Using the Physical Layer to Detect Attacks on Building Automation Networks

Andreas Zdziarstek(✉), Willi Brekenfelder, and Felix Eibisch

University of Rostock, Rostock, Germany
{andreas.zdziarstek,willi.brekenfelder,felix.eibisch}@uni-rostock.de

Abstract. This work investigates possible methods of adding security features to building automation networks in the form of intrusion or tamper detection by using the physical layer. This is a concept that is widely known in the field of wireless communications but is—as of now—less prevalent in wired environments. We propose three distinct and complementary methods which rely on electrical fingerprinting of devices and the communication medium, as well as active radio-frequency probing of the network. To assess their effectiveness, we conduct a series of experiments in a building automation system test environment.

Keywords: Physical layer security · Network security · Building automation · Network intrusion detection and prevention

1 Introduction

In modern and large public buildings such as hospitals, universities or schools, a high amount of complex electrical signaling may be required to control its internal appliances such as lights, heating, air-treatment, or access controls. To alleviate some of this complexity, there are standards for building automation systems (BASes) which enable the use of multiplexed control lines in a network topology.

Unfortunately, there are major security flaws present within widely adopted BAS protocols [9]. Messages on the network medium are either sent using obsolete forms of encryption or none at all. This enables an attacker with physical access to the cabling to read traffic and potentially send harmful commands to connected devices. In a public building, it is not easy to prevent this kind of access as reaching a BAS endpoint can be as easy as popping off a light switch and connecting to the exposed wires. Even if there are no critical devices connected to that particular bus line, previous research has shown that logging seemingly harmless lighting data over an extended period of time provides information that may raise privacy and security concerns [16]. This is especially true if the building employs motion detectors for lighting control.

N. Park et al. (Eds.): SecureComm 2020, LNICST 336, pp. 372–390, 2020.
https://doi.org/10.1007/978-3-030-63095-9_24

These concerns are now beginning to be addressed by newer standards, examples being KNX Secure [15] and BACnet/SC [7]. But, given the large number of insecure devices already installed and still being sold, widespread practical adoption of secure protocols cannot be expected in the near future.

In this situation, it becomes necessary to think about adding new security features to existing network installations. Generally, it can be assumed that any device-specific modifications—for example adding an encryption layer—would require prohibitive effort and may interfere with other communications on the bus. This limits the possible interventions to adding special devices which monitor the network and either intervene in and/or report any suspicious activity. Such devices can be seen as being a part of an intrusion detection system (IDS) which would safeguard the building network against malicious activity. Generally, these systems operate by analyzing network traffic, matching patterns and detecting unusual activity, i.e. they tend to operate on OSI layer 2 and up.

It is common to just assume the inherent security of BAS networks when they are not directly connected to other wider area networks such as local area networks or the internet. But as mentioned above, especially in publicly accessible buildings this assumption may be treacherous. Nevertheless, a physically isolated and reasonably static network structure does offer unique possibilities of detecting an intrusion. In this paper, we explore ideas on how to make use of the physical layer characteristics of the network and the connected devices for attack detection. Our general assumption is that every network and device has uniquely detectable features when analyzing the bus on an electrical level. At least in the field of wireless networking, previous research has shown the feasibility of such physical layer approaches, for example using radio-frequency (RF) fingerprinting in WLAN networks [12].

2 Related Work

2.1 Device Fingerprinting

As a starting point, it makes sense to look at Physical Layer Security research in wireless networks in the hopes of transporting some of the concepts over to wired systems. For example, experiments by Brik et al. [4] show promising results for their method of distinguishing WLAN transceivers which are of the same model. They are using vector signal analyzer (VSA) hardware devices to capture WiFi signals. Subsequently, they perform signal analysis and machine learning on the data using specific WLAN signal characteristics as classifying features. The authors ascribe the electrical differences between devices to hardware imperfections in the transmitting parts of the chips. While WiFi-specific signal features are not applicable in our case, the general approach of using passive signal measurements and extracting device-specific quirks from the data is relevant to wired use-cases as well.

A method proposed by Wang et al. [21] by contrast does not rely on protocol- or technology-specific characteristics but uses more abstract mathematical features of the captured waveforms. Their method could theoretically be applied to

any time-series input. For our use-case, this poses the question of how significant the differences of those features are for technologies other than WLAN.

Also of particular interest is a technique described in a publication by Gerdes et al. [8] as the authors aim to identify wired Ethernet devices according to a waveform comparison of their synchronization signals by implementing a matched filter. On the hardware level they use oscilloscopes with high sample rates to be able to accurately capture the Ethernet waveforms.

The measured device fingerprints in these methods can be seen as a form of physically unclonable function (PUF). In these cases, the PUF would be an intrinsic characteristic of the transmitters in the network stemming from manufacturing variations. This bears some similarity to ring-oscillator PUFs which can be used to uniquely identify FPGAs or ASICs by measuring the variations of internal delay lines [14]. In contrast, though, in the above fingerprinting methods it is not necessarily clear how the variations may present themselves in the measurements.

2.2 Environment Fingerprinting

As described by Campos and Lovisolo in [5], it is also possible to use these kinds of physical layer methods to verify a location. This means the fingerprint is tied to an environment instead of a device. An example of this would be to compare the received RF spectrum of different locations thereby telling them apart.

A US Patent by Bevan et al. [3] specifies a WLAN localization approach by mapping the complex channel frequency response between a network node and multiple static base stations. In a way, this can be applicable to wired environments as cable length, cable type and physical connection topology can have an effect on the frequency response of a wire thereby altering signals traveling on them.

2.3 Tamper Detection

If both environment and device fingerprinting are brought together, then this can be used to detect if the physical network configuration—that is, both physical topology and device hardware—has been tampered with. An example of this for WiFi-networks is proposed by Bagci et al. in [2].

3 Suggested Methods

Given the publications mentioned above, it is safe to say that there is a healthy research interest in physical layer security in wireless communications with a lot of well-developed methods in existence. Research is less easy to find, though, in the area of wired networks and especially communication buses like those used in building automation. Due to the security problems mentioned in Sect. 1, it is worthwhile to explore this avenue.

3.1 Threat Model Definition

In order to come up with specific ideas on how to guard against security threats, we need to define what the scope of possible attacks may be. As a general rule we assume that the network systems our methods are applied in do not have sufficient higher-level security to guarantee secure authentication and/or confidentiality. Also, we assume that in the age of small battery-powered single-board computers with universal internet access, a permanent air gap between any large building network installation and the outside world is near unenforceable, at least in public and semi-public buildings where a potential attacker is likely to find some accessible network tap to install an illicit device.

The most obvious kind of attack on such a system is what we will call an active one, meaning the attacker interacts with other devices in the network by sending messages. In the initial situation, there is no way to disprove the authenticity of those messages, meaning the attacker may instruct connected devices to perform potentially harmful acts. Potential high risk examples would environmental controls in laboratories, medical facilities, or similar installations. During a medical procedure, even only turning off the lights may cause a catastrophic situation. While this kind of threat clearly has a high potential for serious harm, we also believe it to be the most detectable as the attacker's device needs to openly communicate, thereby revealing its physical layer characteristics.

On the other hand, the possible harmful impact of having an attacker passively listening to the messages on a BAS network may not be as obvious but is a serious issue. An example could be a sensitive facility where an attacker might map guard routes and occupation times of specific offices using the building's motion detectors in preparation for a break-in. In contrast to an active attack, the potential for harm is less direct and momentary but detection may be more complicated as the attacker does not need to openly send data through the network. Nevertheless, it is possible that a connected device does change the electrical characteristics of the medium enough to be detectable.

With the above in mind, we developed three different Physical Layer Security approaches which can be implemented on BAS networks.

3.2 Passive Waveform Fingerprinting

The most basic, general, and non-invasive approach to look at is passive fingerprinting, i.e. using a method like the ones classifying WiFi transmitters and applying it to BAS devices. For this, it is necessary to capture the communication on the bus as analog waveforms. This is possible with a modern digital oscilloscope. As the fingerprint is generally unique with respect to the currently active transmitter, such methods are suited to verifying the identities of known devices in the network. Nevertheless, the cable as the transmission medium does alter the signal due to its impedance characteristic and possible external interference. Moreover, even connected but inactive devices represent an alteration of the network circuit's electrical characteristics that may be large enough to have a measurable impact on the fingerprint of all transmitters. This means, given

suitable equipment and data analysis, passive fingerprinting is potentially able to detect all possible physical network changes.

As for feature extraction from the measured data, it is a good starting point to look at very general approaches, such as the *RF-DNA* method [21] mentioned above. It just requires a time-series representation of the signal which is converted to instantaneous phase and instantaneous frequency data using the Hilbert transform. Then the authors calculate the mean, variance, skewness and kurtosis for each of the signal's amplitude, frequency and phase. This results in a 12-dimensional feature vector for each device and physical makeup of the network that is measured.

3.3 Device Noise Characteristics

While the above approach relies on very general signal analysis, it may be beneficial to make use of BAS physical layer specifics. Generally, their local directly-connected parts rely on some form of bit-based serial bus protocol for simplicity and robustness, typical examples being KNX-TP [19] and BACnet MS/TP [13]. This means there is a well-defined bit timing which we can use to determine exactly when a specific device is sourcing or sinking current to alter the bus wire's potential for data transmission. Our assumption is that during these active phases the transmitting device will—in addition to the data signal—emit a certain amount of noise that can be measured. There are multiple forms of noise emissions, some of them caused by semiconductor impurities and doping errors [10] which become apparent when current flows through them. This means each device has a unique noise level and signature that can be detected given sensitive enough measuring equipment. To enhance resolution, a form of oversampling can be employed by using repeated measurements. For example, it is possible to repeatedly request a serial number or diagnostic information from a device and capture the responses which can then be overlaid.

Formally, we assume each measurement is represented by a time-series

$$\{x_0, \ldots, x_M\}$$

of numerical samples captured with a constant sample rate. This is true for digital oscilloscopes and other analog-digital-conversion-based data capture devices. If N measurements of messages containing the same bit patterns coming from the very same device are then aligned by cross-correlation, each sample becomes a statistical set X_i of measurements where $i, i \leq M$ denotes the sample position in the time-series and $|X_i| = N$. It is possible to compute the sample averages

$$\overline{X}_i = \frac{1}{N} \sum_{j=1}^{N} x_j \text{ with } x_j \in X_i \tag{1}$$

in order to obtain a finer quantization of the data. Now that we have a set of measurements and its mean for each sample point, it is possible to also calculate the sample standard deviations

$$s_{X_i} = \sqrt{\frac{1}{N-1}\sum_{j=1}^{N}\left(x_j - \overline{X}_i\right)^2} \text{ with } x_j \in X_i \tag{2}$$

which gives an indication of the overall noise level during the sampling timepoint i. This measured noise is a sum of every type of noise generated in the measuring circuit, including the oscilloscope's internal components and artifacts of digital quantization. In order to counteract the latter, it is possible to now superimpose all equal symbols within a message and calculate the average standard deviation for every sample point of the active transmission phase, i.e. where the device is sinking or sourcing current.

Formally, find the smallest i_0 where an active transmission phase begins. Depending on the protocol, this could for example be a falling or rising edge which can be detected by calculating the slope between $\overline{X}_{i_0}$ and $\overline{X}_{i_0-1}$. Then find the width of an active transmission phase

$$I = i_e - i_0 \tag{3}$$

where i_e is the endpoint of the active phase following i_0. Now, find the indices of all remaining active phase beginnings $i_m, m \in \mathbb{N}$. Then build a series of sets

$$S_j = \left\{s_{X_{i_0+j}}, \ldots, s_{X_{i_m+j}}\right\} \text{ with } j \leq I, j \in \mathbb{N} \cup \{0\} \tag{4}$$

and calculate the means $\overline{S}_j$. The resulting time-series $\left\{\overline{S}_0, \ldots, \overline{S}_I\right\}$ can offer a more accurate picture of a device's noise characteristics than simply analyzing a single capture of a bit transmission. If the assumption of detectable differences in noise characteristics due to manufacturing imperfections holds true, this average standard deviation time-series is unique to any given device and can be used for identity verification.

3.4 Active Measurements

The methods described so far all rely on passive measurements as this is straightforward and does not interfere with normal bus communications. It may be of interest, though, to look into active probing of a network. Specifically, it may be worthwhile to measure the wideband frequency response of the connected bus line as it should be more sensitive to changes in physical network topology, i.e. how the cables are interconnected and where. This is generally done by emitting a test frequency in the RF range at one point in the network and measuring the received magnitude of the signal at another. By sweeping through a range of frequencies, a response curve can be mapped out. As this reveals the characteristic of the whole network medium, this method may be better suited than the others to guard against passive listening attacks.

As the measurements are meant to be done on an active bus, this poses the question of how the RF path can safely coexist with the BAS communication bus on the same medium without impairing either. The technical solution to this problem will likely be unique to a specific bus system.

4 Experiments

Now, to demonstrate the viability of the above methods we present three experiments as proofs-of-concept. As we intend to show applicability to currently relevant BASes, we decided on KNX-TP as the physical-layer medium.

4.1 KNX

KNX is a well-known BAS standard which is partly derived from the earlier European Installation Bus (EIB). The standard has been widely adopted, with a 2019 press release by the KNX Association [11] claiming over 300 million KNX-ceritified products to be in use around the world. The protocol is designed to work across different physical media including standards for IP-tunneled, wireless, and powerline communication. For a regular wired installation on a BAS-specific medium though, KNX retains the EIB physical layer which is essentially a shared serial bus using twisted-pair cabling called KNX-TP [19].

4.2 Technical Background

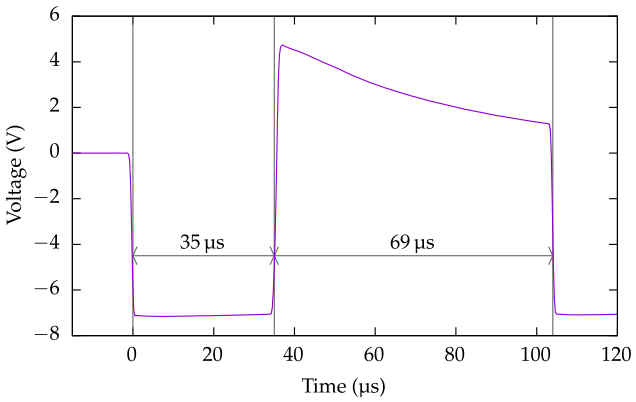

Fig. 1. A KNX-TP zero-bit waveform measured with a digital oscilloscope, averaged over $N = 171$ measurements, with added timing markers. Note that 0 V in this case denotes the resting potential of about 30 V DC as the oscilloscope was set to AC coupling.

The KNX-TP physical layer can be described as a multiplexed serial line including direct-current (DC) power delivery [19]. In order to provide the latter, the bus has a resting voltage of 30 V DC supplied by a KNX-specific active transformer. The current capability of this power supply is usually around a few hundred milliamperes while KNX devices draw a few tens. This is enough for

most switching, control and communication tasks. Devices meant for more power hungry applications will need an additional external source.

As per the data signaling protocol, the resting potential is interpreted as a digital 1, while a 0 is actively generated by the transmitting device. At the start of every zero bit, the transmitter pulls the bus voltage lower by more than 6 V for exactly 35 µs then stops pulling and waits for 69 µs until the start of the next bit [20].

In this waiting phase, a choke coil in the power supply will generate an equalization pulse which overshoots the resting potential by several volts and then drops off slowly. The reason for this implementation is to ensure that the average bus potential does not drop significantly if the bus carries a lot of data traffic as this might interfere with power delivery. Figure 1 shows a plot of a KNX-TP zero symbol with timing markers superimposed.

4.3 Passive Waveform Fingerprinting

In our first experiment we adapted the method used in [21] for KNX-TP. Our test setup is a single KNX line consisting of a KNX push button sensor controlling a DATEC 1630.03160/61100 switch actuator connected to a MEAN WELL KNX-20E-640 power supply. For the push button sensor we used two identical models (MDT BE-TA5508.01), designated as A and B, only one of which is connected to the line during a test run. To control the amount of variables, the switches are only connected at a single specific spot in the line. This yields 2 possible experiment configurations.

We captured $N = 320$ KNX message waveforms for each configuration with a PicoScope 2206A USB-oscilloscope connected to a single endpoint of the network. Using the Hilbert transform we calculated mean, variance, kurtosis, and skewness for the signal amplitude, instantaneous phase, and instantaneous frequency. For visual inspection of the resulting 12-dimensional data we decided to use principal component analysis (PCA) to find a suitable projection to 2D without having to discard whole dimensions arbitrarily. Practically, we relied on a freely available Python implementation included in the *scikit-learn*[1] toolkit. The aim of PCA is to generate new dimensions (principal components) from the data as linear combinations of the original ones [1]. Ideally, PCA will maximize the explained variance in the first principal components in such a way that the rest may be discarded without much information loss. In our case, two principal components were enough to retain 99.9994% of the explained variance.

Figure 2 shows a scatter plot of the results where each point represents a recorded waveform. It is necessary to note that we discarded one of the data points from the set belonging to Switch B as the measurement seems to be a result of a spurious trigger of the oscilloscope rather than an actual message from the switch.

[1] *scikit-learn* is freely available at scikit-learn.org.

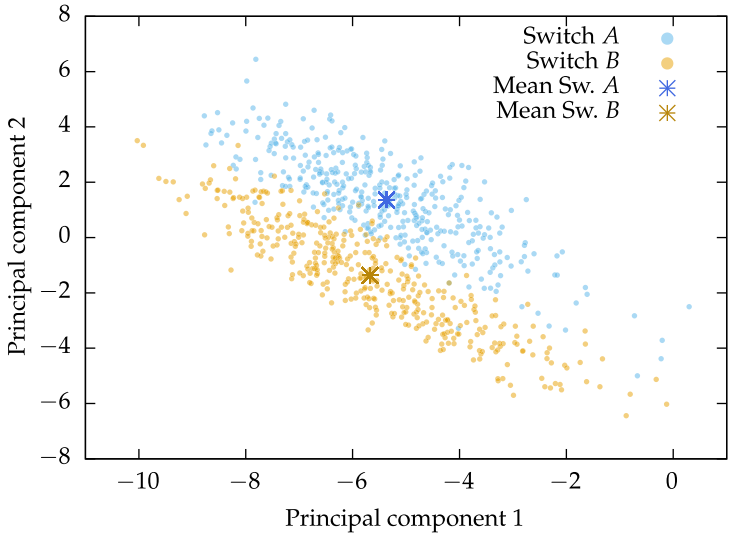

Fig. 2. Dimensionality-reduced scatter plot of the passive device identification test

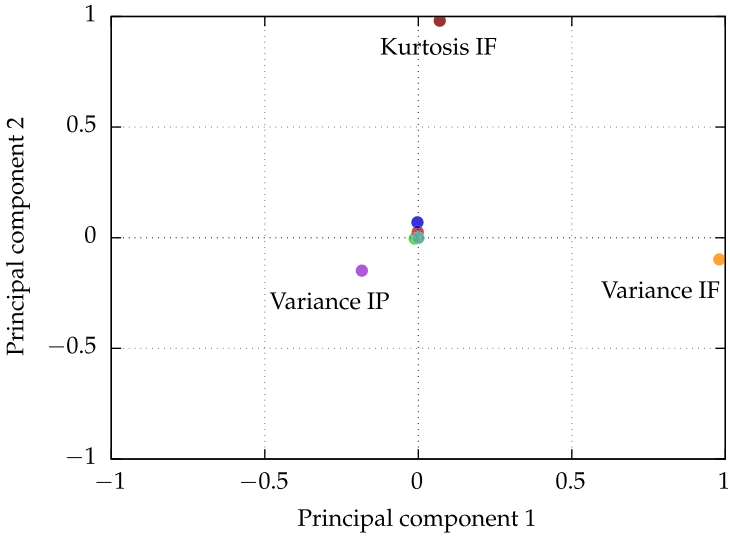

Fig. 3. PCA loadings plot of the passive device identification test

In the diagram, it is easy to see that the point clouds are scattered around two distinct averages and that both clouds are largely separate from each other with only a few outliers. This means, given a high enough N, the two switches are distinguishable using this method.

To explain the meaning of the two principal components (PCs), Fig. 3 shows the corresponding loadings plot of the transform. It indicates that the first PC is mostly influenced by the variance of the instantaneous frequency (IF) and to a much lesser extent by the Kurtosis of the if and the Variance of the instantaneous phase (IP). In contrast, the kurtosis of the IF has the most weight within the second PC together with small influences from IP variance and IF kurtosis. All other original dimensions do not have a sizable effect on either.

4.4 Device Noise Characteristics

Table 1. Configurations of the 2nd experiment

Switch	Experiment	Linecoupler	N
A	E_1	✓	170
	E_2	✗	171
	E_3	✓	349
	E_4	✗	351
B	E_1	✓	178
	E_2	✗	165
	E_3	✓	379
	E_4	✗	325

For the next experiment, we used the technique of comparing active transmission noise characteristics described in Sect. 3.3. The devices and test network were the same as before with the addition of a third device (a KNX linecoupler) that could be connected and disconnected using a switch. This addition allows to check if the method is sensitive to network changes. To also rule out temperature influences on our data, we used a temperature-controlled lab environment and performed two sets of measurements for each possible configuration with different lengths to account for internal device heating. Table 1 lists all 8 test run configurations.

In the KNX-TP physical layer protocol, active transmission phases occur during the first 35 µs of each bit with a value of 0. Throughout that phase the transmitting device needs to sink enough current to drop the bus voltage by about 6V or more (see Sect. 4.2). It is possible to find the indices i_0, i_e by looking for the corresponding rising and falling edges visible in Fig. 1. Formally, find the smallest i_0 such that

$$\overline{X}_{i_0} - \overline{X}_{i_0-1} \leq s \tag{5}$$

where s is a suitable slope threshold value. Now, find the smallest i_e such that

$$\overline{X}_{i_e} - \overline{X}_{i_e-1} \geq s \tag{6}$$

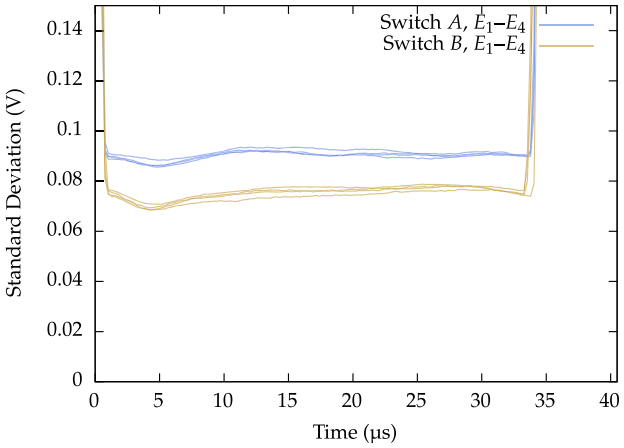

Fig. 4. Results of the transmission noise experiments (Color figure online)

Then use Eq. 3 to determine I and Eq. 4 to build the time series $\{\overline{S}_0, \ldots, \overline{S}_I\}$. Figure 4 shows the resulting plots of all experiments. The differences in N and the addition of the linecoupler evidently had only a very small effect on the data when compared to the differences between the two switches. To improve readability, the plots corresponding to the same switch have the same color and are drawn semi-transparently, making superimposed lines visible. From the plotted results, it is easily possible to tell both switches apart graphically.

4.5 Active Measurements Using a NanoVNA

For the third experiment we used a NanoVNA vector network analyzer to measure the frequency response curve of three different test setups, one with a variable device, one with a variable cable length and a third one also with a variable device but connected to a larger network. The NanoVNA is a device which is able to successively transmit a series of radio frequencies and at the same time measure the amplitude of the returned signal in order to map the spectrum. To deal with the problem of interfacing to an active communication bus without functional interference, we developed an RF transmitter amplifier as well as an attenuator with AC coupling on the receiving end. Figure 5 shows the schematics of both circuits.

The TX amplifier is capable of subtracting a small analog waveform from the KNX voltage by modulating a bus load current small enough to not be registered by any KNX devices. To control the load current, we used two bipolar-junction transistors (NTE2633) originally designed for high-frequency video amplification. The transistors are operating in parallel Class A mode with resistor dividers providing biasing. An AC coupling capacitor and a resistor at the input roughly provide a 50 Ω impedance.

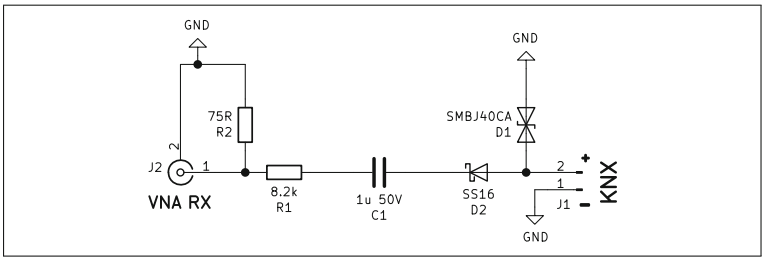

(a) Receiver protection, attenuation and AC coupling circuit

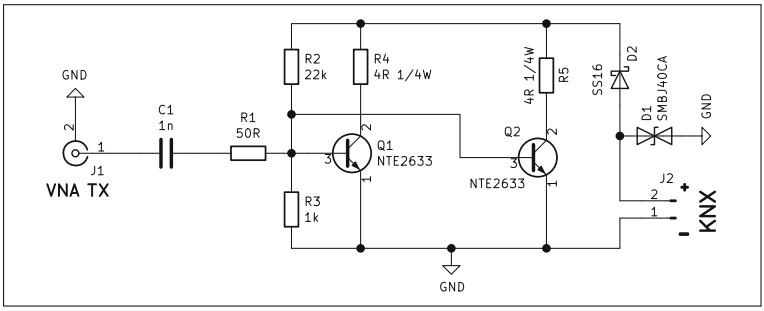

(b) Transmitter amplifier circuit

Fig. 5. Circuit diagrams of the NanoVNA/KNX-Adapters

The receiver-side circuit is a simple passive AC coupling capacitor together with a resistor divider providing at least 40 dB of attenuation or more, depending on the input impedance of the connected receiver. Together this is enough to limit the maximum voltage seen by the receiver to safe levels even during KNX communication causing large AC swings.

To explore if both device and topology changes are significantly reflected in the system's frequency response curve, we performed three sets of experiments. Figure 6 shows their respective setups. As can be seen, the first two have the same general outline, with the TX and RX circuits connected across a 10 m length of KNX-TP wire with a center tap. In the first one (Fig. 6a), two different devices were connected at that position using a 1 m lead. Specifically, the devices were a custom KNX interface board based on NCN5121 transceiver chip by ON Semiconductor [17], and one of the MDT push-button switches from before. In the second one (Fig. 6b) the MDT switch is connected using either a 1 m or a 6 m lead.

The third setup (Fig. 6c) is in essence an extension of the first experiment where we connected the devices to a KNX test line containing 4 permanently installed devices (2 KNX-USB dongles, a quad relay, and a line coupler interconnected with about 1.5 m of cabling in total) and attached both VNA input and output to taps at the far end. Our intention is to determine if a more complex network and a larger amount of connected devices will make small changes like

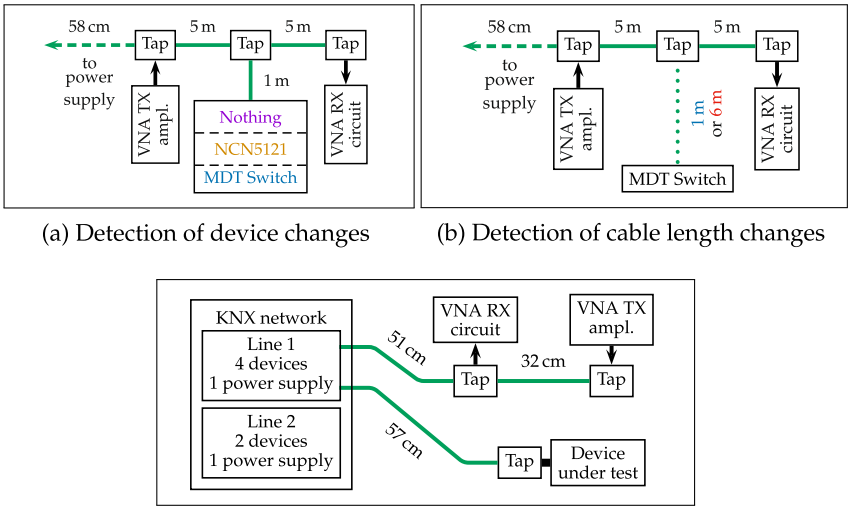

(a) Detection of device changes (b) Detection of cable length changes

(c) Detection of device changes in a complex network

Fig. 6. The NanoVNA experimental setups

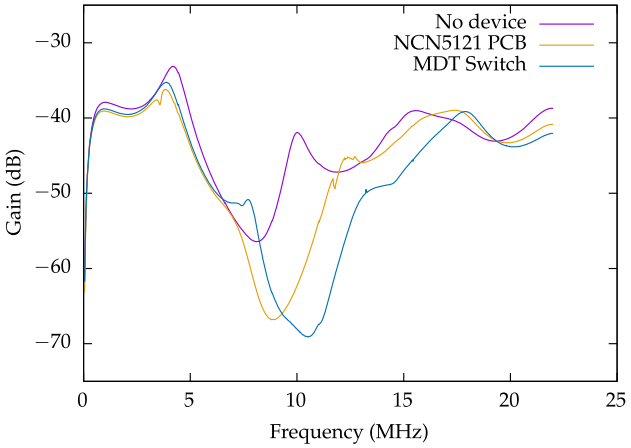

Fig. 7. Results of the NanoVNA frequency sweep with device changes (Color figure online)

swapping out a device less detectable. We also decided to connect both the input and output of the VNA close together as this most closely mirrors a situation where a single intrusion detection device is added at a single point in a network.

In all three parts of the experiment, the NanoVNA was configured to produce a frequency sweep between 50 KHz and 22 MHz in 1010 discrete steps, yielding a resolution of 21.73 KHz/step.

To achieve this resolution, the NanoVNA does in fact perform 10 successive partial sweeps of the spectrum with 101 samples each. The measured quantity is gain in decibels, i.e. the logarithmic power ratio of the signal transmitted by the NanoVNA on its TX side versus what it receives back at its RX side. The signal path includes the damping and amplification circuits as well as the connected KNX network.

Figure 7 shows the results of the device change experiments. Each of the three measurements was taken $N = 100$ times and then averaged per frequency step to improve accuracy. Average sample standard deviations were 0.02 dB for "no device", 0.15 dB for the MDT switch and 0.04 dB for the NCN5121 interface board. Evidently, there is a significant difference in frequency response between devices, i.e. their response curves are unique and—as evidenced by the low standard deviation between measurements—virtually static.

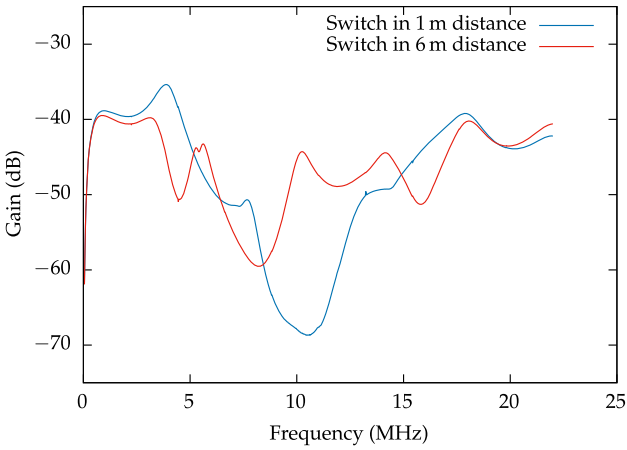

Fig. 8. Results of the NanoVNA frequency sweep with cable length changes

The results of the second experiment plotted in Fig. 8 show that also cable length has an easily detectable impact on the system's frequency response with both plotted curves being largely dissimilar. As before, sample standard deviations are low with an average of 0.12 dB for the 1 m-lead and 0.04 dB for 6 m. The blue line in this plot is nearly congruent with the blue line of Fig. 7 which is expected because the experimental setup is the same (see Fig. 6). Interestingly, some features of the red curve (6 m distance) coincide with the "no devices" plot from Fig. 7. Namely, the local gain minima at about 8 MHz and 12 MHz exist in both cases. Yet, the red curve also has local minima at about 4.5 MHz and 15.8 MHz which are nonexistent in the other. A possible reason for this partial similarity may be that the device's influence on the frequency response was damped by the added resistance of the 5 m lead between it and the measurement path.

The third experiment's results can be seen in Fig. 9, this time with $N = 99$ per curve. Note that in this experiment we have connected two specimens of each device, labeled A and B. As can be seen, differences between two specimens of the same device model are visually undetectable. Differences between the switches, the NCN5121 PCBs, and leaving the tap unconnected, though, are less obvious than before but still easily discernible when looking at the range from 10 MHz to 15 MHz and the location of the minima around 6 MHz. The average sample standard deviations are similarly low to the former experiments with 0.07 dB being the highest. This means, while the visible differences look small, they are orders of magnitude larger than the sample standard deviations and thereby statistically significant.

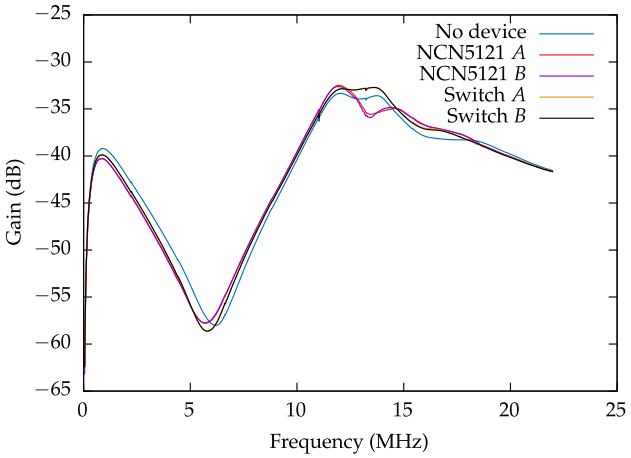

Fig. 9. Results of the NanoVNA frequency sweep with device changes in a complex setup

5 Discussion

Given these results, it is clear that our proposed approaches incorporating passive waveform fingerprinting and device noise characteristics are sensitive to device changes in the network. Furthermore, we have demonstrated that it is possible to detect a swapped-out device even if it is from the same make and model as the original one. This would be a highly valuable feature for tamper and intrusion detection as no attacker would be able to assume a fake device identity without fully taking over a known, present device in the network—a significantly higher hurdle than adding a fake device of their own. With regards to the threat model specified in Sect. 3.1, this method seems to be suited to counteract an active attack.

Similarly, the results also show the active probing approach to be sensitive to both changes in connected devices as well as cabling changes. This means the

method can be used to counteract the passive kind of attack defined in Sect. 3.1 as this would at the very least entail tampering with a cable endpoint. It is, though, worth noting that given our limited dataset it is not easily possible to match a change in frequency response to any specific event. This contrasts with the other approaches where the offending device can be pinpointed.

Given their respective limitations, it is advisable to not singularly use any of the methods for attack detection but to combine them with each other. A combined system would be able to alert the user to a wide range of possible suspicious events while also delivering a form of device authentication by hardware identification. Of course, any or all of the concepts can also be added to existing higher level intrusion detection systems to provide additional data points indicating the current network state.

Also, as our practical experiments were limited to KNX networks and devices, the question remains how easily transferable the methods are to other bus systems. Nevertheless the concepts introduced in Sect. 3 are sufficiently generalized to guide the implementation work for any wired electrical bus system. Of course, depending on the systems' technical features this can be more difficult in some cases.

For passive signal analysis, it is necessary to have measuring equipment that is able to sample the waveforms fast enough. With a constant speed of 9600 bit/s, KNX-TP is quite slow when compared to other building, industrial and vehicular serial buses such as RS485/Profibus (up to 12 Mbit/s) and CAN-bus (up to 1 Mbit/s) [6,18]. Speeds like that would require a considerably higher investment in measurement hardware.

While bus speed is of lesser importance for the active probing method, its practical implementation is very dependent on the electrical specifications of the physical medium. Our self-designed RF amplifier could use the KNX DC level to power itself and modulate an AC waveform on top of it. For other bus systems this might be possible to do in a similar fashion if there is a positive resting potential to exploit such as in the CAN bus [6]. If a bus system uses active-high signaling though the amplifier will need to be powered from an external source which would lead to a more complicated design.

6 Conclusion

In this work, we introduced three different methods for anomaly or attack detection intended for the physical layer of BASes fieldbus networks. Two of those are based on passive observation of the electrical signals on the bus while the third method relies on actively mapping the high-frequency response spectrum of a network. To demonstrate their viability, we performed a series of experiments on a widely adopted BAS fieldbus system. The results comprehensively support our assumption that physical layer security can be a valuable addition to existing BASes to help ensure their integrity.

7 Future Work

The above being said, it is clear that more testing is needed to conclusively verify that the approaches are universally usable among different BAS network types and topologies. Also, it would be useful to have a larger database of network setups and devices with their respective measured characteristics. This could allow us to find feature patterns correlating with certain specific kinds of network changes leading to better insight into what data anomaly indicates which actual event or class of events. We plan to conduct further experiments on more complex setups as well as different physical layer technologies used in fieldbus-type networks.

As these are problems related to classification and pattern matching, we also intend to investigate the benefits of applying different machine-learning algorithms to our data. Especially the *RF-DNA*-derived passive fingerprinting approach may be well complemented by high-dimensional classification methods.

Another area of interest is the methods' practical implementation as a self-contained physical layer intrusion detection device. At least the first two experiments as described in Sect. 4 use comparatively expensive hardware, specifically the oscilloscope. Also all three methods so far rely on offline data analysis using a personal computer. For practical usage, it would be beneficial to limit these costs. Examples for improvements in this area could be to substitute the oscilloscope with cheaper and less versatile analog-digital converters and to use less expensive computation hardware, e.g. embedded controllers or single-board computers. Given such improvements, it may be possible to devise low-cost physical layer security devices that can be easily added to existing networks.

Supplementary Material

The measurements and results of the above experiments can be found at https://opsci.informatik.uni-rostock.de/repos/datasets/bas-pls/bas-pls-res.zip.

Acknowledgment. This research was funded by a grant from the German Federal Ministry for Economic Affairs and Energy in accordance with a resolution passed by the German federal parliament.

References

1. Abdi, H., Williams, L.J.: Principal component analysis. WIREs Comput. Stat. **2**(4), 433–459 (2010)
2. Bagci, I.E., Roedig, U., Martinovic, I., Schulz, M., Hollick, M.: Using channel state information for tamper detection in the internet of things. In: Proceedings of the 31st Annual Computer Security Applications Conference, pp. 131–140. Association for Computing Machinery, Los Angeles (2015)
3. Bevan, D.D., Averin, I., Lysyakov, D.: RF fingerprinting for location estimation. US Patent 8,170,815 B2 (USA). R.B. LP. May 1, 2012. http://patft1.uspto.gov/netacgi/nph-Parser?patentnumber=8170815. Accessed 14 Feb 2020

4. Brik, V., Banerjee, S., Gruteser, M., Oh, S.: Wireless device identification with radiometric signatures. In: Proceedings of the 14th ACM International Conference on Mobile Computing and Networking, pp. 116–127. Association for Computing Machinery, San Francisco (2008)
5. Campos, R.S., Lovisolo, L.: RF fingerprinting location techniques (chap. 15). In: Handbook of Position Location, pp. 487–520. Wiley (2011). https://doi.org/10.1002/9781118104750. ISBN 9781118104750
6. Corrigan, S.: Introduction to the Controller Area Network (CAN). SLOA101B. Application Report. Texas Instruments Incorporated, May 2016. https://www.ti.com/lit/an/sloa101b/sloa101b.pdf. Accessed 07 Feb 2020
7. Fisher, D., Isler, B., Osborne, M.: BACnet Secure Connect. A Secure Infrastructure for Building Automation. White Paper, version 15. ASHRAE SSPC 135 IT Working Group (2019). https://www.ashrae.org/File%20Library/Technical%20Resources/Bookstore/BACnet-SC-Whitepaper-v15_Final_20190521.pdf. Accessed 03 Apr 2020
8. Gerdes, R.M., Mina, M., Russell, S.F., Daniels, T.E.: Physical-layer identification of wired ethernet devices. IEEE Trans. Inf. Forensics Secur. **7**(4), 1339–1353 (2012)
9. Granzer, W., Praus, F., Kastner, W.: Security in building automation systems. IEEE Tran. Ind. Electron. **57**(11), 3622–3630 (2010)
10. Gray, P.R., Hurst, P.J., Lewis, S.H., Meyer, R.G.: Noise in integrated circuits (chap. 11). In: Analysis and Design of Analog Integrated Circuits, 5th edn., pp. 736–795. Wiley, January 2009. ISBN 978-0-470-24599-6
11. KNX on track for success again in 2019: Sector Coupling and IoT in focus. Press Release, KNX Association Cvba (2019). https://media.knx.org/feed/file/1050. Accessed 13 Feb 2020
12. Lackner, G., Payer, U., Teufl, P.: Combating wireless LAN MAC-layer address spoofing with fingerprinting methods. Int. J. Netw. Secur. **9**(2), 164–172 (2009)
13. Leach, T.: Implementing a BACnet network. ASHRAE J. **59**(3), 40–48 (2017)
14. Maes, R., Verbauwhede, I.: Physically unclonable functions: a study on the state of the art and future research directions. In: Sadeghi, A.R., Naccache, D. (eds.) Towards Hardware-Intrinsic Security. ISC, pp. 3–37. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-14452-3_1. ISBN 978-3-642-14452-3
15. Maximum data protection for smart buildings. Press Release, KNX Association Cvba (2017). https://media.knx.org/feed/file/918. Accessed 03 Apr 2020
16. Mundt, T., Krüger, F., Wollenberg, T.: Who refuses to wash hands? Privacy issues in modern house installation networks. In: 2012 Seventh International Conference on Broadband, Wireless Computing, Communication and Applications, pp. 271–277 (2012)
17. NCN5121. Transceiver for KNX Twisted Pair Networks. Rev. 2. Datasheet. Semiconductor Components Industries, LLC, August 2019. https://www.onsemi.com/pub/Collateral/NCN5121-D.PDF. Accessed 17 Feb 2020
18. PROFIBUS System Description. Technology and Application. 4.332. PROFIBUS Nutzerorganisation e. V. (PNO), April 2016. https://www.profibus.com/index.php?eID=dumpFile&t=f&f=52380&token=4868812e468cd5e71d2a07c7b3da955b47a8e10d Accessed 07 Feb 2020
19. Sokollik, F., Helm, P., Seela, R.: KNX für die Gebäudesystemtechnik in Wohnund Zweckbau. VDE Verlag GmbH (2017)

20. Szmulewicz, D.: Using MSP on KNX Systems. SWRA497. Application Report. Texas Instruments Incorporated. December 2015. http://www.ti.com/lit/an/swra497/swra497.pdf. Accessed 17 Feb 2020
21. Wang, X., Zhang, Y., Zhang, H., Wei, X., Wang, G.: Identification and authentication for wireless transmission security based on RF-DNA fingerprint. EURASIP J. Wirel. Commun. Netw. **2019**, 230 (2019). https://doi.org/10.1186/s13638-019-1544-8