

On False Data Injection Attack against Building Automation Systems

Michael Cash*, Christopher Morales-Gonzalez[†], Shan Wang^{††}, Xipeng Jin*, Alex Parlato*,
Jason Zhu[§], Qun Zhou Sun*, Xinwen Fu*[†]

* University of Central Florida. Email: {mcash001, xipeng_jin}@knights.ucf.edu, {alex.parlato, qz.sun}@ucf.edu

[†]Southeast University. Email: shanwangsec@gmail.com

[§]Email: jjasonzzhu@gmail.com

[†]University of Massachusetts Lowell. Email: christopher_moralesgonzalez@student.uml.edu, xinwen_fu@uml.edu

Abstract—KNX is one popular communication protocol for a building automation system (BAS). However, its lack of security makes it subject to a variety of attacks. We are the first to study the false data injection attack against a KNX based BAS. We design a man-in-the-middle (MITM) attack to change the data from a temperature sensor and inject false data into the BAS. We model a BAS and analyze the impact of the false data injection attack on the system in terms of energy cost. Since the MITM attack may disturb the KNX traffic, we design a machine learning (ML) based detection strategy to detect the false data injection attack using a novel feature based on the Jensen Shannon Divergence (JSD), which measures the similarity of KNX telegram inter-arrival time distributions with *attack* and with *no attack*. We perform real-world experiments and validate the presented false data injection attack and the ML based detection strategy. We also simulate a BAS, and show that the false data injection attack has a huge impact on the BAS in terms of power consumption. Our results show an increase in overall energy cost during a false data injection attack. Other the examined ML models, the Support Vector Machine (SVM) classifier achieved the best results with a 100% detection rate with our proposed similarity features compared to mean and variance related features.

I. INTRODUCTION

A building automation system (BAS) is one type of cyber-physical system (CPS), which manages and automates a building's mechanical and electrical systems such as heating, ventilation, and air conditioning (HVAC), sensors and actuators. Devices in a BAS follow a specific communication protocol, can be linked to a local area network (LAN) and may be reachable from the Internet. KNX is one popular BAS communication protocol with products market size of \$4.381 billion in 2019 and projected to reach \$10.15 billion by 2026 [1]. In this paper, we focus on the KNX based BAS.

A BAS may be subject to both cyber attacks and physical attacks. A cyber attack may be used to break the networks, pollute, and steal data, while a physical attack may be used to manipulate and damage physical components [2].

We are the first to study the false data injection attack against a KNX based BAS. Our major contributions can be summarized as follows. We explore physical attacks and cyber attacks to craft a man-in-the-middle (MITM) attack and inject false data to a KNX based BAS. In particular, two Raspberry Pis with KNX adaptors perform as the man in the middle, and relay and change data in KNX telegrams such as temperature

sensor data. We also model a HVAC system and analyze the impact of the fake data injection attack against the BAS in terms of energy cost.

We design a machine learning based detection strategy to detect the false data injection attack based on statistic features of the traffic, which is disturbed by the MITM attack. The MITM attack cannot be detected by injected false data, which is changed in a minor way. We design a novel feature which can measure the similarity of the probability distributions of the KNX telegram inter-arrival times with *attack* and with *no attack* based on Jensen Shannon Divergence (JSD). A ML model is trained with the JSD feature to detect the MITM attack and thus the false data injection attack.

We perform simulation and real-world experiments to validate the false data injection attack and the ML-based intrusion detection strategy. We simulate the HVAC system under two types of false data injection attacks, and show that the attacks have severe impacts on power consumption. Our experiments show that the proposed JSD feature works better than common statistics such as mean and variance which are sensitive to outliers inherently. With the SVM-based ML model with the JSD feature, the detection rate reaches 100% with short KNX telegram segments such as 5 minutes of data.

Responsible disclosure: We have reported all our findings to Siemens and its development teams. We hope the BAS industry puts more efforts on BAS security.

II. RELATED WORK

No existing work focuses on the false data injection attack against a KNX based BAS and analyses its impact with regards to energy costs while there are attacks against the BAS. Attackers can launch a replay attack against actuators by sniffing traffic messages, so as to manipulate actuators [3]. An adversary may use a hotel room's iPad to access the room's KNX network and identify addresses of KNX devices in hotel rooms. Then attackers can control the KNX devices remotely without the need of the hotel iPad [4]. Rios presents several password retrieval attacks to exploit and gain access to facility management systems [5]. A fuzzing tool is designed to discover vulnerabilities in the KNX protocol [6].

III. BACKGROUND

In this section, we introduce the communication mediums, address types and telegrams in the KNX protocol.

A. Communication Medium

The KNX protocol supports multiple physical communication mediums such as twisted pair (TP1), KNXnet/IP, radio frequency (RF), and powerline (PL) for connected devices to exchange information via a KNX network. The KNX network can be configured in either a tree, star, or line (daisy-chained) topology. With the TP1 backbone communication medium, the KNX network is a bus network. Any message that is sent onto the bus network will be received by every device on the network.

B. Address Type

The KNX protocol supports two general types of 16-bit addresses: individual addresses and group addresses, which can be used to access devices and data in a KNX network.

Individual Address. An individual address uniquely identifies a KNX device on a KNX network. It has three parts in the notation "area.line.deviceAddress". The individual address whose deviceAddress is 0 is reserved for a line coupler. If deviceAddress is not 0, the address represents a regular KNX device. Individual addresses allow devices to access other devices on the network. For example, when a device sends a *read* request, the source address in the request data frame contains the individual address of the device. In the response of the request, the destination address will be the individual address of the device which initially sent the request.

Group Address: A group address is a logical entity that is linked to a specific piece of data, called a KNX object, which a KNX device has to provide. A group address can be either 2 levels "main/sub", or 3 levels "main/middle/sub". Please note that the addresses "0/0" and "0/0/0" are reserved for the broadcast addresses. The linking of a KNX object to a group address allows the object to be read from or written to through the address. For example, when a device reads a KNX object, the destination address of the *read* request will be the group address which is linked to the KNX object.

C. Telegram

The data frames that are sent throughout a KNX network are known as *telegrams*. The structure of these telegrams will change dependent on both the communication medium and the programming mode in which the KNX network is configured in. A KNX network with the TP1 communication medium supports several types of programming modes such as standard mode (S-mode) and easy mode (E-mode). S-mode is the most common mode, in which devices and objects are manually added and configured. A variation of E-mode called logical tag extended mode (LTE-mode) is specifically developed for HVAC applications. In LTE-mode, the configuration of a device is automatically performed by a pre-programmed controller. Figure 1 shows the structures of standard telegrams in S-mode

and extended telegrams in LTE-mode. A specific bit in the *control field* of a telegram indicates the telegram type.

Standard Telegrams in S-mode. Standard telegrams in S-mode are the most common telegrams. The *control field* in a standard telegram as shown in Figure 1 contains a plethora of information such as the repetition status, priority level and telegram type, which will be used in the interpretation of the telegram. The *source address* details the origin of the telegram while the *destination address* specifies the desired recipient. The data link layer service data unit (*LSDU*) contains the information about the purpose of the telegram such as *groupRead* or *groupWrite*, which will read from or write to a group address respectively. In the case of a *groupWrite*, the *LSDU* will also contain the data that is to be written. The last field *checksum* is used to verify the integrity of the telegram.

Extended Telegrams in LTE-mode. Extended telegrams in LTE-mode are larger than standard telegrams due to the inclusion of the *extended control field* and a larger *LSDU* as shown in Figure 1.

The extended control field in an extended telegram can be divided into three parts: the address type, hop count, and extended frame format (EFF). The address type indicates the type of the destination address, which may be an individual address or a group address. The hop count limits the telegram distance to ensure that the telegram does not infinitely traverse the network. If the address type is a group address, then the EFF will describe how to interpret the group address. Otherwise, it will interpret the destination address as a normal individual address.

A LTE device is presumed to have multiple objects that can be of the same type. As a result, a larger *LSDU* is needed to access a specific object of the LTE device. When reading a KNX object of a LTE device, the object type (OT), object index (OI), and a property ID (PID) all must be provided. These are additional fields located in the *LSDU* field of an extended telegram. These three fields indicate a particular object and describe a property that a device desires to read. A tag address is also needed to read a KNX object of a LTE device. The tag address can be interpreted from a group address based on the EFF described previously.

IV. FALSE DATA INJECTION ATTACK

In this section, we first introduce physical attacks against KNX that will allow an attacker to gain access into a KNX network. We then present an eavesdropping attack, which can be used to learn detailed protocols of the KNX based BAS. Next, we present the false data injection attack via a man-in-the-middle (MITM) attack. Finally, we model a KNX based BAS, i.e. a HVAC system, and formally analyze the impact of false data injection attacks on the HVAC system in terms of energy cost.

A. Physical Attack

A KNX device such as a temperature sensor is often placed in publically accessible locations and thus is subject to a physical attack. For a KNX device utilizing the communication

Standard Telegram															
Octet 0 Control Field	Octet 1 Source Address	Octet 2 Source Address	Octet 3 Destination Address	Octet 4 Destination Address	Octet 5 LSDU/Length	Octet 6 LSDU	Octet 7 LSDU	Octet 8 LSDU	...	Octet 21 LSDU	Octet 22 (max) Checksum				
Extended Telegram															
Octet 0 Control Field	Octet 1 Extended Control Field	Octet 2 Source Address	Octet 3 Source Address	Octet 4 Destination Address	Octet 5 Destination Address	Octet 6 Length	Octet 7 LSDU	Octet 8 LSDU	Octet 9 (OT)	Octet 10 LSDU (OT)	Octet 11 LSDU (OI)	Octet 12 LSDU (PID)	...	Octet N LSDU	Octet N + 1 Checksum

Fig. 1. Standard vs Extended Telegrams

medium TP1, attackers can remove the device from the KNX network, i.e. a bus network, by removing its twisted pair cables from either the bus end or the device end. If the cables are disconnected from the device end, the same cables can be used by a malicious device to gain access to the KNX network. A KNX device may have four terminal ports and each port can be used to connect to another KNX device. If one of the four ports is available, a complete disconnect of the TP1 cables is not required. A malicious device can utilize the available port to connect and access to the KNX network. Then the malicious device can see any telegrams being transmitted in the network, or send its own telegrams to the network.

B. Eavesdropping Attack

In the eavesdropping attack, the attacker attaches their own device to the KNX network through the described physical attack above, and then passively listens on the bus for any telegrams that are sent. It provides a means of understanding how the KNX telegrams of a device are formed and the detailed communication protocol between devices. After performing a data dump, the attacker is presented with a plethora of information such as: source and destination addresses and individual/group addresses. In the case of extended telegrams, tag addresses, OTs, OIs and PIDs are also presented. With this gained information, the attacker can craft their own telegrams with respect to the victim device.

C. False Data Injection via MITM Attack

With the information gained from eavesdropping and properly crafted telegrams, attackers can deploy MITM attacks as shown in Figure 2 against a KNX based BAS with the goal of false data injection. Without loss of generality, we use the temperature sensor as a victim device in the MITM attack example. Two Raspberry Pis are used to perform the MITM attack. The left Pi is connected to the KNX network that the temperature sensor is connected to via a KNX adaptor, KNX Pi Hat. It then changes the data within the *LSDU* of received telegrams from the sensor, and forwards these changed telegrams to the right Pi via Ethernet or wirelessly. The right Pi is connected to the KNX network that the DXR2 controller is on, and sends the changed telegrams to the DXR2. The two Pis also forward telegrams generated by the DXR2 to the victim temperature sensor and other types of telegrams, excluding the temperature readings generated by the victim sensor, to the DXR2 without modification.

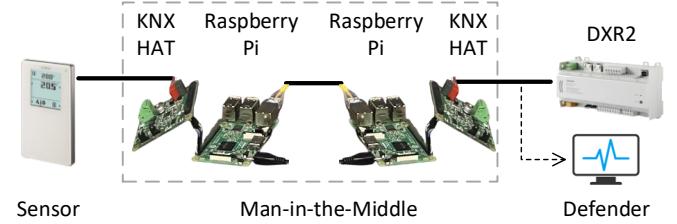


Fig. 2. MITM Attack against Sensors

Using two Raspberry Pis is necessary for stealthy data injection. One Pi can use only one KNX adaptor in out setup. If one Pi and one KNX adaptor are used in the MITM attack, the victim temperature sensor will be connected to the DXR2 directly through the KNX adaptor. Although the Pi can still gather the temperature sensor's telegrams, change them, and forward them to the DXR2, the DXR2 will receive both the original and changed telegrams; This would lead a defender monitoring the DXR2's KNX network to easily detect the attack. Using two Pis and two KNX adaptors separates the temperature sensor and the DXR2 into two different KNX networks and addresses the issue above.

D. Impact of False Data Injection

We can utilize the MITM attack to control various KNX devices such as a temperature sensor, and to inject fake data into a HVAC system. In this paper, we focus on injecting false room temperature values by manipulating a temperature sensor, which may incur severe energy costs. Next, we model a HVAC system. Then, we formalize the power consumption incurred by room temperature changes in terms of cooling in hot weather, so as to quantify the energy cost incurred by injecting false room temperature values.

We model a HVAC system as shown in Figure 3, which consists of a chiller that removes heat from water—the coolant, the chilled water pump, air handling unit (AHU) with three dampers and two fans which circulate the air, and the constant boundary conditions which represent a typical cooling tower system. The AHU serves one thermal zone such as a room. These components work together to maintain the room temperature at a predefined setpoint. Assume the room temperature increases. Then components in the HVAC system will work together to cool down the room temperature to maintain it at the setpoint, which incurs power consumption of fans, the pump and the chiller.

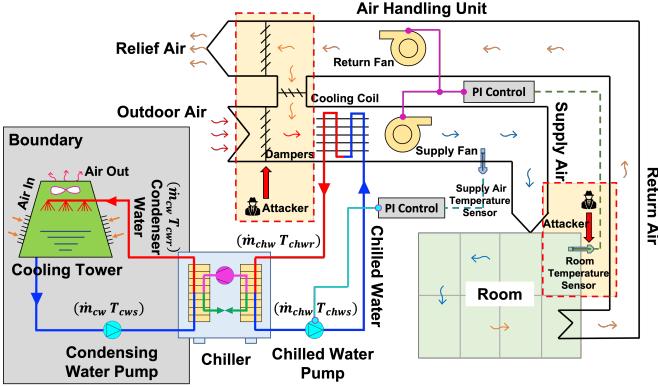


Fig. 3. HVAC system model

The total power consumption of the HVAC system is the summation of the power consumed by fans, the chilled water pump, and the chiller as follows [7]–[9].

$$\begin{aligned} P_{total}(t) = & P_{fan}(T_r(t)) \\ & + P_{pump}(T_{sa}(t)) \\ & + P_{chiller}(\dot{m}_{chw}(t)). \end{aligned} \quad (1)$$

$T_r(t)$ is the room temperature and may change with time. The power consumption of fans $P_{fan}(\cdot)$ increases with the room temperature. $T_{sa}(t)$ is the supply air temperature and increases with $T_r(t)$. The power consumption of the pump $P_{pump}(\cdot)$ increases with $T_{sa}(t)$ and thus the room temperature. $\dot{m}_{chw}(t)$ is the chilled water mass flow rate and increases with $T_{sa}(t)$. So the power consumption of the chiller $P_{chiller}(\dot{m}_{chw}(t))$ increases with the room temperature.

It can be observed from Equation (1) and discussion above that when the room temperature increases, the power consumption by fans, the chilled water pump, and the chiller increases and thus the total cooling power consumption increases. Therefore, when an attacker uses the false data injection attack to change the room temperature reading, the cooling system may be activated and the attack may incur energy waste and cost.

V. DEFENSE

In this section, we first present the detection problem definition, and show the traffic patterns with *attack* and with *no attack* are different. We then discuss selection of features and classifier for machine learning. At last, we introduce the attack detection with trained machine learning models.

A. Problem Definition for False Data Injection Attack Detection

The false data injection attack presented in Section IV utilizes the MITM attack. We carefully code the MITM attack to reduce potential delays in forwarding KNX telegrams by Pis. However, the attack may still change the KNX network traffic pattern due to network delays and jitters. In order to detect the false data injection, we can detect the MITM attack. For the detection, we train a machine learning (ML) model utilizing traffic pattern features. We assume that a defender collects telegram data and

performs detection at the DXR2 controller as shown in Figure 2.

The raw data we collect is telegrams with timestamps, from which we derive the telegram arrival time data series. From two consecutive telegrams arriving at t_i and t_{i+1} , we can derive the telegram inter-arrival time $I_i = t_{i+1} - t_i$. Therefore, from a telegram arrival time data series, we can derive a telegram inter-arrival data series.

We use a machine learning model to detect the false data injection attack. For training purposes, the defender collects two labeled telegram inter-arrival time training data sets : $A = \{I_{A,i} \mid 0 < i \leq l\}$, which is the telegram inter-arrival time data set when there is an attack; $B = \{I_{B,i} \mid 0 \leq i < m\}$, which is the telegram inter-arrival time data set when there is no attack. We define the *detection time window* t as the time duration of the telegram inter-arrival time data segment, which the defender collects for attack detection.

Problem Definition: Given training data A and B and a collected test telegram inter-arrival time segment \mathcal{T} of a time duration of t , does \mathcal{T} belong to *attack* or *no attack*?

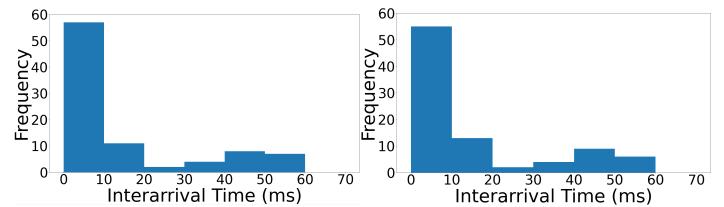


Fig. 4. Histogram of one segment in *attack* dataset A with 20min detection time window

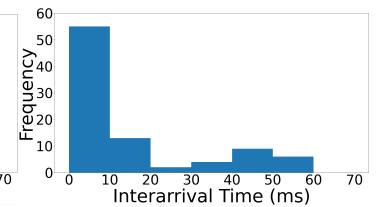


Fig. 5. Histogram of another segment in *attack* dataset A with 20min detection time window

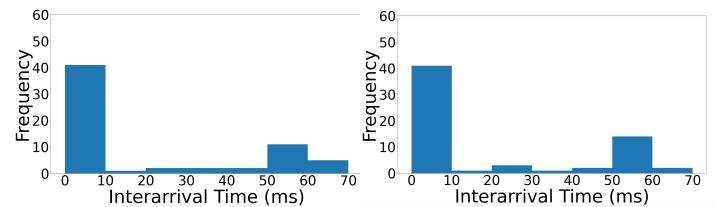


Fig. 6. Histogram of one segment in *no attack* dataset B with 20min detection time window

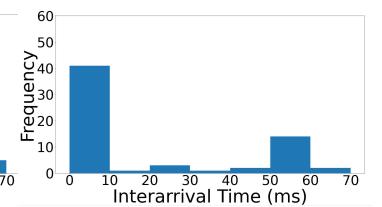


Fig. 7. Histogram of another segment in *no attack* dataset B with 20min detection time window

B. Intuition of Attack Detection

We want to observe the patterns of telegram inter-arrival times with attack and with no attack, and see if there is any difference. We draw the histograms of telegram inter-arrival time segments of a detection time window in training data sets A and B , and observe the distribution of the inter-arrival time. For a detection time window of 20 minutes, Figures 4 and 5 show the histograms of two segments in A with attack. Figures 6 and 7 show the histograms of two segments in B with no attack. It can be observed that the histograms from the same data set are similar while the histograms from A and B look different. For example, there are more small telegrams in the attack dataset histograms compared to the no attack dataset histograms.

inter-arrival times in A than in B . Therefore, the telegram inter-arrival time with attack has different patterns from the telegram inter-arrival time with no attack.

C. Feature Selection

We divide the telegram inter-arrival time training data sets A and B into segments of a time duration of t . For A , assume we obtain u segments $\{S_{A,i} \mid 0 \leq i < u\}$. For B , we derive v segments, $\{S_{B,j} \mid 0 \leq j < v\}$.

1) *Mean and Variance as Features*: Since Figures 4-7 show telegram inter-arrival times with attack and with no attack have different patterns, the first features we consider are the *mean*, *variance* of the telegram inter-arrival time segment and the vector (*mean*, *variance*). For our telegram inter-arrival time data, we derive the feature *mean* as follows. For each telegram inter-arrival time segment in A and B , we can derive its mean. Therefore, we derive the training data sets of the feature *mean* as follows, $\{\mathcal{M}_{A,i} \mid 0 \leq i < u\}$ and $\{\mathcal{M}_{B,j} \mid 0 \leq j < v\}$. We also derive the mean for the test data \mathcal{T} as $\mathcal{M}_{\mathcal{T}}$. Similarly, we can derive the training data sets of the feature *variance* as $\{\mathcal{V}_{A,i} \mid 0 \leq i < u\}$ and $\{\mathcal{V}_{B,j} \mid 0 \leq j < v\}$ and the variance for the test data \mathcal{T} as $\mathcal{V}_{\mathcal{T}}$. For the feature (*mean*, *variance*), we then have $\{(\mathcal{M}_{A,i}, \mathcal{V}_{A,i}) \mid 0 \leq i < u\}$, $\{(\mathcal{M}_{B,j}, \mathcal{V}_{B,j}) \mid 0 \leq j < v\}$ and $(\mathcal{M}_{\mathcal{T}}, \mathcal{V}_{\mathcal{T}})$.

2) *Probability Distribution Similarity as Feature*: Since mean and variance are inherently sensitive to outliers and not robust as features, we want to find a feature based on the probability distribution given that the outliers have a small probability to occur and features considering probabilities are more robust. We propose to use the Jensen Shannon Divergence (JSD) to measure the similarity between probability distributions of telegram inter-arrival time segments.

We derive the JSD similarity of two telegram inter-arrival time segments as follows. From the two segments, we can derive their corresponding telegram inter-arrival time distributions P and Q . \mathcal{X} is the range of the telegram inter-arrival time.

$$KL(P||Q) = \sum_{x \in \mathcal{X}} P(x) \log \frac{P(x)}{Q(x)} \quad (2)$$

$$JSD(P||Q) = \frac{1}{2}KL(P||M) + \frac{1}{2}KL(Q||M), \quad (3)$$

where $M = \frac{1}{2}(P + Q)$, and $JSD \in [0, 1]$. The closer to 0 JSD is, the more similar these two segments are.

We can now derive JSD features for our training data sets A and B , and the test data \mathcal{T} .

- For each segment $S_{A,i}$ ($0 \leq i < u$) of A , we measure its JSD similarity with all segments of B and derive a feature vector of v elements as follows

$$(JSD(P_{A,i}||P_{B,0}), \dots, JSD(P_{A,i}||P_{B,v-1})). \quad (4)$$

Therefore, we have u feature vectors as training data. Those feature vectors are the *attack* feature vectors, which measure the similarity between *attack* data and *no attack* data.

- We now derive the *no attack* feature vectors for training. For each segment $S_{B,j}$ ($0 \leq j < v$) of B , we measure its JSD similarity with all segments of B and derive a feature vector of v elements as follows

$$(JSD(P_{B,j}||P_{B,0}), \dots, JSD(P_{B,j}||P_{B,v-1})). \quad (5)$$

Therefore, we have v *no attack* feature vectors.

- We derive the feature vector for test data \mathcal{T} as follows

$$(JSD(P_{\mathcal{T}}||P_{B,0}), \dots, JSD(P_{\mathcal{T}}||P_{B,v-1})). \quad (6)$$

D. Classifier

For each candidate feature including *mean*, *variance*, (*mean*, *variance*) and JSD similarity, we have derived two training data sets of feature vectors and the test data feature vector. For simplicity, we denote *mean* and *variance* as single-element vectors. Therefore, for each type of feature vector, we can train a classification model to classify the test data as *attack* or *no attack*. In our experiments, we use 70% of the data we collect for the training purpose and the other 30% as test data. We label both the training data and the test data with a class based on the ground truth.

Two classic ML algorithms, i.e. decision tree and support vector machine (SVM), are used to train binary classification models. The decision tree classifier is a tree-like model of decisions. Each internal node is a test on a feature in a feature vector, and each branch of this internal node is an output of the test. A leaf node is a *no attack* class or an *attack* class. Given an input feature vector, the decision tree model finds a path from the root to a leaf node based on the value of each feature, and uses the class of the leaf node as the output. The SVM based classifier is a function which splits a space into parts with different classes. Given an input feature vector, the function maps the feature vector to a subspace and outputs the class of the subspace.

E. Attack Detection

We use a trained ML model to detect the MITM attack, thus the false data injection attack. For real-time detection, we can just collect a telegram arrival time series in a period of t that the ML model uses, calculate its inter-arrival time segment \mathcal{T} , derive the feature vector for \mathcal{T} per the features that the ML model uses, and then input the feature vector to the trained ML model which will output a class for \mathcal{T} . We use the feature vectors of segments in the testing data set to evaluate the detection rate. By comparing the output class with the true class label of a testing feature vector, we can know if the classification, i.e. attack detection, is correct, and obtain the detection rate. In our experiments, we set different detection time windows t and evaluate the detection rate of different ML models with different features versus t .

VI. EVALUATION

In this section, we first present the experiment setup. Then we show the impact of the false data injection attack on the HVAC system. Finally, we demonstrate the effectiveness of the proposed ML-based defense scheme.

A. Experiment Setup

Figure 8 shows the experiment setup using a Siemens BACnet Field Panel located at the University of Central Florida. We use the Siemens DXR2.E12P-102B (DXR2) BAS controller and the Siemens QMX3.P74B-1WSB Room Operator Unit, which can work as a KNX temperature sensor in our experiments. The KNX device data from the DXR2 can be viewed and monitored in real time by the Siemens Desigo CC building management software platform. We use two Raspberry Pi 3s with 16 GB of flash memory running Raspbian OS version 10-Buster. We use two KNX Raspberry Pi HATs (PiHAT), each of which is connected to a Pi's universal asynchronous receiver/transmitter (UART) on-board pins using 4 Dupont wires.

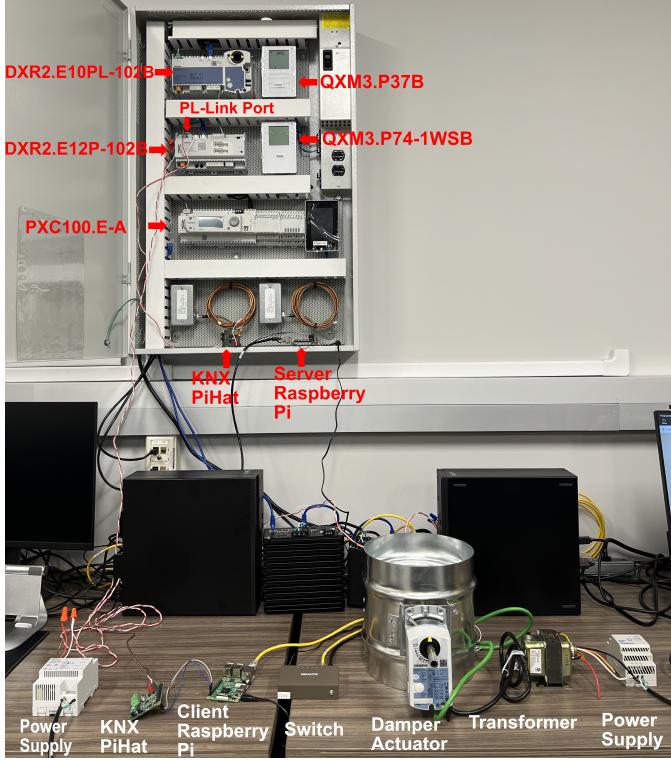


Fig. 8. Experiment setup for MITM attack against a temperature sensor

We use twisted pair (TP1) cables to connect one Pi HAT to the DRX2, forming one KNX network, and connect the second Pi HAT to the temperature sensor to form another KNX network. An Ethernet switch is used to connect the two Pis together so that the two Pis may communicate with each other to perform the MITM attack. We use the Calimero Java library for KNX to construct telegrams, perform eavesdropping and forward telegrams.

We validate the MITM attack against the temperature sensor as shown in Figure 8. Please note that the MITM attack against the damper is also performed but not presented in this paper.

B. Attack Impact on HVAC System

We use a simulation tool named *Dymola* and the Modelica buildings library [10] to simulate the HVAC system model

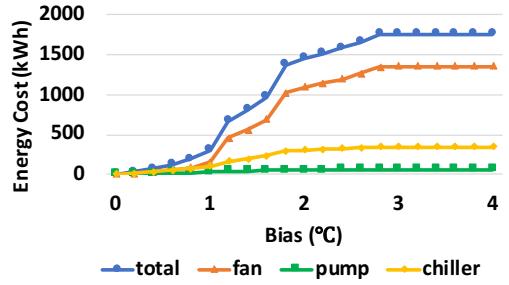


Fig. 9. Additional energy cost incurred by Attack-i

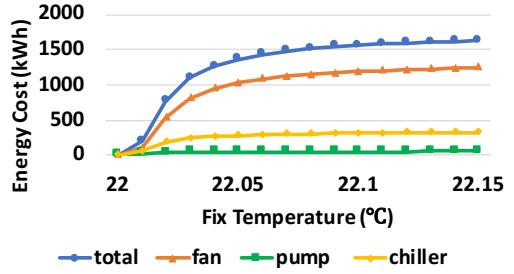


Fig. 10. Additional energy cost incurred by Attack-ii

as introduced in Section IV-D and mimic the real-time attacks to evaluate the impact of the false data injection attack.

Simulation Settings. In the simulation, a built-in solar radiation module of *Dymola* is used to simulate solar radiation (i.e., sun) influence to the building. We use the TMY3 weather data of Orlando, FL on July 29, 2005 to obtain a realistic ambient (i.e., surrounding) air temperature, humidity and wind speed. The room temperature setpoint, the supply air temperature setpoint and the supply water temperature setpoint are 22°C, 14°C and 6°C respectively. The damper opening level is set as 30%. The room size is 29,750 cubic meters. In the simulation, the outside temperature changes along with time, which affects the room temperature and incurs energy cost. Without attack, the energy cost incurred by the outside temperature changes is 1,303.70 kWh from 7:00 am to 7:00 pm. We use the 1,303.70 kWh as the energy cost benchmark. We evaluate two types of false data injection attacks as follows. (Attack-i) Constantly adding a fixed bias such as 1°C to the original room temperature sensor reading. (Attack-ii) Overwriting the original room temperature sensor reading to a constant false value such as 22.005°C which is higher than the setpoint. We launch each attack at 7:00 am and cease the attack after 12 hours.

Energy Cost. Figures 9 and 10 show the additional energy cost incurred by the two false data injection attacks in comparison to the energy cost benchmark 1,303.70 kWh without attack. For each attack strategy, we measure the total additional energy cost and the additional energy cost of the fans, the chilled water pump and the chiller. It can be seen that the total power consumption has risen significantly. Our proposed attack strategies can have a huge impact on HVAC

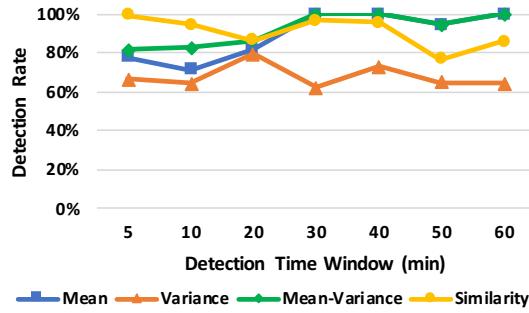


Fig. 11. Detection Rates vs. Detection Time Window for J48

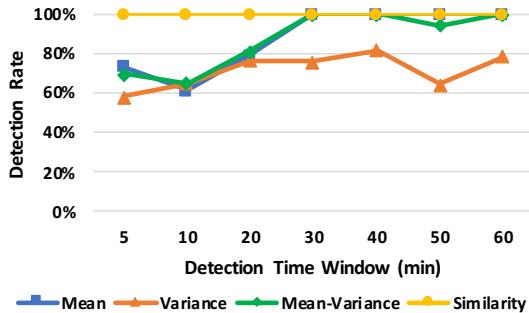


Fig. 12. Detection Rates vs. Detection Time Window for SMO SVM

system power consumption.

C. Defense Effectiveness

We use a popular ML tool Weka V3.8.6 to train the ML model and evaluate the effectiveness of the proposed defense scheme on the false data injection attack against the temperature sensor. Two classic ML algorithms, Weka J48 decision tree and Weka SMO support vector machines (SVM), are used. We gathered the attack dataset *A* and baseline dataset *B* over a 24-hour period each from our real-world testbed. We set the detection time window *t* as 5min, 10min, 20min, 30min, 40min, 50min, and 60min, respectively. The detection rate is defined as the overall accuracy of recognizing feature vectors labeled as *attack* and *no attack* in the testing dataset.

Figure 11 shows the detection rate for each detection time window with the J48 decision tree based model. The proposed probability distribution similarity feature works better than mean, variance and (mean, variance) features when the detection time window is less than or equals to 20 min. The defender who detects the attack would prefer a small detection time window so as to discover the attack quickly. An attack may not last long either. In the case of the SVM-based ML model, the probability distribution similarity feature always outperforms other features, and its corresponding detection rate reaches 100% with each detection time window as shown in Figure 12. It can be observed that our defense scheme is effective, and the proposed probability distribution similarity feature works better than the mean and variance related features.

VII. CONCLUSION

In this paper we demonstrate the feasibility and impact of a false data injection attack against KNX based building automation systems. In such an attack, an attacker may physically remove a KNX sensor, and launch a Man-In-the-Middle attack with the intent to inject false data into the BAS. We carefully analyze the impact of false data injection attacks on HVAC systems, and quantify the additional energy cost incurred by the false data injection attack by simulating two attack strategies. We are the first to study the false data injection attack against a BAS and its impact. Machine learning-based strategies are presented to detect the false data injection attack based on features of the telegram inter-arrival time. The SVM classifier can achieve a detection rate of 100% using the proposed probability distribution similarity feature as compared to the mean and variance related features with small detection time windows.

ACKNOWLEDGMENT

This research was supported in part by US Department of Energy (DOE) Award DE-EE0009152, and US National Science Foundation (NSF) Awards 1931871 and 1915780. Any opinions, findings, conclusions, and recommendations in this paper are those of the authors and do not necessarily reflect the views of the funding agencies.

REFERENCES

- [1] “Global knx products market insights and forecast to 2027,” <https://www.360marketupdates.com/global-knx-products-market-17971150>, Apr 2021, accessed on Oct. 23, 2021.
- [2] G. D. Wyss, P. E. Sholander, J. L. Darby, and J. M. Phelan, “Identifying and defeating blended cyber-physical security threats,” <https://www.osti.gov/servlets/purl/1427000>, Jan 2007.
- [3] A. Antonini, F. Maggi, and S. Zanero, “A Practical Attack Against a KNX-based Building Automation System,” in *Proceedings of the 2nd International Symposium on ICS & SCADA Cyber Security Research 2014*, ser. ICS-CSR 2014. UK: BCS, September 2014, pp. 53–60. [Online]. Available: <https://doi.org/10.14236/ewic/ics-csr2014.7>
- [4] J. Molina, “Learn how to control every room at a luxury hotel remotely: The dangers of insecure home automation deployment,” *BlackHat USA*, 2014.
- [5] B. Rios, “Exploiting access control and facility management systems,” in *Blackhat Asia*, 2014.
- [6] C. Vacherot, “Sneak into buildings with KNXnet/IP,” in *Sneak into buildings with KNXnet/IP*, Lyon, France, Nov. 2020, présenté aux conférences :- GreHack 2020- Defcon 2021. [Online]. Available: <https://hal.archives-ouvertes.fr/hal-03022310>
- [7] M. Stewart, *Surface Production Operations: Volume IV: Pumps and Compressors*. Gulf Professional Publishing, 2018.
- [8] W. Wang, G. Tian, Q. Z. Sun, and H. Liu, “A control framework to enable a commercial building hvac system for energy and regulation market signal tracking,” *IEEE Transactions on Power Systems*, 2022.
- [9] M. Hydeman and K. L. Gillespie Jr, “Tools and techniques to calibrate electric chiller component models/discussion,” *ASHRAE transactions*, vol. 108, p. 733, 2002.
- [10] M. Wetter, W. Zuo, T. S. Nouidui, and X. Pang, “Modelica buildings library,” *Journal of Building Performance Simulation*, vol. 7, no. 4, pp. 253–270, 2014.