$ stores the weighting factor and helps the attacker to identify the likelihood of i being N_s . Each element $Position[i][j]$ indicates whether the time slot with offset j is scheduled for transmission by assuming $i = N_s$.

The two-level nested loop computes the $Density$ value for each possible value of N_s according to $TSUR[]$ (line 2 – 9). The outside loop traverses all possible values of N_s (from 1 to N_u), where N_u is the upper bound of N_s . The inside loop traverses all records in $TSUR[]$ (from 1 to N_r). At each iteration, Algorithm 1 marks the time slots with scheduled transmissions with 1 in the $Position[i][j]$ array by assuming $i = N_s$ (line 4 – 7). $TSUR[j] \% i$ indicates the corresponding slot offset for the record $TSUR[j]$ when applying it to a slotframe that consists of i slots. If the offset ($Position[i][TSUR[j] \% i]$) has not been previously labeled by any record, Algorithm 1 marks it with 1 (line 5) and increases $Density[i]$ by $1/i$ (line 6). As an example, Figure 3 illustrates the first five iterations of the inside loop when $i = 4$ in Algorithm 1. In the first

iteration, $TSUR[1] = 1$, so the condition $Position[4][1\%4] == 0$ is met. $Position[4][1]$ is then set to 1 and $Density[4]$ increases by $1/i = 1/4$. In the second iteration, $TSUR[2] = 5$, so the condition $Position[4][5\%4] == 0$ is not met. $Density[4]$ does not change. Similarly, $Density[4]$ increases by $1/4$ in the third iteration and does not change in the fourth and fifth iterations. After the outside loop exits, Algorithm 1 outputs the array index with the smallest value in $Density$ array. The index is either N_s or a multiple of N_s ⁵. In other words, the output of Algorithm 1 $\in \{m * N_s | m \in \mathbb{N}^+\}$. We prove the statement by contradiction.

PROOF. We assume that there exists n ($n\%N_s \neq 0$) and $Density[n] < Density[m * N_s]$ ($\forall m \in \mathbb{N}^+$) and separate the proof into two cases: (1) N_s and n do not share any common factor and (2) N_s and n share at least one common factor.

Case 1: We assume that the first element in the $Position[][]$ array marked as 1 by Algorithm 1 when $i = n$ is $Position[n][p]$. Since N_s and n do not share any common factor, Algorithm 1 marks $Position[n][p]$, $Position[n][(p + N_s)\%n]$, ..., $Position[n][(p + (n - 1)N_s)\%n]$ as 1 after executing the line 3–8 if there are sufficient observations in $TSUR[]$. After Algorithm 1 exits, the first n elements of $Position[n][]$ are all marked as 1. So we have $Density[n] = 100\% \geq Density[N_s]$, contradicting the assumption.

Case 2: We assume that there exists n that shares at least one common factor with N_s and the largest common factor (LCF) of N_s and n is $LCF(N_s, n)$. We divide the slotframe consisting of N_s slots into $\frac{N_s}{LCF(N_s, n)}$ blocks, each of which has $LCF(N_s, n)$ slots. We then divide the slotframe into $\frac{n}{LCF(N_s, n)}$ blocks, each of which has $LCF(N_s, n)$ slots. We define the densities of the blocks in the slotframe with N_s slots as: $\rho_1, \rho_2, \rho_3, \dots, \rho_{\frac{N_s}{LCF(N_s, n)}}$. Similarly, we define the densities of the blocks in the slotframe with n slots as: $\rho'_1, \rho'_2, \rho'_3, \dots, \rho'_{\frac{n}{LCF(N_s, n)}}$. We now treat each block as a single unit.

$\frac{N_s}{LCF(N_s, n)}$ and $\frac{n}{LCF(N_s, n)}$ do not share any common factor. So we can convert Case 2 into Case 1. In the fourth line of Algorithm 1, when $i = n$, it maps each element in $TSUR[]$ to a block in the slotframe consisting of n slots ($TSUR[]\%n$). According to the proof for Case 1, all $\frac{n}{LCF(N_s, n)}$ blocks are eventually marked by all blocks in the slotframe consisting of N_s slots after Algorithm 1 finishes executing the line 3–8 if there are sufficient observations in $TSUR[]$. So we have

$\rho'_x \geq \max\{\rho_1, \rho_2, \rho_3, \dots, \rho_{\frac{N_s}{LCF(N_s, n)}}\}$ ($\forall x \in [1, 2, 3, \dots, \frac{n}{LCF(N_s, n)}]$) and then the density of the slotframe when assuming the slotframe has n slots is

$$\begin{aligned} Density[n] &= \frac{\rho'_1 + \rho'_2 + \rho'_3 + \dots + \rho'_{\frac{n}{LCF(N_s, n)}}}{\frac{n}{LCF(N_s, n)}} \\ &\geq \frac{n}{LCF(N_s, n)} \times \frac{\max\{\rho_1, \rho_2, \rho_3, \dots, \rho_{\frac{N_s}{LCF(N_s, n)}}\}}{\frac{n}{LCF(N_s, n)}} \\ &= \max\{\rho_1, \rho_2, \rho_3, \dots, \rho_{\frac{N_s}{LCF(N_s, n)}}\} \\ &\geq \frac{\rho_1 + \rho_2 + \rho_3 + \dots + \rho_{\frac{N_s}{LCF(N_s, n)}}}{\frac{N_s}{LCF(N_s, n)}} = Density[N_s] \end{aligned}$$

⁵Using N_s or a multiple of N_s to generate the channel offset table are functionally equivalent for the attacker. The only difference is the size of channel offset table.



Figure 4: Testbed consisting of 50 TelosB motes.

So we have $Density[n] \geq Density[N_s]$, contradicting the assumption.

With the proof for both cases, we finish the proof by contradiction. ■

The time complexity of Algorithm 1 to derive N_s is $O(N_u * N_r)$. If the attacker uses a constant N_u , the time complexity becomes $O(N_r)$. Please note that there is no need to rerun Algorithm 1 when the attacking program obtains new records. The attacking program can only take the new records as input and process them (executing line 2 – 9) based on the existing $Density$ array and $Position$ array.

After obtaining N_s (or its multiple), the attacker can identify the repetition cycle $LCM(N_s, S_{length})$ by exploring all possible S_{length} values. If every two transmissions with a time interval of $LCM(N_s, M)$ slots always use the same channel, $LCM(N_s, M)$ can be used as the repetition cycle. The attacker then follows the same methods presented in Section 3.1 to identify the time slots with scheduled transmissions, and generate the channel offset table.

4 CASE STUDY ON ORCHESTRA

In this section, we present our case study on cracking the channel hopping sequence of the TSCH implementation [9] in Contiki operating system [8] developed for Orchestra [11] and 6TiSCH networks [14]⁶. Orchestra proposes an autonomous transmission scheduling method running on top of RPL [37] and TSCH networks. Each node computes its transmission schedule locally based on its routing state and MAC address. All nodes running Orchestra change the channels together following the TSCH channel hopping method (Eq. 2). Orchestra employs three types of slotframes for three different kinds of traffic: application, routing, and time synchronization. Different types of slotframes are assigned with different lengths. Orchestra allows S_{length} to be larger than $N_{channel}$.

4.1 Experimental Methodology

We run the experiments on a testbed that consists of 50 TelosB motes [42] deployed on a single floor of an office building [38]. Figure 4 plots the testbed topology. We configure the network to have a single access point and 49 network devices operating on four channels (the default value in Orchestra). The slotframe lengths for application, routing, and time synchronization are 47, 31, and 61, respectively. The combined slotframe has 88,877 time slots in total. Each network device generates a packet every 20s. The

⁶The implementation is provided by Duquennoy et al. and is publicly accessible [10].




Figure 5: Cracking performance on Orchestra with different snooping periods.

attacking program runs on a Raspberry Pi equipped with a 1.2GHz 64-bit quad-core processor and 1.0 GB memory. We perform three sets of experiments. We first measure the prediction performance and cracking time when the attacker snoops different amount of time before launching the attack. We then examine the impact of slotframe length on the cracking performance. Finally, we study the cracking performance when employing a tailored version of the attacking program to Orchestra. We record all the channel activities during the experiments and use them as the ground truth.

4.2 Cracking Performance with Different Snooping Periods

We configure the attacking program to start cracking after snooping the channel activities during a certain number of time slots (snooping period). We vary the length of snooping period from 88,877 slots (1 combined slotframe) to 2,133,048 slots (24 combined slotframes). The channel usage during the snooping period is used as the training set and the channel usage of the next 1,599,786 slots (18 combined slotframes) is taken as the validation set. Our cracking program provides predicted N_s and $LCM(N_s, S_{length})$, identifies the future slots with scheduled transmissions, and predicts the channels used by future transmissions. We compare the predicted transmission activities and their channels against the ground truth in the validation set. If the predicted time slots with transmissions (and corresponding channels) and the ones without transmissions match the ground truth, they are labeled as True Positive (TP) and True Negative (TN), respectively. The wrong predictions are marked as False Positive (FP) and False Negative (FN). After labeling all predictions, we compute the True Positive Rate ($TPR = TP/(TP+FN)$), True Negative Rate ($TNR = TN/(TN+FP)$), Accuracy ($Accuracy = (TP + TN)/(TP + FN + FP + TN)$), and Precision ($Precision = TP/(TP + FP)$).

Figure 5 plots TPR, TNR, and Accuracy of the predictions with different amount of training data (snooping period). As Figure 5 shows, TPR and Accuracy are small (9.65% and 15.55% for TPR, 29.03% and 33.22% for Accuracy) when the eavesdropped number of slots are 88,877 and 177,754 (first two sets of bars). Without enough observations, the cracking program provides a wrong N_s , making the predictions very inaccurate. TPR and Accuracy increase sharply to 60.91% and 90.65%, respectively, when there are 266,631 slots (3 combined slotframes) in the training set. Although the predicted N_s provided by the cracking program is still incorrect, it shares a common factor with the actual value, resulting in some correct prediction on the future channel usage. TPR and Accuracy then




Figure 6: Cracking time with different length of snooping.

increase slowly when the training set is increasing from 266,631 to 711,016 eavesdropped time slots. TPR and Accuracy reach 85.15% and 94.18% with 799,893 eavesdropped slots, providing accurate prediction on the channel usage. This is because the training set includes enough observations for the attacking program to produce the correct N_s leading to accurate channel usage prediction. After that, the increases of TPR and Accuracy become moderate when the training set is larger than 1,777,540 slots (TPR ranging from 96.09% to 97.20% and Accuracy ranging from 97.03% to 97.31%). We observe a similar trend on TNR.

Figure 6 shows the time consumed by the attacking program to crack the channel hopping sequence⁷. The time consumption increases linearly from 924s (88,877 slots) to 13430s (2,133,048 slots), which accords with the $O(N_r)$ time complexity of Algorithm 1.

Observation 1: An attacker can predict the TSCH channel hopping sequences accurately under realistic traffic.

Observation 2: A sudden increase on prediction accuracy does not warrant a correct predicted N_s .

Observation 3: The attacker can predict the channel usage accurately when the observations are large enough to derive N_s . After that, more observations are desired to further improve the prediction accuracy with small additional value.

4.3 Impact of Slotframe Length

To explore the impact of slotframe length on the cracking difficulty, we perform five sets of experiments and increase the length of the combined slotframe roughly 10 times for each set (from 54 slots to 578,429 slots). Table 2 lists the number of slots in each type of slotframe as well as the product and LCM of them. In each set of experiments, we run the experiments three times with different

⁷The snooping period is not added into the result.




Figure 7: Cracking performance with different slotframe length.

Table 2: Slotframe Composition.

No.	Syn	Routing	App	Product	LCM
1	27	9	18	4374	54
2	21	10	20	4200	420
3	23	11	17	4301	4301
4	53	23	37	45103	45103
5	397	31	47	578429	578429

amount of training data, namely eavesdropped with 5,000, 30,000, and 150,000 time slots. We use the next 45,000 slots for validation.

Figure 7(a) plots TPR and Precision of the predictions with different combined slotframe lengths. As Figure 7(a) shows, TPR and Precision decrease when the slotframe length becomes larger. For example, with 5,000 eavesdropped time slots in the training data, TPR and Precision are high (89.19% for TPR, 100% for Precision) while the slotframe length is 54 slots. The slotframe length is so short that such few observations are enough to identify the channel usage repetition and provide a correct N_s . TPR decreases sharply to 40.16% when the slotframe length is 420 slots. Since the predicted N_s shares a common factor (210) with the correct one, Precision still can reach 100%. With the slotframe length increasing from 4,301 to 578,429 slots, TPR and Precision decrease significantly (TPR dropping from 35.75% to 13.12%, Precision dropping from 99.49% to 89.04%), indicating that the eavesdropped slots are too short to identify a complete channel usage repetition cycle. Similarly, with 30,000 eavesdropped slots in the training data, both TPR and Precision are 100% when the slotframe includes 54 or 420 slots. This is because the eavesdropped activities are enough for Algorithm 1 to produce the correct N_s and identify all scheduled slots. After the slotframe length reaches 4,301 slots, TPR and Precision experience a significant decrease (TPR ranging from 79.48% to 50.28%, Precision ranging from 83.69% to 61.13%), indicating the eavesdropped slots are too short to identify a complete repetition cycle. A similar trend is observed when there are 150,000 eavesdropped slots in the training data.

Within each group of bars in Figure 7(a), TPR increases with the training data size. For example, TPR increases from 35.75% (5,000 eavesdropped slots) to 79.48% (30,000 eavesdropped slots), and finally becomes 100% (150,000 eavesdropped slots), when the slotframe length is 4,301 slots. 5,000 eavesdropped slots are too

short for Algorithm 1 to pinpoint a repetition cycle, while 30,000 eavesdropped slots are enough to provide the correct N_s . When the training set includes 150,000 eavesdropped slots, it is large enough to identify all slots with scheduled transmissions. Due to insufficient observations, Precision decreases while TPR increases for some group of bars. Precision drops from 89.04% (5,000 slots) to 61.13% (30,000 slots), reaches 38.31% (150,000 slots) when the slotframe length is 578,429 slots. The observations are insufficient for the cracking program, so TPR increases at the cost of generating more FP, making Precision decrease dramatically.

Figure 7(b) presents the time consumed by the attacking program to crack the channel hopping sequences. In each set of experiments, the time consumption increases approximately linearly with the increase of eavesdropped time slots, confirming the time complexity of Algorithm 1. For instance, the time consumption increases from 0.55s (5,000 slots) to 16.01s (30,000 slots), then to 122.20s (150,000 slots) when the slotframe length is 54 slots.

Observation 4: The combined slotframe length plays an important role in keeping the channel hopping sequence unpredictable. A larger slotframe significantly increases the cracking difficulty.

Observation 5: It is beneficial to use a prime number for each slotframe length, which effectively enlarges the combined slotframe.

Observation 6: TPR may increase at the cost of decreasing Precision when the snooping period is smaller than the slotframe length. A low precision caused by insufficient observations may expose the attacker during jamming.

4.4 Impact of Strong Transmission Pattern

When performing the above experiments, we observe that there exist some time slots showing strong cyclic patterns of transmissions, which help the attacking program to identify the repetition cycles. The transmissions cyclic behavior is introduced by the scheduling design in Orchestra. For example, a fixed and shared slot in the routing slotframe is assigned for all network devices to exchange routing related packets including the DODAG Information Object (DIO) and Destination Advertisement Object (DAO) messages and the device i uses the i th slot in the synchronization slotframe to broadcast beacons and j th slot to receive beacons from its parent (device j). After observing the patterns, the attacking program can first extract the channel activities with these patterns from the




Figure 8: Cracking performance after removing the strong repetitions.

observations and then perform the cracking. We repeat the experiments when applying this tailored attacking method. Figure 8(a) compares TPR and Precision between the original attacking program (Origin) and the tailored version (Acceleration) when the length of the training set is 30,000 slots. TPR and Precision are 100% for both methods when the slotframe length is 54 and 420 slots. This is because these observations are enough for both methods to derive the correct N_s . TPR and Precision of Acceleration are much higher than Origin (100% and 100% for Acceleration, 79.48% and 83.69% for Origin) when the slotframe includes 4,301 slots. TPR and Precision of Acceleration are also much higher when the slotframe lengths are 45,103 and 578,429 slots. Figure 8(b) compares the time consumption used by each method with 30,000 eavesdropped slots. For each slotframe length, the time spent for cracking decreases by more than 60% (e.g., from 5.55s to 1.85s for 578,429 slots) after acceleration. Benefiting from the extraction, the actual combined slotframe length decreases dramatically. The attacking program therefore produces better cracking performance with less time consumption.

Observation 7: The strong cyclic behavior of packet transmissions significantly reduces the cracking difficulty.

4.5 Suggestions to Orchestra

Here are some suggestions gathered from our case study:

- It is beneficial to set three different prime numbers as the lengths of the three slotframes in Orchestra. So the length of the combined slotframe is large (LCM of those numbers), significantly increasing the cracking difficulty. The larger a combined slotframe is, the more time an attacker have to spend on snooping and cracking. Orchestra's decision on using prime numbers as the lengths of the slotframes significantly enhances the network security.
- It is beneficial to use multiple slotframes for application traffic. So the length of the combined slotframe is enlarged (LCM of those numbers), which significantly increases the cracking difficulty.
- It is beneficial to randomize the distribution of the slots with transmissions in each slotframe for routing and time synchronization. The strong cyclic behavior of packet transmissions greatly reduces the difficulty of cracking.

Table 3: Six data flows configured in WirelessHart network.

Flow	Sensor	Actuator	Period	Priority
1	147	146	320ms	1
2	144	143	640ms	2
3	105	104	1280ms	3
4	149	102	2560ms	4
5	136	135	5120ms	5
6	137	108	10240ms	6

5 CASE STUDY ON WIRELESSHART

In this section, we present our case study on cracking the channel hopping sequence of the TSCH implementation in TinyOS operating system [18] developed for WirelessHART networks⁸ [19–21]. Typically, a WirelessHART network consists of a gateway, multiple access points, and a set of field devices (sensors and actuators) forming a multi-hop mesh network. The network is managed by a centralized network manager, a software module running on the gateway, which is responsible for generating routes and transmission schedules and maintaining the operation of the network. Different from Orchestra, all devices in the WirelessHART networks follow the channel hopping sequences generated by the network manager. To enhance the network utilization, the network manager assigns different *ChannelOffsets* to different links, allowing up to $N_{channels}$ packets to be transmitted simultaneously in each time slot. WirelessHART supports both source and graph routing. Source routing provides a single route for each data flow, whereas graph routing generates a reliable routing graph in which each device should have at least two neighbors to which they can forward packets.

5.1 Experiment Methodology

We run the experiments on the same testbed and configure the network to have two access points and 48 field devices operating on four channels. As Table 3 lists, we set up six data flows with different sources, destinations, data periods, and priorities and employ graph routing as well as the rate monotonic policy for transmission scheduling. The slotframe lengths are 32, 64, 128, 256, 512 and 1,024 as suggested in the WirelessHART standard [49]. Therefore, the

⁸The implementation is provided by Li et al. and is publicly accessible [46].




Figure 9: Cracking performance on WirelessHART with different snooping periods.

combined slotframe includes 1,024 time slots in total. A maximum of three transmission attempts are scheduled for each packet. The first two attempts go through the main route and the final attempt uses the backup route in the routing graph. WirelessHART [49] specifies that S_{length} is equal to $N_{channel}$ and every channel appears one time in F . We perform two sets of experiments. First, we measure the prediction performance and cracking time when the attacker snoops different amount of time before launching the attack. Second, we vary S_{length} and investigate its impact on the cracking difficulty. We use the same device to execute the attacking program.

5.2 Cracking Performance with Different Snooping Periods

In this set of experiments, we vary the size of the training set (number of eavesdropped time slots) from 171,008 slots (167 combined slotframes) to 2,048,000 slots (2,000 combined slotframes) and use the channel activities during the next 1,707,008 slots (1667 combined slotframes) as the validation set. Figure 9(a) plots TPR, TNR, and Accuracy. As Figure 9(a) presents, TPR and Accuracy are low (80.30% and 86.09% for TPR, 90.95% and 93.38% for Accuracy) when 171,008 and 342,016 time slots are eavesdropped (167 and 334 combined slotframes as presented as the first two sets of bars). Without enough observations, the cracking program fails to derive the correct N_s . However, the N_s produced by the attacking program shares a common factor with the actual value, resulting in some correct predictions on the future channel usage (TPRs higher than 80%). Comparing Figure 9(a) and Figure 5, we observe that the attacking program achieves much higher TPR and Accuracy on WirelessHART than those on Orchestra with a similar amount of eavesdropped time slots which is insufficient for Algorithm 1 to derive the correct N_s . This is because the number of slots in each slotframe is specified to be 2^n in WirelessHART, resulting in a smaller combined slotframe and a short repetition cycle. In contrast, the default slotframe lengths of Orchestra are prime numbers, leading to a significant larger slotframe. Therefore, it is harder for the attacker to capture the repetition cycle. As Figure 9(a) shows, TPR and Accuracy experience a quick rise when 512,000 slots has been eavesdropped. TPR and Accuracy reach 96.98% and 98.54%, providing very accurate predictions on the channel usage. The training set is large enough for the attacking program to derive

the correct N_s and predict the channel usage. TPR and Accuracy then increase slowly when the training set becomes larger (TPR ranging from 96.98% to 99.78% and Accuracy ranging from 98.54% to 99.80%). We observe consistent high TNRs, since most time slots in the slotframe are not scheduled with transmissions.

Figure 9(b) shows the time consumed by the attacking program to crack the channel hopping sequence⁹. The time consumption increases linearly from 83.18s (171,008 slots) to 1037.47s (2,048,000 slots), which accords with the $O(N_r)$ time complexity of Algorithm 1.

Observation 8: The cracking difficulty depends highly on the length of combined slotframe (LCM of different slotframe lengths). The attacking program consumes more time when cracking larger combined slotframe but provides less accurate predictions.

5.3 Impact of the Length of Sequence S_{length}

In this set of experiments, we increase S_{length} from 1 to 16 and repeat the experiments. For all experimental executions, we configure the attacking program to crack after snooping the channel activities for 1,366,016 time slots (1,334 combined slotframes) and use the following 1,366,016 slots for validation. Figure 10(a) plots False Positive Rate ($FPR = FP / (FP + TN)$). As Figure 10(a) shows, FPRs are 5.88%, 4.08%, 6.49%, and 8.33% when 13–16 channels are available in the network, higher than the ones with less channels. Figure 10(b) plots the time consumed by the attacking program to crack the channel hopping sequences with different S_{length} . The time consumption increases from 768.53s with 1 channel, to 973.15s with 9 channels, and finally reaches 1465.03s with 16 channels. The results present that the cracking becomes more difficult when more channels are used in the network. This is because the data flows involve more hops when more channels are available for use [12], resulting in more transmissions in each slotframe.

Observation 9: It is difficult to crack the channel hopping sequences when a large S_{length} is used in WirelessHART networks.

5.4 Suggestions to WirelessHART

Here are some suggestions gathered from our case study:

- The specification on using slotframes with 2^n slots in WirelessHART makes the channel hopping sequences easier to

⁹The snooping period is not added into the result.




Figure 10: Cracking performance with different S_{length} in the network.

be derived. It is beneficial to use a prime number as the slotframe length.

- It is beneficial to use more channels (increasing S_{length}), which increases the cracking difficulty.

6 LESSONS LEARNED

In this section, we provide a series of insights on how to secure the TSCH channel hopping based on our analysis and case studies.

6.1 Slotframe Length

As Figure 7(a) shows, the length of combined slotframe makes a significant effect on the cracking difficulty. The larger the combined slotframe is, the more difficult the cracking is. It is beneficial to set the number of slots in different slotframes to be co-prime integers, maximizing the length of combined slotframe. As an example, the individual slotframes in the first three settings (listed in Table 2) have similar lengths, but the cracking difficulty under the third setting is much higher than the others, as Figure 7(a) shows. The default slotframe lengths in Orchestra are prime numbers, while the slotframes in WirelessHART include 2^n slots, sharing common factors with each other. Comparing Figure 9(a) against Figure 5, cracking the channel hopping sequences in WirelessHART is much easier than cracking Orchestra. Therefore, we would suggest employing multiple slotframes for different types of traffic, and even for different data flows belonging to the same type of traffic, and configuring the number of slots in each slotframe to be a prime number. For instance, if 31, 61, 127, 257, 509, and 1,021 are used as the slotframe lengths for the six data flows in WirelessHART (replacing the setup in Section 5.1), the combined slotframe includes more than 3.20×10^{13} time slots. The attacker has to spend more than 15,256 years to snoop a complete combined slotframe.

6.2 Repetition Pattern

The network device running Orchestra makes scheduling decisions based on its MAC address with a fixed offset in each slotframe, significantly making the channel usage repetition cycle detectable. As Figure 8(a) and Figure 8(b) show, the strong cyclic behavior of packet transmissions significantly reduces the cracking difficulty. As a comparison, there is no strong pattern observed in WirelessHART, which can be used by the attacker to speed up the

cracking. Therefore, we would suggest the designer of transmission scheduler avoid strong repetition pattern and randomize the transmissions. For example, Orchestra can employ pseudo-random numbers to randomize the transmission slots in the routing and time synchronization slotframes.

6.3 Channel Diversity

Using more channels not only improves the network performance but also enhances the channel hopping security. As Figure 10(a) shows, the cracking difficulty increases when using more channels. Moreover, the channel hopping sequence used by each device repeats in every $LCM(N_s, S_{length})$ time slots (Section 3.1). Hence, a large S_{length} without any common factor with N_s significantly extends the repetition cycle, making it hard for an attacker to identify the channel repetition pattern. Therefore, we would suggest using all available channels and choose N_s without having a common factor with S_{length} .

6.4 Link Setting

Orchestra specifies that all links associating with the same slotframe uses a single *ChannelOffset*. This design not only limits the network capacity but also significantly reduces the size of channel offset table which is created and maintained by the attacker. Because of the small number of *ChannelOffset* (up to three), the attacker can perform the cracking very memory-efficiently. Therefore, we would suggest using available *ChannelOffset* (S_{length}).

7 RELATED WORKS

Jamming attacks have been extensively studied in the WSN and wireless mesh network literature. Simply jamming a channel or the whole spectrum continuously can be easily detected and located by a WIPS [23, 34, 50, 51, 54]. Many countermeasures have been developed in the literature to minimize the damage. For instance, countermeasure strategies (e.g., adapting frequencies/codes to enforce spread-spectrum techniques) can be implemented in the physical layer to make jamming too complicated to carry out [5, 30, 39, 41, 54, 55]. Adjusting routing [7, 16, 25, 52], adapting transmission power [53], hopping channel [17, 52], adding redundancy [52], increasing randomness on channel access [3, 6, 44] have been shown effective against jamming attacks. Compared to continuous jamming, selective (reactive) jamming has been shown to be

much harder to detect [32, 33, 36, 39, 41, 48]. Selective jammers jam wireless channels only when their target devices are transmitting or specific packets of “high” importance are being transmitted, thus making them active for a short period of time and expending orders of magnitude with less energy. Recent studies have shown that the selective jammers can be implemented on inexpensive commercial off-the-shelf (COTS) platforms, making it a realistic threat to wireless communications [32, 33, 36, 48]. However, the existing solutions may fail to distinguish the damage caused by attacks from the normal signal fluctuations, because the transmission failures caused by the attacks happen occasionally and are buried in the normal fluctuations of low-power links. In this paper, we consider a specific kind of selective jamming, tailored to attack TSCH based wireless networks, where jamming is selectively performed against specific communication channels in specific slots. In contrast to the previous studies [40, 45] that only analyzes the ideal case and assumes S_{length} to be equal to $N_{channel}$, this paper represents the first systematic study that investigates the security vulnerability of TSCH channel hopping in IEEE 802.15.4e under realistic traffic. More important, this paper details a step-by-step attacking approach and presents two case studies with real-world TSCH implementations running on a physical testbed. The experimental results show that an attacker can reverse engineer the channel hopping sequences by silently observing the channel activities, making selective jamming attack a realistic threat to the TSCH networks.

The current industrial WSN standards (e.g., WirelessHART and ISA100) equip many security features to protect the network against such attacks as denial of service (DoS), MAC spoofing, man in the middle (MITM), and authentication and encryption cracking. There has been an increasing interest in investigating security issues in those standards. For instance, Raza et al. analyzed the potential attacks to WSNs and proposed a series of security enhancement mechanisms for WirelessHart [35]. Alcazar et al. presented a series of vulnerabilities on the routing of WirelessHart and ISA100 and proposed some countermeasures [2]. Pietro et al. developed a distributed self-healing protocol to enhance the intrusion-resilience [27]. Chakrabarty et al. proposed a software defined networking (SDN) architecture to mitigate traffic analysis and data gathering attacks [4]. IEEE 802.15.4 also provides security features such as data confidentiality, data authenticity, and replay protection of MAC frames. For example, the standard includes a security suite based on the Advanced Encryption Standard (AES) 128 bits symmetric-key cryptography and supports three different security modes: encryption only (CTR), authentication only (CBC_MAC), and both encryption and authentication (CCM). Unfortunately, the above security features cannot prevent an attacker from cracking the channel hopping sequences by silently observing the channel activities. Our work is therefore orthogonal and complementary.

8 CONCLUSIONS

Industrial IoT, underlying the Fourth Industrial Revolution (or Industry 4.0), promises to be one of the largest potential economic effects of IoT. IEEE 802.15.4-based WSNs technology is appealing for use in industrial IoT applications because it does not require wired infrastructure and can be manufactured inexpensively. Battery-powered wireless modules easily and inexpensively retrofit existing

sensors and actuators in industrial facilities without running cables for communication and power. To address the stringent real-time and reliability requirements posed by industrial IoT applications, WSNs made a set of unique design choices such as employing TSCH that distinguish themselves from traditional WSNs that require only best effort services. The function-based channel hopping in TSCH simplifies the network operations at the cost of security. Our study shows that an attacker can reverse engineer the channel hopping sequences by silently observing the channel activities and put the network in danger of selective jamming attacks, bringing severe threats into WSNs. This paper represents the first systematic study that investigates the security vulnerability of TSCH channel hopping in IEEE 802.15.4e under realistic traffic, demonstrates the cracking process, and presents two case studies using publicly accessible TSCH implementations. Finally, this paper provides a set of insights gathered from our analysis and case studies to secure the TSCH channel hopping by increasing the cracking difficulty.

For future work, we plan to study how an attacker identifies the key messages delivering in the network and launches the selective jamming attacks. Once we fully understand the problem from an attacker's point of view, then we can develop strategies to efficiently detect the selective jamming attacks and defense solutions in future.

ACKNOWLEDGMENT

The authors thank the anonymous reviewers and the shepherd Brad Campbell for their insightful comments. This work was supported by the NSF through grant CRII-1657275 (NeTS).

REFERENCES

- [1] 802.15.4e. 2013. IEEE802.15.4e WPAN Task Group. Retrieved September 28, 2018 from <http://www.ieee802.org/15/pub/TG4e.html>
- [2] Cristina Alcaraz and Javier Lopez. 2010. A Security Analysis for Wireless Sensor Mesh Networks in Highly Critical Systems. *IEEE Transactions on Systems, Man, and Cybernetics* 40, 4 (July 2010), 419–428. <https://doi.org/10.1109/TSMCC.2010.2045373>
- [3] Farhana Ashraf, Yih-Chun Hu, and Robin H. Kravets. 2012. Bankrupting the jammer in WSN. In *Proceedings of the 2012 IEEE 9th International Conference on Mobile Ad-Hoc and Sensor Systems (MASS) (MASS '12)*. IEEE, Washington, DC, USA, 317–325. <https://doi.org/10.1109/MASS.2012.6502531>
- [4] Shaibal Chakrabarty, Daniel W. Engels, and Selina Thathapudi. 2015. Black SDN for the Internet of Things. In *Proceedings of the 2015 IEEE 12th International Conference on Mobile Ad-Hoc and Sensor Systems (MASS) (MASS '15)*. IEEE, Washington, DC, USA, 190–198. <https://doi.org/10.1109/MASS.2015.100>
- [5] Jerry T. Chiang and Yih-Chun Hu. 2011. Cross-Layer Jamming Detection and Mitigation in Wireless Broadcast Networks. *IEEE/ACM Transactions on Networking* 19, 1 (Feb. 2011), 286–298. <https://doi.org/10.1109/TNET.2010.2068576>
- [6] Roberta Daidone, Gianluca Dini, and Marco Tiloca. 2014. A Solution to the GTS-based Selective Jamming Attack on IEEE 802.15.4 Networks. *Wireless Networks* 20, 5 (July 2014), 1223–1235. <https://doi.org/10.1007/s11276-013-0673-y>
- [7] Jing Deng, Richard Han, and Shivakant Mishra. 2003. A Performance Evaluation of Intrusion-Tolerant Routing in Wireless Sensor Networks. In *Proceedings of the 2nd international conference on Information processing in sensor networks (IPSN'03)*. Springer-Verlag Berlin, Heidelberg, 349–364. https://doi.org/10.1007/3-540-36978-3_23
- [8] Adam Dunkels. 2002. Contiki: The Open Source OS for the Internet of Things. Retrieved September 28, 2018 from <http://www.contiki-os.org/>
- [9] Simon Duquennoy, Atis Elstsz, Beshr Al Nahasx, and George Oikonomou. 2017. TSCH and 6TiSCH for Contiki: Challenges, Design and Evaluation. In *2017 13th International Conference on Distributed Computing in Sensor Systems (DCOSS)*. IEEE, Piscataway, NJ, USA. <https://doi.org/10.1109/DCOSS.2017.29>
- [10] Simon Duquennoy, Beshr Al Nahas, and Atis Elstsz. 2018. 6TiSCH Implementation. Retrieved September 29, 2018 from <https://github.com/contiki-ng/contiki-ng/wiki/Documentation:-TSCH-and-6TiSCH>
- [11] Simon Duquennoy, Beshr Al Nahas, Olaf Landsiedel, and Thomas Watteyne. 2015. Orchestra: Robust Mesh Networks Through Autonomously Scheduled TSCH. In *Proceedings of the 13th ACM Conference on Embedded Networked Sensor Systems*

- (SenSys '15). ACM, New York, NY, USA, 337–350. <https://doi.org/10.1145/2809695.2809714>
- [12] Dolvara Gunatilaka, Mo Sha, and Chenyang Lu. 2017. Impacts of Channel Selection on Industrial Wireless Sensor-Actuator Networks. In *IEEE INFOCOM 2017 - IEEE Conference on Computer Communications*. IEEE, Piscataway, NJ, USA. <https://doi.org/10.1109/INFOCOM.2017.8057049>
 - [13] HART. 2019. HART Communication Protocol and Foundation (Now the FieldComm Group). <https://fieldcommgroup.org/>
 - [14] IETF. 2019. 6TiSCH: IPv6 over the TSCH mode of IEEE 802.15.4e. Retrieved September 28, 2018 from <https://datatracker.ietf.org/wg/6tisch/documents/>
 - [15] ISA100. 2018. ISA100. <http://www.isa100wci.org/>
 - [16] Chris Karlof, Naveen Sastry, and David Wagner. 2004. TinySec: a Link Layer Security Architecture for Wireless Sensor Networks. In *Proceedings of the 2nd international conference on Embedded networked sensor systems (SenSys '04)*. ACM, New York, NY, USA, 162–175. <https://doi.org/10.1145/1031495.1031515>
 - [17] Loukas Lazos, Sisi Liu, and Marwan Krunz. 2009. Mitigating Control-channel Jamming Attacks in Multi-channel Ad Hoc Networks. In *Proceedings of the second ACM conference on Wireless network security (WiSec '09)*. ACM, New York, NY, USA, 169–180. <https://doi.org/10.1145/1514274.1514299>
 - [18] Philip Levis. 2013. TinyOS Documentation Wiki. Retrieved September 28, 2018 from http://tinyos.stanford.edu/tinyos-wiki/index.php/TinyOS_Documentation_Wiki
 - [19] B. Li, Y. Ma, T. Westenbroek, C. Wu, H. Gonzalez, and C. Lu. 2016. Wireless Routing and Control: a Cyber-Physical Case Study. In *Proceedings of the 7th International Conference on Cyber-Physical Systems (ICCPs '16)*. IEEE, Piscataway, NJ, USA. <https://doi.org/10.1109/ICCPs.2016.7479131>
 - [20] B. Li, L. Nie, C. Wu, H. Gonzalez, and C. Lu. 2015. Incorporating Emergency Alarms in Reliable Wireless Process Control. In *Proceedings of the ACM/IEEE Sixth International Conference on Cyber-Physical Systems (ICCPs '15)*. ACM, New York, NY, USA, 218–227. <https://doi.org/10.1145/2735960.2735983>
 - [21] B. Li, Z. Sun, K. Mechitov, G. Hackmann, C. Lu, S. Dyke, G. Agha, and B. Spencer. 2013. Realistic Case Studies of Wireless Structural Control. In *Proceedings of the ACM/IEEE 4th International Conference on Cyber-Physical Systems (ICCPs '13)*. ACM, New York, NY, USA, 179–188. <https://doi.org/10.1145/2502524.2502549>
 - [22] Chenyang Lu, Abusayed Saifullah, Bo Li, Mo Sha, Humberto Gonzalez, Dolvara Gunatilaka, Chengjie Wu, Lanshun Nie, and Yixin Chen. 2016. Real-Time Wireless Sensor-Actuator Networks for Industrial Cyber-Physical Systems. *Proceedings of the IEEE, Special Issue on Industrial Cyber Physical Systems* 104, 5 (May 2016), 1013–1024. <https://doi.org/10.1109/PROC.2015.2497161>
 - [23] Zhuo Lu, Wenye Wang, and Cliff Wang. 2014. Modeling, Evaluation and Detection of Jamming Attacks in Time-Critical Wireless Applications. *IEEE Transactions on Mobile Computing* 13, 8 (Aug. 2014), 1746–1759. <https://doi.org/10.1109/TMC.2013.146>
 - [24] James Manyika, Michael Chui, Jacques Bughin, Richard Dobbs, Peter Bisson, and Alex Marrs. 2013. Disruptive Technologies: Advances that will Transform Life, Business, and the Global Economy. <http://www.mckinsey.com/business-functions/digital-mckinsey/our-insights/disruptive-technologies>
 - [25] Hossein Mustafa, Xin Zhang, Zhenhua Liu, Wenyan Xu, and Adrian Perrig. 2012. Jamming-Resilient Multipath Routing. *IEEE Transactions on Dependable and Secure Computing* 9, 6 (Nov. 2012), 852–864. <https://doi.org/10.1109/TDSC.2012.69>
 - [26] Raspberry Pi. 2019. Raspberry Pi. <https://www.raspberrypi.org/>
 - [27] Roberto Di Pietro, Gabriele Oliveri, Claudio Soriente, and Gene Tsudik. 2010. Intrusion-Resilience in Mobile Unattended WSNs. In *Proceedings of the 29th conference on Information communications (INFOCOM'10)*. IEEE, Piscataway, NJ, USA, 2303–2311. <https://doi.org/10.1109/INFOCOM.2010.5462056>
 - [28] Kris Pister. 2010. Smart Dust: Autonomous Sensing and Communication in a Cubic Millimeter. <https://people.eecs.berkeley.edu/~pister/SmartDust/>
 - [29] Kristofer S. J. Pister and Lance Doherty. 2008. TSMP: Time Synchronized Mesh Protocol. In *IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*. IEEE, Piscataway, NJ, USA, 391–398.
 - [30] Christina Popper, Mario Strasser, and Srdjan Capkun. 2010. Anti-jamming Broadcast Communication Using Uncoordinated Spread Spectrum Techniques. *IEEE Journal on Selected Areas in Communications* 28, 5 (June 2010), 703–715. <https://doi.org/10.1109/JSAC.2010.1006088>
 - [31] Michael E Porter and James E Hoppmann. 2014. How Smart, Connected Products are Transforming Competition. *Harvard Business Review* 92, 11 (Nov. 2014), 64–88. <https://hbr.org/2014/11/how-smart-connected-products-are-transforming-competition>
 - [32] Alejandro Proaño and Loukas Lazos. 2010. Selective Jamming Attacks in Wireless Networks. In *2010 IEEE International Conference on Communications*. IEEE, Piscataway, NJ, USA, 1–6. <https://doi.org/10.1109/ICC.2010.5502322>
 - [33] Alejandro Proaño and Loukas Lazos. 2012. Packet-hiding Methods for Preventing Selective Jamming Attacks. *IEEE Transactions on Dependable and Secure Computing* 9, 1 (Jan. 2012), 101–114. <https://doi.org/10.1109/TDSC.2011.41>
 - [34] David R. Raymond and Scott F. Midkiff. 2008. Denial-of-Service in Wireless Sensor Networks: Attacks and Defenses. *IEEE Pervasive Computing* 7, 1 (Jan. 2008), 74–81. <https://doi.org/10.1109/MPRV.2008.6>
 - [35] Shahid Raza, Adriaan Slabbert, Thiemo Voigt, and Krister LandernÅds. 2009. Security considerations for the WirelessHART protocol. In *Proceedings of the 14th IEEE international conference on Emerging technologies and factory automation (ETFA'09)*. IEEE, Piscataway, NJ, USA, 242–249. <https://doi.org/10.1109/ETFA.2009.5347043>
 - [36] AndrÁla Richa, Christian Scheideler, Stefan Schmid, and Jin Zhang. 2013. An Efficient and Fair MAC Protocol Robust to Reactive Interference. *IEEE/ACM Transactions on Networking* 21, 3 (June 2013), 760–771. <https://doi.org/10.1109/TNET.2012.2210241>
 - [37] RPL. 2012. RFC 6550: RPL: IPv6 Routing Protocol for Low-Power and Lossy Networks. Retrieved September 28, 2018 from <https://tools.ietf.org/html/rfc6550>
 - [38] Mo Sha. 2016. Testbed at the State University of New York at Binghamton. Retrieved September 28, 2018 from <http://www.cs.binghamton.edu/%7Emsha/testbed>
 - [39] Michael Spuhler, Domenico Giustiniano, Vincent Lenders, Matthias Wilhelm, and Jens B. Schmitt. 2014. Detection of Reactive Jamming in DSSS-based Wireless Communications. *IEEE Transactions on Wireless Communications* 13, 3 (March 2014), 1593–1603. <https://doi.org/10.1109/TWC.2013.013014.131037>
 - [40] Spase Stojanovski and Andrea Kulakov. 2014. Efficient Attacks in Industrial Wireless Sensor Networks. In *International Conference on ICT Innovations (ICT)*. Springer, Cham, Cham, Switzerland, 289–298. https://doi.org/10.1007/978-3-319-09879-1_29
 - [41] Mario Strasser, Boris Danev, and Srdjan Ånapkun. 2010. Detection of Reactive Jamming in Sensor Networks. *ACM Transactions on Sensor Networks* 7, 2 (Aug. 2010), 16:1–16:29. <https://doi.org/10.1145/1824766.1824772>
 - [42] TelosB. 2013. TelosB Datasheet provided by MEMSIC. Retrieved October 2, 2018 from http://www.memsic.com/userfiles/files/Datasheets/WSN/telosb_datasheet.pdf
 - [43] Adam Thierier and Andrea Castillo. 2015. Projecting the Growth and Economic Impact of the Internet of Things. <https://www.mercatus.org/publication/projecting-growth-and-economic-impact-internet-things>
 - [44] Marco Tiloca, Domenico De Guglielmo, Gianluca Dini, Giuseppe Anastasi, and Sajal K. Das. 2017. JAMMY: a Distributed and Self-Adaptive Solution against Selective Jamming Attack in TDMA WSNs. *IEEE Transactions on Dependable and Secure Computing* 14, 4 (July 2017), 392–405. <https://doi.org/10.1109/TDSC.2015.2467391>
 - [45] Marco Tiloca, Domenico De Guglielmo, Gianluca Dini, Giuseppe Anastasi, and Sajal K. Das. 2019. DISH: Distributed SHuffling against Selective Jamming Attack in IEEE 802.15.4e TSCH Networks. *ACM Transactions on Sensor Networks (TOSN)* 15, 1 (Feb. 2019). <https://doi.org/10.1145/3241052>
 - [46] Wireless Cyber-Physical Simulator (WCPS). 2018. Wireless Cyber-Physical Simulator (WCPS). Retrieved October 2, 2018 from http://wsn.cse.wustl.edu/index.php/WCPS:_Wireless_Cyber-Physical_Simulator
 - [47] Wi-Spy. 2018. Wi-Spy USB Spectrum Analyzer. <http://www.wi-spy.co.uk/index.php/products>
 - [48] Matthias Wilhelm, Ivan Martinovic, Jens B. Schmitt, and Vincent Lenders. 2011. Short Paper: Reactive Jamming in Wireless Networks How Realistic is the Threat?. In *Proceedings of the fourth ACM conference on Wireless network security (WiSec '11)*. ACM, New York, NY, USA, 47–52. <https://doi.org/10.1145/1998412.1998422>
 - [49] WirelessHART. 2019. WirelessHART. <https://fieldcommgroup.org/technologies/hart/hart-technology>
 - [50] A.D. Wood, J.A. Stankovic, and S.H. Son. 2003. JAM: a Jammed-area Mapping Service for Sensor Networks. In *Proceedings of the 24th IEEE International Real-Time Systems Symposium (RTSS '03)*. IEEE, Washington, DC, USA, 286–297. <https://doi.org/10.1109/REAL.2003.1253275>
 - [51] Anthony D. Wood and John A. Stankovic. 2002. Denial of Service in Sensor Networks. *Computer* 35, 10 (Oct. 2002), 54–62. <https://doi.org/10.1109/MC.2002.1039518>
 - [52] Anthony D. Wood, John A. Stankovic, and Gang Zhou. 2007. DEEJAM: Defeating Energy-Efficient Jamming in IEEE 802.15.4-based Wireless Networks. In *2007 4th Annual IEEE Communications Society Conference on Sensor, Mesh and Ad Hoc Communications and Networks*. IEEE, Piscataway, NJ, USA, 60–69. <https://doi.org/10.1109/SAHCN.2007.4292818>
 - [53] Wenyan Xu, Ke Ma, W. Trappe, and Yanyong Zhang. 2006. Jamming Sensor Networks: Attack and Defense Strategies. *IEEE Network* 20, 3 (May 2006), 41–47. <https://doi.org/10.1109/MNET.2006.1637931>
 - [54] Wenyan Xu, Wade Trappe, Yanyong Zhang, and Timothy Wood. 2005. The Feasibility of Launching and Detecting Jamming Attacks in Wireless Networks. In *Proceedings of the 6th ACM international symposium on Mobile ad hoc networking and computing (MobiHoc '05)*. ACM, New York, NY, USA, 46–57. <https://doi.org/10.1145/1062689.1062697>
 - [55] Wenyan Xu, Timothy Wood, Wade Trappe, and Yanyong Zhang. 2004. Channel Surfing and Spatial Retreats: Defenses Against Wireless Denial of Service. In *Proceedings of the 3rd ACM workshop on Wireless security (WiSe '04)*. ACM, New York, NY, USA, 80–89. <https://doi.org/10.1145/1023646.1023661>