

Exploring Potential 6LoWPAN Traffic Side Channels

Yan Yan Elisabeth Oswald Theo Tryfonas

April 12, 2017

University of Bristol,
{y.yan elisabeth.oswald theo.tryfonas}@bristol.ac.uk

Abstract

The Internet of Things (IoT) has become a reality: small connected devices feature in everyday objects including childrens' toys, TVs, fridges, heating control units, etc. Supply chains feature sensors throughout, and significant investments go into researching next-generation healthcare, where sensors monitor wellbeing. A future in which sensors and other (small) devices interact to create sophisticated applications seems just around the corner. All of these applications have a fundamental need for security and privacy and thus cryptography is deployed as part of an attempt to secure them. In this paper we explore a particular type of flaw, namely side channel information, on the protocol level that can exist despite the use of cryptography. Our research investigates the potential for utilising packet length and timing information (both are easily obtained) to extract interesting information from a system. We find that using these side channels we can distinguish between devices, different programs running on the same device including which sensor is accessed. We also find it is possible to distinguish between different types of ICMP messages despite the use of encryption. Based on our findings, we provide a set of recommendations to efficiently mitigate these side channels in the IoT context.

1 Introduction

The expression 'Internet of Things' (IoT) can refer to a multitude of objects and protocols, which share that they have been purposefully designed for resource constraint environments. Whereas the typical TCP/IP network stack produces considerable overhead to achieve quality of service for applications that are based on it, the nature of many IoT 'things' is such that a full implementation of it would not be practical. Often 'things' are sensor, which are devices that have to function on little resources (most importantly power). Thus a whole

host of new networking protocols have been developed over the years to cater for such resource constrained devices: 6LoWPAN is the ‘tiny’ version of IPv6, UDP tends to be used instead of TCP/IP, DTLS can be used for end-to-end security or one can directly invoke 802.15.4 security which is part of 6LoWPAN, and finally CoAP(s) is the replacement for HTTP(s). Thus there are two options (802.15.4, and DTLS) to secure communications between the ‘things’ and a server/gateway.

Implementing cryptography correctly and securely has proven to be a massive challenge as evidenced by the multitude of implementation attacks over the years. Triggered off by research that showed how to utilise additional information via timing and power side channels [23], many different flavours of side channel attacks were discovered over the last decade. Many attacks use physical information (such as low level execution timings or power consumption) to recover secret keys, but many other attacks use protocol level information (such as packet lengths, types of packets or protocol messages) to recover information about plaintexts, devices in the network, or the network itself. There exists a considerable body of work in the context of conventional, i.e. HTTPs over TCP/IP network, but the applicability of (some) of these attacks in the context of a typical IoT protocol stack is lacking. This is the gap that we would like to address with this work.

This paper is structured as follows: after reviewing some relevant attack paths for HTTPs over TCP/IP in the following subsection, we provide a brief introduction to the necessary protocol and network features in Section 2. We discuss the impact of packet length leakage in Section 3, followed by an analysis of the response time leakage in Section 4. We summarise our work in Section 5.

1.1 Related Work

Traffic Analysis is well studied in the context of encrypted Internet traffic, especially for web applications based on HTTPs and TCP/IP. The landmark study by Chen et al. [8] discussed different side channel attacks against web applications and [41] studied the practicability of an attack specifically targeted Google and Bing search boxes. Later work by Mather and Oswald [30] proposed the use of Mutual Information to pinpoint the potential leakage points in web traffic. For non-HTTPs applications, the papers [11], [50] and [9] described attacks against encrypted text, voice and video traffic respectively. Machine learning is widely used to analyse the traffic, and behaviours of different classifiers are studied by [19] and [15]. Based on all these published works we can conclude that two features, the packet length and response time, are the most exploited ones among all attacks. Different countermeasures were studied by [51], [29] and [16].

Reflecting on IoT applications, we stipulate that most of these attacks may still be applicable, as we intend to demonstrate in this paper. Considering the future vision that IoT devices could be indeed connected to the Internet with even more sensitive data flowing over different networks, the task of designing secure IoT applications becomes increasingly challenging.

With regard to the aspect of protocol design, the recent paper [36] summarised some known flaws of 6LoWPAN, including its susceptibility to the Fragmentation Attack [21], Sinkhole Attack [26], Hello Flood Attack [44], Wormhole Attack [20] and Blackhole Attack [47]. In addition, [40] reported certain problematic designs in 802.15.4 security [18]. However we do not discuss further these particular design flaws as they touch on a different aspect of the security issues in 6LoWPAN compared to what we address in this paper.

2 A Typical IoT Protocol Stack

There are large number of protocols, which have been proposed for different IoT applications adapting to various requirements. For example, some smart houses simply use WiFi for connectivity and VANETs¹ may adopt DSRC [22].

In this paper we focus on 6LoWPAN [32] which is based on the 802.15.4 [18] standard. These standards are designated for constrained environments such as Wireless Sensor Networks, but other competing standards exist at different layers. Bluetooth Low Energy(BLE) [1] is a strong competitor to 802.15.4 as well as the LiFi [38] technology. Zigbee [6] was originally intended as a collective protocol over 802.15.4 but it has been recently adapted to IP-based network in ZigbeeIP [7]. The RIME stack [13] proposes a set of non-layered primitives over 802.15.4 but it is likely to be phased-out due to the lack of interoperability with the TCP/IP protocol stack.

6LoWPAN thus is the most popular standard for low power networks, and thus it is supported by several competing IoT Operating Systems, including Contiki OS [33], OpenWSN [2], FreeRTOS [3] and the recent RIOT [4]. We chose Contiki OS for our experiments because it is easy to customise.

2.1 Our experimental network

Our experimental network is constructed using two different devices. These are a TelosB [34] and a CC2538 [35]. The TelosB is a low cost sensor powered by an MSP430 with an AES co-processor. It represents typical low-end devices. The CC2538 is the high end device powered by an ARM Cortex-M3 with multiple cryptographic processors including AES, RSA, SHA-2 and ECC, suggesting that it is suitable to develop secure applications.

Both devices are supported by the Contiki OS. We adopted the default settings of the Contiki OS, except for enabling 802.15.4 security [18] for some experiments. Note that the Contiki MAC [12] is chosen by default over TSCH [45]. For Layer 4 [17] and above protocols, we went with the widely accepted combination of CoAP [43], and DTLS [39](optional) over UDP [37]². Table 1 summarises our choice of protocol stack.

¹Vehicular ad hoc networks

²CoAPs is equivalent to CoAP over DTLS.

Physical	802.15.4
Link	
Network	6LoWPAN
Transmission	UDP
	DTLS*
Application	CoAP / CoAPs*

Table 1: Protocol stack for our experiments(* is optional)

2.1.1 802.15.4 and DTLS

In our setting, there are two standards available for packet encryption, namely 802.15.4 security [18] and DTLS [39]. 802.15.4 security is provided by the noncoresec [25] API, which implements 802.15.4 authenticated encryption with AES-128 CCM* [14] using a hard-coded key shared by the whole 6LoWPAN network. We chose tinyDTLS [5] as library for the DTLS protocols, because it provides a minimum DTLS implementation that supports two ciphersuites which are TLS_PSK_WITH_AES_128_CCM.8 [31] and TLS_ECDHE_ECDSA_WITH_AES_128_CCM.8 [31] respectively. Evidently, they both utilise AES-128 CCM* as the packet encryption method.

3 Exploiting Packet Length Information

As our brief survey of traffic analysis via exploiting packet lengths showed in Section 1.1, the packet length has proven to be a powerful side channel for the classical Internet protocols. It is worth noting that this side channel is ‘noisy’ in the classical Internet setting: websites or web applications in this setting typically feature advertisements, which impact on packet lengths; TCP/IP allows to fragment packets and then reassembles them, a feature which is not presented in UDP. Thus, due to the nature of UDP exploiting the packet length as side channel should be easier in the IoT setting.

Clearly then, any web application style implementations involving an IoT device will thus be extremely vulnerable to attacks such as [8]. In the absence of this scenario for state-of-the art IoT applications, it still sends a cautionary warning to developers: binary responses (e.g. ‘yes’ vs. ‘no’, or ‘on’ vs. ‘off’) must always be coded via a binary variable and not via strings because these will have different lengths, which are directly visible via the packet length.

In the remainder of this section we will highlight further problems that arise if packet lengths leak information.

3.1 Distinguishing ICMP Messages

The Internet Control Message Protocol(ICMP) [10] performs the management tasks in a network, such as link establishment and routing information exchange. As explained before we utilise the open source system Contiki, which supports

a (sub)set of the ICMP standard (we list the supported ICMP messages further below). Many ICMP messages are ideal for network discovery and exploration, although the purpose of ICMP is to send error messages to the source IP address if standard IP packets fail to be transmitted correctly.

- **DAG Information Object (DIO)**
DIO contains the 6LoWPAN global information. It could be periodically broadcasted for network maintenance, or unicast to a new joining node as a reply to DIS (see below).
- **DAG Information Solicitation (DIS)**
DIS is sent by a newly started node to probe any existing 6LoWPANs. A DIO would be replied if the DIS is received by any neighbour nodes.
- **Destination Advertisement Object (DAO)**
DAO is sent by a child node to its precedents (The 6LoWPAN DODAG topology is defined in [49]) to propagate its routing information.
- **Neighbour Solicitation (NS) and Neighbour Advertisement (NA)**
NS and NA are the ARP replacement in IPv6, where NS queries a translation and NA answers one. In addition, they are also used for local link validity check.
- **Echo Request and Echo Response (PING)**
Echo Request and Echo Response are well known as the PING packets. They are mostly used for diagnostic purposes, such as connectivity test or Round Trip Time (RTT) estimation. Echo Request may contain arbitrary user defined data and Echo Response simply echoes its corresponding request.

Generally, ICMP messages can be protected by either using the secure ICMP messages as described in [10], or relying on the lower layer encryption provided by 802.15.4. Contiki OS does not have the former implemented, hence 802.15.4 security is the only option currently. We simulated a 6LoWPAN network with 802.15.4 security enabled (with strongest encryption and authentication). We configured the nodes to also generate random UDP packets. Despite the fact that all ICMP messages were encrypted, our experiments show that several ICMP messages can be identified by their packet size and MAC destination. Table 2 summarises the packet features. The value x denotes the size of user defined data in bytes.

Among the unicast packets, PING and UDP have at least 101 and 108 bytes³. Therefore, DAO can be uniquely identified as the shorter unicast packet of 97 bytes. For the same reason NA and unicast NS can also be distinguished from other packets by filtering packets of 87 bytes. Considering that NA is sent as a response to NS according to the protocol, one can always identify the first being NS and second being NA.

³PING can be sent without user defined data and UDP packets requires at least 1 byte.

	Packet Size (bytes)	Type of MAC Destination
DIS	85	broadcast
DIO	118/123	broadcast/unicast
DAO	97	unicast
NS	87	broadcast/unicast
NA	87	unicast
PING	$101 + x$	unicast
UDP Multicast	$85 + x$	broadcast
UDP Unicast	$107 + x$	unicast

Table 2: 6LoWPAN Packet Features

Similarly, unicast DIO can be identified as the 123 bytes packet followed by DIS, where the later has a unique 85 byte size. However, there is a potential of false positive induced by PING or UDP packets with user defined data crafted to have the same packet length⁴. PING could be recognised by its pair-wised appearance, as the response would have nearly the same meta data as the original request, except the exchanged source and destination. For broadcast packets, DIS can be easily identified by its unique 85 bytes packet size. Others like broadcast NS can be identified by the followed characteristic NA response; and packets of 118 bytes those are periodically broadcasted are likely to be DIOs.

In summary, among all the packets, DAO, NA, NS, DIS can be identified with certainty. DIO and PING cannot be certainly identified but they both have significant characters. Notice that the above contained all ICMPv6 messages supported by Contiki; therefore UDP packets can be reversely filtered, although in some cases they get mixed with DIO and PING.

Although leakage in ICMP messages does not directly lead to any breach of application data, it would still be harmful by providing the adversary with information about the state of the network, including which nodes recently joined etc. Specifically DAO is always sent from a child to its parent and can be uniquely identified; therefore together with MAC addresses the adversary may exploit it to draw a graph that shows the parental relations in the network. In addition, these information can also be exploited by attacks as in [28].

3.2 Distinguishing Different Devices

In the classical Internet world, ICMP has been well known for its use for OS fingerprinting [46]. In the case of the IoT, this could be possible as well (as different OS support different subsets of ICMP), however an additional attack vector exists. This is because different IoT devices have different hardware limitations or drivers. We noticed that our TelosB [34] discards all packets exceeding 127 bytes⁵ whereas our CC2538 handles packets even up to 160 bytes.

⁴22 bytes for PING and 16 bytes for UDP.

⁵MTU specified by 802.15.4 standard.

Therefore an adversary can immediately rule out TelosB whenever a packet larger than 127 bytes processed by the target.

4 Exploiting Response Time Information

The response time is another major feature that has been previously exploited in Internet traffic analysis attacks. Like in the case of exploiting packet lengths, we would expect that the same attacks (as in the classical Internet setting) can be applied to 6LoWPAN traffic. Indeed, like in the previous section, we would expect that they will work even better because the accuracy of timing measurements can be greatly improved for 6LoWPAN traffic: this is because there are fewer noise sources in the traffic, the devices are physically close to each other and uses RF to communicate, the adversary can remove the RTT noises by measure the packets on the server side, and the performance of the constrained devices is low and hence gives a better resolution of the execution time.

4.1 Distinguishing Different Sensors

The first application of timing analysis that we describe is to distinguish between different sensors that are accessed on a device. For this purpose we set up an experiment on a CC2538, which has three on-board sensors: Vdd, temperature, and an Ambient Light Sensor (short ALS). We access these via CoAP [43], which is a protocol designed for constrained devices that provides an universal interface for accessing resources. CoAPs is the secure version which stands for CoAP with DTLS.

Due to the different physical characteristics of the sensors, there could be a variance of time that is required for reading the measurements. We investigated whether such variances could be observed through the packet response latency. If this was the case, then an adversary could learn the nature/purpose of sensors on a network by observing their response time.

We thus set up an experiment on CC2538, using all three sensors from “cc2538-demo”. We used CoAP from the “er-rest-example” in the Contiki OS source code, as there is no CoAPs implementation available. Although DTLS processing would definitely have an impact on the response latency, we argue that such impact would be independent to the sensors being accessed; hence similar result can be equally expected for CoAPs. We carefully controlled other factors, including URIs, data representation and code flow, to be uniform for all three sensors in order to guarantee a controlled environment.

Table 3 summarises the result. It shows that ALS takes about 2ms longer and hence can be easily distinguished. Vdd and temperature have much more strongly overlapping distributions, and thus are more difficult to distinguish. Nevertheless these results confirm our hypothesis: different sensors have different latencies and these leak through the response time. An adversary who is interested in finding out information about devices on a network might thus

	Average (ms)	Range(ms)
Vdd	9.622	[9.388, 10.318]
Temperature	9.835	[9.525, 10.318]
ALS	11.651	[11.338, 12.031]

Table 3: CoAP Response Latency for Sensor Readings on CC2538

	CC2538	TelosB
Average(ms)	9.56	17.03
Range(ms)	[9.16, 10.06]	[16.49, 17.68]

Table 4: PING Response Latency

be able to match the (known) behaviour of ‘interesting’ sensors to what they observe on the network. We remark that this could be useful even in the setting where the sensors transmit their data unencrypted: after all they might return only some reading without a unit of measurement; thus seeing their return data might not as such reveal their nature.

4.2 Distinguishing Different Devices

As we observed before, different devices have different underlying hardware and thus different computational power. This implies that there could be the potential that different devices take different amounts of time to process the same message. Because ICMP messages are standardised, they are particularly suitable for this purpose. Among the different ICMP messages, PING is especially ideal for two reasons:

1. It is mandatory in the ICMP standard.
2. It only swaps the source and destination address of the packet; thus minimises different code path in protocol processing.

Table 4 shows the PING response latency on CC2538 and TelosB. The result confirms that these devices can be distinguished by PING response latency.

4.3 Distinguishing Programs

We remarked before that the functionality of a sensor is potentially valuable information. For instance some sensors might be predominantly passive, e.g. they might read the temperature and report it back periodically, whereas some sensors might control something upon receiving commands. Thus knowing the functionality enables an adversary to make (more) sense of the observed traffic in the network. This could be done if a ‘fingerprint’ could be produced for different programs. From an adversary’s perspective a positive result would imply that they could ‘fingerprint’ products which are on the market and thus use this information to infer what program is running on a target device.

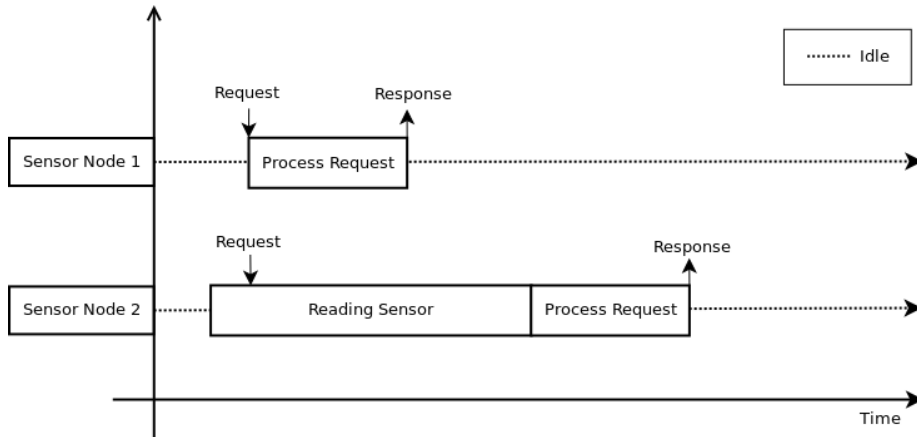


Figure 1: Variations in Response Time

To illustrate why this might work, we now look at Figure 1. It illustrates two sensors receiving the same service request. In our example, at the time of receiving the request, Sensor Node 1 was idle and hence responded immediately, whilst Sensor Node 2 postponed the request for reading a sensor. Clearly, the response time on Sensor Node 2 would appear longer than that of Sensor Node 1.

In real life, most sensors are programmed in a loop; therefore the same code fragments are repeated through the life time of a sensor. Each code fragment takes different time to execute and hence the response times vary. This behaviour could be statistically analysed and the resulting distribution could be stored as a ‘fingerprint’.

For this fingerprinting scenario, we must assume the adversary has the pre-knowledge of potential programs and can fingerprint them (or that they have access to a database that contains this information). To identify an unknown program running on target sensor, the adversary collects a new fingerprint and then matches it to available fingerprints. Clearly, to effectively launch the attack, the adversary needs to be able to send the request to a targeted sensor (requests with short predictable processing time are preferable as they induce less noise).

In practice, the request can be instantiated by several messages defined in the sensor network protocols. PING is exceptionally ideal as it is mandatory in the ICMP standard [27] and has only negligible computation. Other options but not excluded are Heartbeat in DTLS [42], Reset in CoAP [43], etc.

4.3.1 Extracting Fingerprints

We explored the feasibility of fingerprinting programs on an CC2538 running Contiki OS by using the PING command.

Figure 2 shows an example of captured packets. Contiki MAC [12] sends

duplicated PING requests. The response time, which refers to PING Response Interval, PRI, is defined to be the time between a PING response and its last paired request. The highlighted Packets 205 and 203 shows such an example.

No.	Time	Source	Destination	Protocol	Length	Info
198	4.667274	aaaa::1	aaaa::212...	ICMPv6	80	Echo (ping) request id=0x64c4, seq=16,
199	4.670572	aaaa::1	aaaa::212...	ICMPv6	80	Echo (ping) request id=0x64c4, seq=16,
200	4.674060	aaaa::1	aaaa::212...	ICMPv6	80	Echo (ping) request id=0x64c4, seq=16,
201	4.677277	aaaa::1	aaaa::212...	ICMPv6	80	Echo (ping) request id=0x64c4, seq=16,
202	4.680601	aaaa::1	aaaa::212...	ICMPv6	80	Echo (ping) request id=0x64c4, seq=16,
203	4.684369	aaaa::1	aaaa::212...	ICMPv6	80	Echo (ping) request id=0x64c4, seq=16,
204	4.684724			IEEE 8...	5	Ack
205	4.701468	aaaa::...	aaaa::1	ICMPv6	80	Echo (ping) reply id=0x64c4, seq=16, h
206	4.701962			IEEE 8...	5	ACK
207	5.632173	aaaa::1	aaaa::212...	ICMPv6	80	Echo (ping) request id=0x64c4, seq=17,
208	5.635445	aaaa::1	aaaa::212...	ICMPv6	80	Echo (ping) request id=0x64c4, seq=17,

Figure 2: Example PRI

Figure 3a shows the histogram of PRIs collected on the “helloworld” example from Contiki OS. Values ≥ 12 ms are collected at 12ms. The result shows that most PRIs are clustered around 9.5ms which consists with our result in Table 4. The majority, roughly ranged [9.0, 10.3]ms, corresponds to the usual response time as depicted by Sensor Node 1 in Figure 1.

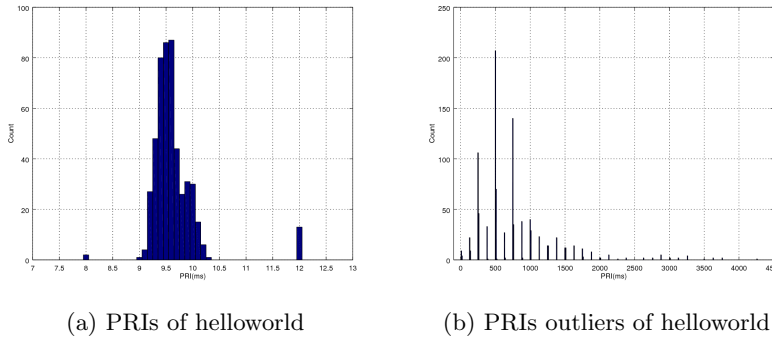


Figure 3: helloworld PRIs

We further plotted the upper outliers, mostly ranged [12, 2000]ms, in Figure 3b. Unfortunately we do not have a solution to investigate the exact cause of such delay, as we were unable to control the code execution that requires environmental interaction within a timing critical context. Nevertheless, we suppose these outliers correspond to the extended response time as depicted by Sensor Node 2 in Figure 1. The distribution described by Figure 3b is the fingerprint of the “helloworld” example.

The result in Figure 3 shows a clear gap between the usual PRIs and extended PRIs. In fact other applications we experimented also showed the same property. This implies that an adversary can easily draw a threshold by observing the whole PRI distribution and then filter out the fingerprint. In our experiments

the threshold is set to 12ms but any other values within the gap would also work.

We collected the fingerprints for three programs taken from the Contiki OS examples:

broadcast This program periodically broadcasts a constant message.

powertrace This program records the power consumption and broadcasts a constant message.

Sensorpayload This program is based on the “er-rest-example” embedded together with sensor accesses taken from “cc2538-demo”. It captures a real case scenario where three different sensors, namely Temperature, Vdd and ALS, are being accessed through CoAP.

Specifically for “Sensorpayload” we collected fingerprints for 8 different scenarios where different sensors are being accessed. For each program we independently collected 2 fingerprints for comparison.

Table 5 summarises the total 20 fingerprints we collected for the experiment.

4.3.2 Fingerprint Matching

During the experiments we realised that most of the fingerprints do not adhere to common distributions; therefore we used a non parametric test, the Kolmogorov-Smirnov Distance [24], as our test statistic. This is a well understood statistic with previous uses in side channel analysis [48]. Table 6 summarises the relative KS distances computed on each pair of fingerprints in our experiments.

Although fingerprints collected on the same application were rejected by the Kolmogorov-Smirnov Test, we noticed that their KS distance tends to be smaller comparing to fingerprints collected on different programs, as the **bold** cells marked in Table 6.

By adapting our distinguisher to utilise the minimum KS distance, we were able to identify 13 out of 20 fingerprints successfully. The ‘overlapping’ fingerprints are mainly due to the “Sensorpayload” program, which access different sensors, but otherwise has identical program code. Thus we did expect that the different instantiations of it would lead to very similar fingerprints.

5 Conclusion

In this paper we explore, for the first time, the use of packet lengths and response times, which are protocol level side channels, as means to recover information about IoT ‘things’. We do this experimentally, which we base on two extremely popular devices running on a popular open source OS, with a typical stack of protocols. Whilst we do not cover a wide range of devices, the fact that two of the most popular devices show the characteristics that we hypothesise, gives credibility to our results. Our results show that it is possible (in principle) to recover information about a device and its function (i.e. the hardware and the

software that runs on it) via inspecting encrypted traffic that it produces. We also point out that ICMP messages can be distinguished from each other despite the use of encryption.

In order to mitigate the leakage that is given by packet lengths, previous works recommend padding [15]. We echo this recommendation. Whilst padding to MTU is considered inefficient for the Internet, it is in fact highly appropriate for 6LoWPAN because:

- It completely hides the length of original plaintext.
- 6LoWPAN has only a low MTU of 127 bytes; therefore the overhead is acceptable.
- It induces negligible computational overhead.

With regard to the leaking information about the device or OS, we suggest strictly applying the standard MTU to eliminate the differences in drivers. Although there is a potential of performance downgrade, it will also improve the compatibility among different devices.

In order to mitigate the leakage given by response times, the natural countermeasure is to write time-constant code, which is known to be notoriously difficult. But two approaches are available to a software developer:

- Randomly delay the response. This essentially adds noise to the measurements of the adversary.
- Use a threshold response time, i.e. a request is either responded at a predefined time or not responded at all.

Within the context of 6LoWPAN the second method is recommended as most 6LoWPAN application would tolerate missing packets and timer is available on most platforms. However, the threshold must be carefully chosen to preserve the functionality of the 6LoWPAN application.

Fingerprint Experiment Programs

Index	Program	Total Size	Fingerprint Size	Note
1	broadcast	6489	593	
2	broadcast	6164	639	
3	powertrace	7142	539	
4	powertrace	7079	561	
5	Sensorpayload	7338	987	Temperature + ALS
6	Sensorpayload	7963	934	Temperature + ALS
7	Sensorpayload	7143	1195	Temperature only
8	Sensorpayload	7316	1096	Temperature only
9	Sensorpayload	7895	827	ALS only
10	Sensorpayload	7867	789	ALS only
11	Sensorpayload	7428	1138	No reading
12	Sensorpayload	7462	833	No reading
13	Sensorpayload	6565	1391	Vdd only
14	Sensorpayload	7193	1111	Vdd only
15	Sensorpayload	7672	955	Temperature, Vdd and ALS
16	Sensorpayload	7790	1023	Temperature, Vdd and ALS
17	Sensorpayload	7864	931	Vdd + ALS
18	Sensorpayload	7936	987	Vdd + ALS
19	Sensorpayload	7217	1222	Temperature + Vdd
20	Sensorpayload	7050	1228	Temperature + Vdd

Table 5: Fingerprint Experiment Programs

Fingerprint Experiment Relative KS-Distances

Index	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
1	N/A	7.8	75	78.9	48.3	58.4	19.2	15.8	62.2	58.5	34.2	34	13.2	10.3	54.1	57	52.8	58.7	19.4	13.3
2	7.8	N/A	81.7	85.6	54.8	65.1	25.6	22.4	68.7	65	40.8	40.4	14.9	16.7	60.7	63.7	59.3	65.4	26	19.3
3	75	81.7	N/A	6.9	30.3	22.9	56.3	59.5	16.2	21.2	41.2	41.5	67	65	21.5	22.8	23.2	29.8	55.8	62.5
4	78.9	85.6	6.9	N/A	32.3	22.5	60.3	63.6	17.9	20.9	45.2	45.6	71	69.3	25.2	22.5	26.7	27.4	59.8	66.7
5	48.3	54.8	30.3	32.3	N/A	11.3	29.4	32.7	15.1	14	14.5	14.7	40.1	38.3	14.8	16.1	17.1	23.8	28.9	35.6
6	58.4	65.1	22.9	22.5	11.3	N/A	39.5	42.9	9	4.9	24.4	24.7	50.3	48.4	14	16	16.8	23.7	39.2	45.8
7	19.2	25.6	56.3	60.3	29.4	39.5	N/A	4.4	43.2	39.5	19.7	17	14.5	9.6	35.3	38.5	33.7	39.7	10.8	9.6
8	15.8	22.4	59.5	63.6	32.7	42.9	4.4	N/A	46.6	42.8	19.5	19.1	14.9	6.2	38.4	41.7	37.1	43.2	10.9	9
9	62.2	68.7	16.2	17.9	15.1	9	43.2	46.6	N/A	5.4	28.2	29	54	52.1	10.3	12.4	12.3	19	42.8	49.7
10	58.5	65	21.2	20.9	14	4.9	39.5	42.8	5.4	N/A	24.5	25.4	50.2	48.4	11.5	13.4	14.4	20.9	39.1	46.4
11	34.2	40.8	41.2	45.2	14.5	24.4	19.7	19.5	28.2	24.5	N/A	7.4	28.2	24.7	20	23.1	19	24.7	15	23.9
12	34	40.4	41.5	45.6	14.7	24.7	17	19.1	29	25.4	7.4	N/A	25.7	23.9	20.7	23.7	19	25	14.9	21.7
13	13.2	14.9	67	71	40.1	50.3	14.5	14.9	54	50.2	28.2	25.7	N/A	11.7	45.9	49.2	44.5	50.6	15.7	8.9
14	10.3	16.7	65	69.3	38.3	48.4	9.6	6.2	52.1	48.4	24.7	23.9	11.7	N/A	44	47.2	42.8	48.7	10.5	7.2
15	54.1	60.7	21.5	25.2	14.8	14	35.3	38.4	10.3	11.5	20	20.7	45.9	44	N/A	3.8	6.7	11.2	34.8	41.5
16	57	63.7	22.8	22.5	16.1	16	38.5	41.7	12.4	13.4	23.1	23.7	49.2	47.2	3.8	N/A	8.8	8.9	37.9	44.7
17	52.8	59.3	23.2	26.7	17.1	16.8	33.7	37.1	12.3	14.4	19	19	44.5	42.8	6.7	8.8	N/A	11.5	33.6	40
18	58.7	65.4	29.8	27.4	23.8	23.7	39.7	43.2	19	20.9	24.7	25	50.6	48.7	11.2	8.9	11.5	N/A	39.4	46.1
19	19.4	26	55.8	59.8	28.9	39.2	10.8	10.9	42.8	39.1	15	14.9	15.7	10.5	34.8	37.9	33.6	39.4	N/A	10.3
20	13.3	19.3	62.5	66.7	35.6	45.8	9.6	9	49.7	46.4	23.9	21.7	8.9	7.2	41.5	44.7	40	46.1	10.3	N/A

Table 6: Relative KS-Distances of Experimented Fingerprints (Multiplied by 100 for readability. Indexes refer to Table 5 Minimum in each row marked as **bold**.)

References

- [1] Available at <https://www.bluetooth.com/what-is-bluetooth-technology/bluetooth-technology-basics/low-energy>.
- [2] Available at <https://openwsn.atlassian.net/wiki/pages/viewpage.action?pageId=688187>.
- [3] Available at <http://www.freertos.org/>.
- [4] Available at <https://www.riot-os.org/>.
- [5] Available at <http://sourceforge.net/projects/tinydtls/>.
- [6] ZigBee Alliance. ZigBee specification. Technical report, June 2005.
- [7] ZigBee Alliance. Zigbee ip and 920ip, 2014.
- [8] Shuo Chen, Rui Wang, XiaoFeng Wang, and Kehuan Zhang. Side-channel leaks in web applications: A reality today, a challenge tomorrow. In *Security and Privacy (SP), 2010 IEEE Symposium on*, pages 191–206. IEEE, 2010.
- [9] Raheem Beyah Chris Wampler, A. Selcuk Uluagac. Information leakage in encrypted ip video traffic. *ieee-globecom 2015*, 2015.
- [10] A. Conta, S. Deering, and M. Gupta. Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification. RFC 4443 (Draft Standard), March 2006. Updated by RFC 4884.
- [11] Scott Coull and Kevin Dyer. Privacy failures in encrypted messaging services: Apple imessage and beyond. *arXiv preprint arXiv:1403.1906*, 2014.
- [12] Adam Dunkels. The contikimac radio duty cycling protocol. *SICS Report*, 2011.
- [13] Adam Dunkels, Fredrik Österlind, and Zhitao He. An adaptive communication architecture for wireless sensor networks. In *Proceedings of the 5th International Conference on Embedded Networked Sensor Systems, SenSys '07*, pages 335–349, New York, NY, USA, 2007. ACM.
- [14] Morris J. Dworkin. Sp 800-38c. recommendation for block cipher modes of operation: The ccm mode for authentication and confidentiality. Technical report, National Institute of Standards & Technology, Gaithersburg, MD, United States, 2004.
- [15] Kevin P. Dyer, Scott E. Coull, Thomas Ristenpart, and Thomas Shrimpton. Peek-a-boo, i still see you: Why efficient traffic analysis countermeasures fail. In *Proceedings of the 2012 IEEE Symposium on Security and Privacy, SP '12*, pages 332–346, Washington, DC, USA, 2012. IEEE Computer Society.

- [16] Kevin P Dyer, Scott E Coull, Thomas Ristenpart, and Thomas Shrimpton. Protocol misidentification made easy with format-transforming encryption. In *Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security*, pages 61–72. ACM, 2013.
- [17] International Organization for Standardization ISO. ISO/IEC 7498-1 Information technology - Open Systems Interconnection - Basic Reference Model: The Basic Model. Technical report, ISO, June 1994.
- [18] IEEE 802.15.4 Working Group. IEEE Standard for Information Technology- Telecommunications and Information Exchange Between Systems- Local and Metropolitan Area Networks- Specific Requirements Part 15.4: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Low-Rate Wireless Personal Area Networks (WPANs). Technical report, IEEE 802.15.4 Working Group, 2006.
- [19] Dominik Herrmann, Rolf Wendolsky, and Hannes Federrath. Website fingerprinting: attacking popular privacy enhancing technologies with the multinomial naïve-bayes classifier. In *Proceedings of the 2009 ACM workshop on Cloud computing security*, pages 31–42. ACM, 2009.
- [20] Yih-Chun Hu, Adrian Perrig, and David B Johnson. Wormhole attacks in wireless networks. *Selected Areas in Communications, IEEE Journal on*, 24(2):370–380, 2006.
- [21] René Hummen, Jens Hiller, Hanno Wirtz, Martin Henze, Hossein Shafagh, and Klaus Wehrle. 6lowpan fragmentation attacks and mitigation mechanisms. In *Proceedings of the Sixth ACM Conference on Security and Privacy in Wireless and Mobile Networks, WiSec '13*, pages 55–66, New York, NY, USA, 2013. ACM.
- [22] John B Kenney. Dedicated short-range communications (dsrc) standards in the united states. *Proceedings of the IEEE*, 99(7):1162–1182, 2011.
- [23] Paul Kocher, Joshua Jaffe, and Benjamin Jun. Differential power analysis. In *Advances in cryptology CRYPTO99*, pages 789–789. Springer, 1999.
- [24] Andrej N Kolmogorov. *Sulla determinazione empirica di una legge di distribuzione*. na, 1933.
- [25] Konrad-Felix Krentz, Hosnieh Rafiee, and Christoph Meinel. 6lowpan security: Adding compromise resilience to the 802.15.4 security sublayer. In *Proceedings of the International Workshop on Adaptive Security, ASPI '13*, pages 1:1–1:10, New York, NY, USA, 2013. ACM.
- [26] Ioannis Krontiris, Thanassis Giannetsos, and Tassos Dimitriou. Launching a sinkhole attack in wireless sensor networks; the intruder side. In *Networking and Communications, 2008. WIMOB'08. IEEE International Conference on Wireless and Mobile Computing,*, pages 526–531. IEEE, 2008.

- [27] M. Kulkarni, A. Patel, and K. Leung. Mobile IPv4 Dynamic Home Agent (HA) Assignment. RFC 4433 (Proposed Standard), March 2006.
- [28] Vijay Kumar, George Oikonomou, and Theo Tryfonas. Traffic forensics for ipv6-based wireless sensor networks and the internet of things. In *Internet of Things (WF-IoT), 2016 IEEE 3rd World Forum on*, pages 633–638. IEEE, 2016.
- [29] Xiapu Luo, Peng Zhou, Edmond WW Chan, Wenke Lee, Rocky KC Chang, and Roberto Perdisci. Https: Sealing information leaks with browser-side obfuscation of encrypted flows. In *NDSS*, 2011.
- [30] Luke Mather and Elisabeth Oswald. Pinpointing side-channel information leaks in web applications. *Journal of Cryptographic Engineering*, 2(3):161–177, 2012.
- [31] D. McGrew and D. Bailey. AES-CCM Cipher Suites for Transport Layer Security (TLS). RFC 6655 (Proposed Standard), July 2012.
- [32] G. Montenegro, N. Kushalnagar, J. Hui, and D. Culler. Transmission of IPv6 Packets over IEEE 802.15.4 Networks. RFC 4944 (Proposed Standard), September 2007. Updated by RFCs 6282, 6775.
- [33] Online: <http://http://www.contiki-os.org/>.
- [34] Online: http://www.willow.co.uk/html/telosb_mote_platform.php.
- [35] Online: <http://www.ti.com/product/CC2538>.
- [36] Pavan Pongle and Gurunath Chavan. A survey: Attacks on rpl and 6lowpan in iot. In *Pervasive Computing (ICPC), 2015 International Conference on*, pages 1–6. IEEE, 2015.
- [37] J. Postel. User Datagram Protocol. RFC 768 (INTERNET STANDARD), August 1980.
- [38] Sridhar Rajagopal, Richard D Roberts, and Sang-Kyu Lim. Ieee 802.15.7 visible light communication: modulation schemes and dimming support. *IEEE Communications Magazine*, 50(3), 2012.
- [39] E. Rescorla and N. Modadugu. Datagram Transport Layer Security Version 1.2. RFC 6347 (Proposed Standard), January 2012. Updated by RFCs 7507, 7905.
- [40] Naveen Sastry and David Wagner. Security considerations for ieee 802.15.4 networks. In *Proceedings of the 3rd ACM Workshop on Wireless Security, WiSe '04*, pages 32–42, New York, NY, USA, 2004. ACM.

- [41] Alexander Schaub, Emmanuel Schneider, Alexandros Hollender, Vinicius Calasans, Laurent Jolie, Robin Touillon, Annelie Heuser, Sylvain Guilley, and Olivier Rioul. Attacking suggest boxes in web applications over https using side-channel stochastic algorithms. In *Risks and Security of Internet and Systems*, pages 116–130. Springer, 2014.
- [42] R. Seggelmann, M. Tuexen, and M. Williams. Transport Layer Security (TLS) and Datagram Transport Layer Security (DTLS) Heartbeat Extension. RFC 6520 (Proposed Standard), February 2012.
- [43] Z. Shelby, K. Hartke, and C. Bormann. The Constrained Application Protocol (CoAP). RFC 7252 (Proposed Standard), June 2014.
- [44] Virendra Pal Singh, Aishwarya S Anand Ukey, and Sweta Jain. Signal strength based hello flood attack detection and prevention in wireless sensor networks. *International Journal of Computer Applications*, 62(15), 2013.
- [45] Pascal Thubert, Thomas Watteyne, Maria Rita Palattella, Xavier Vilajosana, and Qin Wang. Ietf 6tisch: Combining ipv6 connectivity with industrial performance. In Leonard Barolli, Ilsun You, Fatos Xhafa, Fang-Yie Leu, and Hsing-Chung Chen, editors, *IMIS*, pages 541–546. IEEE Computer Society, 2013.
- [46] Franck Veysset, Olivier Courta, Olivier Heen, IR Team, et al. New tool and technique for remote operating system fingerprinting. *Intranode Software Technologies*, 4, 2002.
- [47] Mohammad Wazid, Avita Katal, Roshan Singh Sachan, RH Goudar, and Dhirendra Pratap Singh. Detection and prevention mechanism for black-hole attack in wireless sensor network. In *Communications and Signal Processing (ICCSP), 2013 International Conference on*, pages 576–581. IEEE, 2013.
- [48] Carolyn Whitnall, Elisabeth Oswald, and Luke Mather. An exploration of the kolmogorov-smirnov test as a competitor to mutual information analysis. In *International Conference on Smart Card Research and Advanced Applications*, pages 234–251. Springer, 2011.
- [49] T. Winter, P. Thubert, A. Brandt, J. Hui, R. Kelsey, P. Levis, K. Pister, R. Struik, JP. Vasseur, and R. Alexander. RPL: IPv6 Routing Protocol for Low-Power and Lossy Networks. RFC 6550 (Proposed Standard), March 2012.
- [50] Charles V Wright, Lucas Ballard, Scott E Coull, Fabian Monrose, and Gerald M Masson. Spot me if you can: Uncovering spoken phrases in encrypted voip conversations. In *Security and Privacy, 2008. SP 2008. IEEE Symposium on*, pages 35–49. IEEE, 2008.
- [51] Charles V Wright, Scott E Coull, and Fabian Monrose. Traffic morphing: An efficient defense against statistical traffic analysis. In *NDSS*, 2009.