Classification of Device Behaviour in Internet of Things Infrastructures: Towards Distinguishing the Abnormal From Security Threats

Roman Ferrando
ThingBOOK.IO, roman.f@thingbook.io

Paul Stacey
Institute of Technology, Blanchardstown, paul.stacey@tudublin.ie

Recommended Citation

Ferrando, R., Stacey, P. (2017) Classification of Device Behaviour in Internet of Things Infrastructures: towards distinguishing the abnormal from security threats. *International Conference on Internet of Things and Machine Learning, Liverpool, October 2017*. doi:10.1145/3109761.3109791

Classification of Device Behaviour in Internet of Things Infrastructures

Towards Distinguishing the Abnormal from Security Threats

Roman Ferrando

ThingBook.io
Dublin, Ireland
Roman.f@thingbook.io

Paul Stacey

Department of Engineering
Institute of Technology Blanchardstown
Dublin, Ireland
paul.stacey@itb.ie

ABSTRACT

Increasingly Internet of Things (IoT) devices are being woven into the fabric of our physical world. With this rapidly expanding pervasive deployment of IoT devices, and supporting infrastructure, we are fast approaching the point where the problem of IoT based cyber-security attacks is a serious threat to industrial operations, business activity and social interactions that leverage IoT technologies. The number of threats and successful attacks against connected systems using IoT devices and services are increasing. The Internet of Things has several characteristics that present technological challenges to traditional cyber-security techniques.

The Internet of Things requires a novel and dynamic security paradigm. This paper describes the challenges of securing the Internet of Things. A discussion detailing the state-of-the-art of IoT security is presented. A novel approach to security detection using streaming data analytics to classify and detect security threats in their early stages is proposed. Implementation methodologies and results of ongoing work to realise this new IoT cyber-security technique for threat detection are presented.

Keywords

Internet of Things, Cyber Security, Streaming Analytics, Device Behaviour Classification, Abnormal Behaviour Detection

1 INTRODUCTION

The recent emergence of the Internet of Things (IoT) as a novel and powerful platform for services and decision-making [1, 2], have made it one of the fastest growing technologies Today [3]. This new disruptive paradigm of a pervasive physically connected world, will have a huge impact on social interactions, business and industrial activities. It is predicted that the IoT will gradually permeate all aspects of modern human life [4]. In this new connected world, the fine-grained monitoring of activity and processes involves the storage of vast amounts of sensitive data and information about citizens, organizations and their activities.

Given the growing pervasiveness of connected sensing enabled devices, and consequently the increasing accessibility of highly sensitive data and information, the need for robust security has never been greater [5,6]. The security challenges presented by the deployment of connected resource-constrained devices is well understood, and has been the focus of research for many years [7]. However, off-the-shelf and deployment ready practical solutions to securing constrained and connected sensing devices are not readily available. Prior to the recent IoT revolution, there has not existed an urgency for the employment of robust security measures in similar type devices and systems. This lack of urgency has bred a culture of poor security practices within IoT predecessors such as wireless sensor networks (WSN) and SensorWebs, and consequently their descendant IoT systems.

Already we are beginning to see the net result of poor security in already deployed IoT networks. The Dyn cyberattack on the 21st of October 2016 [8] saw a series of massive Distributed Denial of Service (DDoS) attacks. These attacks were performed using a Mirai-bot based botnet [9]. It is estimated that the attackers used more than 100,000 infected IoT end-points to generate traffic rates of up to 1.2 Tbps to achieve the DDoS attack [10]. This attack highlighted a new urgency for more sophisticated protection systems to secure IoT networks and systems against threats and vulnerabilities. These attacks also highlighted the complacency the IoT community has employed when considering security during IoT deployments.

The Mirai botnet attack has brought the issue of IoT security into the public domain. However, the number of threats had been rising daily [11] prior to the Mirai botnet attack. Also of note is the increase in sophistication of the methods and tools employed by an ever-increasing number of would be attackers [12,13]. These threats now

raise serious questions as to the real dangers faced by individuals and organizations when using IoT technologies. Failure to act could see the vision of a connected world severely limited and represent a missed opportunity for new business models and revenue streams.

In this paper, we present the initial results of ongoing work to address the security challenges presented by an IoT paradigm. Leveraging innovative streaming analytical techniques, we show how detecting events in *traffic feature distributions* can allow the classification of abnormal behaviour within an IoT network.

This paper is organised as follows: Firstly, we present a discussion leading to a definition of an Internet of Things system (Section 2). Section 2.3 provides a brief overview of security considerations for IoT. Section 3 provides a review of the current techniques within the field of unsupervised network anomaly detection, with the state of the art presented in Section 3.1. A novel approach applying these broad techniques to contribute to IoT-appropriate security detection and resolution approaches is presented in Section 4. In Section 5 we draw conclusions based on the results of Section 4 and detail further work.

2   THE INTERNET OF THINGS

2.1   Defining the Internet of Things

The Internet Architecture Board (IAB) states in RFC 7452 the following:

"The term Internet of Things (IoT) denotes a trend where a large number of embedded devices employ communication services offered by the Internet protocols. Many of these devices, often called smart objects, are not directly operated by humans, but exist as components in buildings of vehicles, or are spread out in the environment" [14].

The Internet Engineering Task Force (IETF) notes that a smart object will typically have significant constraints in terms of power supply, memory, communication bandwidth and on-board processing power [15]. References [15-17] all note that the interconnection of the physical world with the virtual world is the focus of IoT specification.

Generally, in the literature what is found are non-contradictory definitions of IoT. However, definition attempts are somewhat high-level and abstract to ensure applicability to diverse use-cases. This disparity among a definition for IoT can confuse any discourse amongst IoT interest groups.   Similar obstacles to meaningful discussions were apparent during the emergence of the concepts of *net neutrality* and *cloud computing*, which hindered community consensus on associated topics of interest [18]. In any case, the arrival of a global consensus on an IoT definition will follow the habitual path of standardization; this is beginning to emerge through the work of International standards organizations such as ISO [19].

Here, when referring to the "Internet of Things" and "IoT" we adopt the following broad definition. IoT refers to:

"the extension of network connectivity and computing capability to objects, devices, sensors and items not normally considered to be computers. These "smart objects" require minimal human intervention to generate, exchange and consume data; they often feature connectivity to remote data collection, analysis and management capabilities" [20].

While many models of IoT include non-IP data flows and thus do not route data via the Internet, the authors assume that any data generated or processed from IoT/smart nodes will pass through an IoT bridge and be hosted along an IP-based publicly accessible network. For example, we assume that a 6LoWPAN Border Router (6LBR) or equivalent would be an intrinsic part of any IoT system.

2.2   IoT Security

IoT network operators and cloud service providers host network flows, which exhibit a myriad of "unusual" behaviors and events. Within the bounds of these unusual events may lie furtive activities with malicious objectives. Eliciting the event patterns of a maligned activity is not a trivial task. The volatile nature of the IoT environment makes discovery difficult. Within an IoT system there may be a high degree of volatile behavior. This volatility merely represents the digital artifacts of a chaotic physical world augmented with sensing and communication technologies. Human-behaviour tends to exhibit volatile behavior, thus a resultant and continuous digital stream from consumers, smart-things and machines may capture a new and normal behaviour as a seemingly unusual event.

While leveraging the ability to analyse IoT behaviour from network traffic, the challenge in an IoT environment is sorting the abnormal (but valid behavior) from that of a security threat. To ensure the correct response to behavioral changes, a sophisticated and dynamic behaviour classification regime must be employed to elicit the true nature of IoT data streams and resultant network flows.

It is important to note that an IoT security environment is not that different from any other non-constrained network. Consequently, many good lessons can be learned from traditional approaches and used as the basis for an IoT relevant and appropriate solution.

3   ANOMALY CLASSIFICATION AND DEVICE DISCOVERY

Given the diversity of IoT sensors, devices and resulting data-streams, the principal challenge in automatically detecting and classifying anomalies is to un-restrict events and activities and rely on the ability to mine these events and identify anomalies that are considered a security threat. Such anomalies can span a vast range of events: from network abuse (examples include denial-of-service attacks, scans, worms) to equipment failures (such as outages) to unusual customer behaviour (e.g., sudden changes in demand, flash crowds, high volume flows), and even to new, previously unknown events.

In the field of anomaly behaviour detection applied to the IoT space, two additional and considerable complications have to be considered. Firstly, IoT covers a huge range of different devices all forming an ecosystem where the line between normal and abnormal is usually blurred. Secondly, anomalies are a moving target. It is difficult to precisely and permanently define a set of anomalies within IoT network behavior, especially in the case of malicious anomalies. New network anomalies will continue to arise over time; so, an anomaly detection system should avoid being restricted to any predefined set of anomalies.

The goal of this paper is to contribute towards a system that fulfills these criteria. We seek methods that can classify sensor traffic in the IoT space and detect a diverse and general set of network anomalies, and to do so with a high detection rate and a low false alarm rate. Furthermore, rather than classifying anomalies into a set of rigid and

static classes (defined historically) we seek to evaluate the anomalies from the data following a fuzzy unsupervised approach that allows the discovery of a comparative similarity index between new and already identified abnormalities.

We base our work on the observation that despite their diversity, most traffic anomalies share a common characteristic: *they induce a change in the distribution of the generated network traffic fields*. Our hypothesis is that examining distributions of network traffic features yields considerable diagnostic power in both the detection and classification of a large set of anomalies.

Next, we present an overview and background to the state-of-art and current trends in the relevant areas that inform our methodologies.

3.1 State of the Art and Related Work

In recent times, the use of IP network flows based anomaly detection has been gaining considerable attention, and has been the focus of increased study. The explosion of data flows and speeds across networks has driven this increased interest. Where previously, individual packet inspection would occur in real-time to aid detection of anomalies, this becomes unwieldly at high data rates. A flow based approach has emerged as a scalable and timely approach. The flow based approach does not seek to replace packet inspection, but work as a complementary approach for anomaly detection. [21] proposes a denial-of-service attack detection architecture for IoT systems. In [21] a packet inspection methodology is employed. Our work can conceivably complement packet inspection approaches.

Arising from recent work, successful implementations of anomaly detection, based on detecting deviations from what is considered a "normal-state" have emerged. Here we present an overview of the ongoing work in this field, referred to as *unsupervised machine learning for network anomaly detection*.

3.2 Network Based Anomaly Detection

The field of anomaly detection within IP networks can be broken down into two main categories:

- Knowledge based detection systems or *supervised* detection systems.
- Knowledge independent or *unsupervised* detection systems.

The focus of this paper is the latter; *unsupervised*. Other approaches prevalent within the network anomaly detection community are not considered here; due to their reliance on previous or historical results of detection activities. We would argue historical data is of limited use, or is not readily available in an IoT context. The nature of IoT systems mandate the need to directly monitor and measure traffic, and react in a dynamic way. In [22] Raza et al. present a hybrid approach to intrusion detection within IoT systems. Their work attempts to balance a signature approach with an anomaly approach. We seek to focus solely on an unsupervised anomaly approach as the need to store historical data for signature analysis is not practical in constrained IoT networks.

Zang et al. [23] present a unified anomaly detection framework for network anomography. They propose to separate anomaly detection into two categories; systems using *temporal* correlation methods, and systems using *spatial* correlation methods to identify normal traffic. In this paper, we adopt Zang et al.'s category definitions within an unsupervised approach; where unknown anomalous behaviour rather than particular signatures is our focus.

3.2.1 Temporal correlation methods describe those techniques where *time* is the main driver in the analytics process. In this technique, a point-in-time represents a reference point for all analysis operations. Anomalous traffic can be separated by performing time-series based temporal analysis on the traffic source. Four types of temporal analysis are identified in [23], which can be split into two groupings: *time-series analysis* and *continuous data observations*.

Time-series analysis associates anomalies with a deviation from a predicted behaviours classifier, and is calculated using a distance metric from that classifier. Two models are usually employed: Auto-Regressive Integrated Moving Average (ARIMA) and deltoids [24]. ARIMA models include Exponentially Weighted Moving Average (EWMA) and linear exponential smoothing.

3.2.2 Spatial Correlation Methods. In spatial correlation methods, data elements in high dimensional data sets such as the network load observations usually have dependencies. The intrinsic dependency structure among the data elements can thus be exploited for filtering anomalous behaviour by discovering data points that violate the normal dependency structure [23]. Here the use of entropy to sum up the feature distribution of networks is employed. By using unsupervised learning, it is shown that anomalies can be clustered to form anomaly classes or cluster vector definitions. The metrics or features successively used in these papers are: byte counts in [25], packets counts, byte counts and IP flow counts in [26] and entropy values for distributions of several features (source IP address, destination IP address, source port, and destination port) in [27]. In [28] it is shown that entropy based approaches are suitable and effective at detecting modern-botnet attacks such as the Mirai-botnet.

What is clear from the literature is that a multi-detector approach is the main conclusion of many of the experiments documented in the literature. Major advancements in supervised approaches using a combination approach have been reported. However, the more complex challenge of unsupervised detection systems using machine learning remains an open question [29].

This work seeks to further the complex field of unsupervised learning to allow for the enabling of appropriate feature distribution clustering and analysis approaches for an IoT eco-systems. Unsupervised advocates work on the assumption that abnormal traffic is fundamentally different to the normal traffic structures. However, the volatile nature of IoT means that this is not necessarily the case. The diagnostic approach described here is intended to ultimately overcome this inherent challenge of using unsupervised approaches within IoT ecosystems.

4 EXPERIMENT

The problem of defining and analysing anomalies purely with data observations (unsupervised learning) and the absence of previously characterized knowledge, is the focus of much attention within the scientific community. The unsupervised school, (the one this work belongs to) works on the assumption that abnormal traffic is fundamentally different to the normal traffic structures. It is assumed that by studying these differences, abnormalities and new knowledge can be discovered. However, the reality of the world surrounding us is rarely pure and never simple. The consideration of abnormality is tied to many circumstances, as the time or the season, that put captivating challenges to those in the search of creating a knowledge acquisition

devices. In the IoT ecosystem, the volatility of the reality push those challenges to the next level. The diagnostic system proposed here is developed to identify abrupt changes in the individual features and in the dependencies of those variables.

The experiment described here uses a *spatio-temporal* methodology to characterise network behaviours. Once characterised, anomalous behaviours can be identified by calculating a similarity/distance metric to previously identified behaviours (e.g. attacks, intrusions or malfunctioning machinery). The thesis under investigation is: through the monitoring of the entropy of variables associated with certain network traffic features, combined with a modified dispersion coefficient for numerical variables, it is possible to generate rich 2D models that capture the nature of the network behaviour, referred to as behavioural shapes. These behavioural shapes contain verbose visual descriptors of individual feature's behaviour and the dependencies that exist between them. We propose that any connected sensor, smart-thing or community of things network flow behaviour can be represented using 2D models/behavioural shapes.

A sliding-window approach is employed to analyse the temporal aspect of our methodology. At each time unit (Ti), a behavioural shape is produced, calculated using the network data within the last n time units.

In this experiment, we configured our system to run in intervals of 30 seconds, however, our system is flexible enough to work with different time horizons. Next, the measurements of each window time are scaled to unit norm to focus on the dependencies of its dispersions rather than its infinite value. The problem is then restricted to find the windows time describing shapes with an abnormal figure compared to the rest. The shape form, area and position in the 2d plane will be defined by the dispersion values and the dependency between the values.

4.1 Behavioral Shape Calculation

The central idea of our work is the characterization of any network behaviour in 2D shapes. Fig 1 shows how a behavioural shape is constructed. A normalisation process is employed to focus on dispersion dependencies.
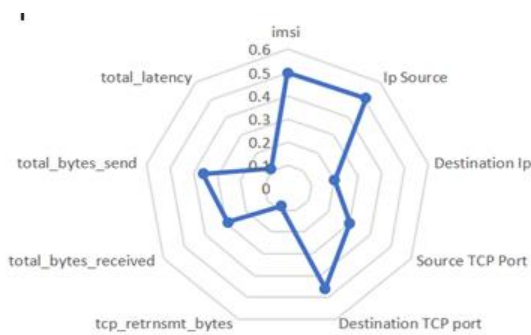
Figure 1: Behavioral Shape. The behavioral shape shown is constructed through a mix of entropy and dispersion measurements. For example, the IP Source entropy in this case is 0.51, the total_bytes_received dispersion coefficient is 0.29. In the example T_0 = 30seconds and normal operation is being observed. The shape form's area and position in the 2D plane is defined by the dispersion values and the dependency between the values.

4.1.1 Variable distribution study. In this study, the data from a telecom mobile operator (4G network) with IoT devices connected to the network is used. Data was read from the SGi interface at the core network (EPC). All traffic coming or going to non-IoT devices were filtered and removed. IoT traffic is not the only traffic flowing on the network therefore, It is safe to assume that changes in the network circumstances (e.g. Equipment malfunctioning, or attacks coming from non IoT-devices) can impact on the behaviour and performance of IoT devices. The following fields, each representing a traffic feature were monitored: International Mobile Subscriber Identity (IMSI), IP Source, IP destination, source TCP port, destination TCP port, tcp_retrnsmt_bytes, total_bytes_received, total_bytes_send and total_latency.

These fields are only a sample of the fields that could be monitored. Also, data acquired is only a sample of the total data. For this experiment, we measured traffic for 4 hours per day over a 7-day period. It should be noted that it was not possible to generate security attacks as a real live network was used. We therefore conjecture that real security attacks could be detected using the proposed method. However, real anomalies were found. Those anomalies could be the result of sensors malfunctioning, sensors software updates or simply, IoT device attacks on a small scale. Lakhina et al. note that:

"The distribution of traffic features is a high-dimensional object and so can be difficult to work with directly. However, we can observe that in most cases, one can extract very useful information from the degree of dispersal or concentration of the distribution the specific variables changing its distribution at the same time, compared with those which remain stable" [27].

Lakhina et al. found in [27] that in some cases, the fact that a group of features were dispersed while others were concentrated is a strong indicator, which should be useful both for detecting an anomaly and identifying it once it has been detected.

4.1.2 Entropy. The formula for Entropy is defined in (1) below.

$$H(x) = \sum_{i=1}^{N} (\frac{ni}{S}) * \log(\frac{ni}{S}) \qquad (1)$$

Where $S = \sum_{i=1}^{N} ni$, and is the total number of observations. The value of sample entropy lies in range $[0, \log(N)]$. The rate of entropy is lesser when the class distribution is pure (poor diversity). The rate of entropy is larger when the class distribution is impure (large diversity). The entropy shows its minimum value 0 when all the items of a feature (e.g IP address or port address) are the same; and its maximum value $\log(N)$, when all the items are different.

Here entropy is used as a convenient summarily statistic for a distribution's tendency in categorical variables to be concentrated or dispersed. We use this metric to build our behavioural shapes. It is important to note that entropy is not the only metric that captures a distribution's concentration or dispersal on categorical variables. However, we have explored other metrics and find that entropy works well for our objectives.

In this work, we used the entropy of feature distributions calculated from network traces counts. However, the temporal approach presented in this paper has some implications on the usage of the entropy calculations. As we propose a fixed temporal length for our sliding window and because the

network will experience network traffic volume fluctuations throughout the day, the value of N (the total number of distinct values seen in a window time) will change accordingly. As the entropy lies in range $[0, \log(N)]$, the value of N will impact the entropy value.

The implications of this effect on our approach are minimal. As we scale each value to the unit norm, our approach focuses on the relationship between entropies rather than their absolute values. We can thus guarantee that similar behaviours will appear to be near to each other in this entropy space regardless of the volume of the traffic.

4.1.3 Dispersion coefficient. In this experiment, we proposed a modified dispersion metric applied to numerical variables. The metric proposed has been modified to the range [0, 1] and is shown in (2) below.

$$D(x) = 0.01 * \sin^{-1}\left(\frac{Avg}{\sqrt{Std^2 + Avg^2}}\right) \qquad (2)$$

Where *Std* is the standard deviation of the sample, *Avg* is the average of the sample the metric proposed here calculates the angle created by the standard deviation and the average of the sample. The bigger the angle, the smaller the standard deviation in respect to the average, and therefore, less disperse the sample. On the contrary, a big standard deviation will generate a small angle and consequently a small *D*.

4.2 Shapes Similarity Concept

The Euclidean distance between 2 shapes seems to be the most obvious resource to measure the distances between shapes and use the resulting metric to determinate a normality/ abnormality score for each new shape. Euclidean distance is a simple method that can measure the distance between 2 individual shapes, however it results in inconsistencies for the purpose of this work. For example, in Fig 2, shape 1 and shape 2 represent two very different behaviours in a 31-dimensional window time. Each axis plots the entropy and the dispersion for each feature. The Euclidean distance between both shapes is 2.6833 (Table 1, 4th column). As can be seen in Fig 2, *both shapes have same area and describe the same form but they are placed in different positions.* When both are compared with a 3rd static reference 31-dimensional shape, the distance remains the same for both (Table 1, columns 1-3).
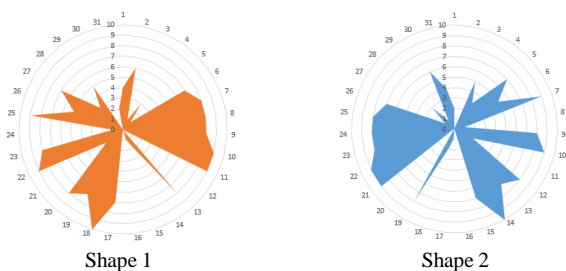
Figure 2: Shown are two different shapes each with 31-dimensions representing different behaviours.

Table 1: Euclidean distances study results

	Ref Point 1	Ref Point 2	Ref Point 3	Point 1
Point 1	2.03	2.54	1.83	0.00
Point 2	2.03	2.54	1.83	2.68

Table 2 Distance obtained by comparing the previously presented 2 shapes with the 3 reference shapes using the angle based projection procedure

	Ref Point 1	Ref Point 2	Ref Point 3
Point 1	9.34	9.69	7.85
Point 2	8.67	10.91	9.19

4.2.1 The angle based projection procedure. To measure the changes produced at the feature and feature dependencies level, we need a method to capture the following aspects of the shapes: *Area*, *Form* and *Position*. We propose a procedure that measures the distances of the angles generated by the projection of each feature to an origin point 0,0. Where the Y axis represents the position of the variable and the X axis represents the dispersion value. Once the *n* length of sequences of angles are generated, the Euclidian distance is calculated with the projected angles generated by the 3 static pre-defined behavioural shapes Fig 3 & Fig 4. Two reference shapes represent antagonistic behaviours with a correlation coefficient of -1; covering all of the behavioural spectrum. In practice, this implies that any behavioural shape scored far from reference point 1 has a good possibility to be similar to reference point 2. Behavioural shapes scored equally distant to reference point 1 & 2 have to describe a behaviour close to reference 3. This procedure can be considered as a part of the family of "projection based dimensionality reduction" procedures, with the peculiarity of using 3 references to measure the distances (Table 2).

In theory, this system could reduce any dimensionality space to 3. As any other projection based dimensionality reduction system, the bigger the dimensionality space, the poorer the accuracy of the resulting space. The projected angle is calculated using (3) below.

$$D(x) = \sin^{-1}\left(\frac{f}{\sqrt{f^2 + e^2}}\right) \qquad (3)$$

Where f represents the position of the feature and e the entropy/dispersion.

Figure 3: Reference point 1 and 2 describe an antagonistic behaviour having a Pearson correlation coefficient of -1.

Figure 4:

4.2.2 Detection of anomalies. In addition to the sliding window time previously defined, the system manages a second, and larger, sliding window time (Macro-Window Time) to measure the sustained level of similarity to the 3 predefined behaviours. We tested using two different length configurations for the second window time of 2 and 3 hours respectively. This means that the system will keep a hard copy of the behavioural shapes for the last 2 or 3 hours; kept in a FIFO (First in First out) manner. The purpose of this is to detect changes in the behaviours of variables that could indicate an anomalous behavioural change.

To do so, once we manage a reduced space dataset to 3 dimensions, the outlier detection becomes simpler and less computational expensive. In our experiments, we used the Chauvenet criterion and Euclidian distance to discover outliers in our behaviours. Chauvenet measures the probability of any point of being spurious given the average and standard deviation values of a data distribution. Chauvenet has main limitation that could origin bias, it assumes an underlying normal distribution of the data. Future work works exploring alternatives should be conducted. 2 main areas are proposed. The first is the application of clustering techniques (e.g. Streaming K means) to discover outliers. This technique presents the challenges of selecting the right number of clusters in an unbounded data set and, the right selection of the distance function (e.g. Eucledian, Mahalanobis). Nevertheless, any distance based procedure as clustering will face the fundamental challenge of deciding what are the acceptable boundaries for the clusters.

The second alternative is the application of autoregressive models (e.g. ARMA, ARIMA) to predict the next behaviour and detect the anomaly based on the discrepancy of the prediction with the reality. The main limitation of this approach is that fundamentally they are not unsupervised and required a training phase before they can predict results. A very promising area is proposed in 2 where the parameter required by ARIMA is automatically discovered from the data.

Regardless, the method used to qualify anomalies, our work demonstrates that by studding the distances to 3 antagonist static behaviours, changes in the traffic topology can be captured and categorized based on the shape similarity to previously identified behaviours.

5 CONCLUSIONS & FUTURE WORK

Network anomaly classification, in the context of IoT systems presents many challenges, and is difficult to achieve in practice. The IoT paradigm means a lack of historical information coupled with a diverse range of deployment scenarios, sensing and connectivity technologies. This leads to the need to employ an unsupervised approach to anomaly detection. Current unsupervised approaches assume that abnormal traffic is fundamentally different to normal traffic structures, this is not always the case in an IoT environment.

This paper demonstrates how treating anomalies as events that alter traffic feature distributions yields considerable diagnostic power in detecting and classifying these new anomalies. The effectiveness of using entropy and dispersion metrics for capturing unusual changes resulting from these events has also been shown. This paper contributes the ability to visualise anomaly structures using the procedures presented by measuring distances to previously defined classes and pre-defined reference classes in a normalised hyperplane.

Chauvenet has the limitation that it assumes an underlying normal distribution of the data. Future work will explore two main alternatives. The first is the application of clustering techniques (e.g. Streaming K means) to discover outliers. This technique presents the challenges of selecting the right number of clusters in an unbounded data set and, the right selection of the distance function (e.g. Eucledian, Mahalanobis). The second alternative is the application of autoregressive models (e.g. ARMA, ARIMA) to predict the next behaviour and detect the anomaly based on the discrepancy of the prediction with the reality.

The methods presented in this paper are not restricted to the monitoring of traffic feature distributions. Currently the authors are

investigating the application of this behavioural profiling procedure to the payload data of the producing IoT node. By focusing on sensed and reported data streams of an IoT node, we aim to classify supra-communities of things based on a model of data-stream/content or topic-of-interest, for building on-the-fly communities.

ACKNOWLEDGMENTS

REFERENCES

[1]   L. Atzori, A. Iera, and G. Morabito. 2010. The internet of things: A survey. *Computer networks*, vol. 54, no. 15, pp. 2787–2805.

[2]   S. Andreev and Y. Koucheryavy. 2012. Internet of things, smart spaces, and next generation networking. Springer, *LNCS*, vol. 7469, p. 464.

[3]   J. Q. Anderson and H. Rainie. 2014. The Internet of Things Will Thrive by 2025.

[4]   M. Abomhara. 2015. Cyber security and the internet of things: vulnerabilities, threats, intruders and attacks. Journal of Cyber Security and Mobility, vol. 4, (1), pp. 65-88.

[5]   J. S. Kumar and D. R. Patel. 2014. A survey on internet of things: Security and privacy issues. *International Journal of Computer Applications*, vol. 90, no. 11, pp. 20–26, published by Foundation of Computer Science, New York, USA.

[6]   A. Stango, N. R. Prasad, and D. M. Kyriazanos. 2009. A threat analysis methodology for security evaluation and enhancement planning. In *Emerging Security Information, Systems and Technologies*, pp. 262–267.

[7]   A. Perrig et al. 2002. SPINS: Security protocols for sensor networks. *Wireless Networks*, vol. 8, pp. 521-534.

[8]   T. Gaskill. 2016. When'Things' Attack. *Qual. Prog.*, vol. 49, pp. 10.

[9]   J. Gamblin. 2016. Leaked Mirai Source Code for Research/IoC Development Purposes. GitHub repository, https://github.com/jgamblin/Mirai-Source-Code.

[10]   S. Mansfield-Devine. 2016. DDoS goes mainstream: how headline-grabbing attacks could make this threat an organisation's biggest nightmare. *Network Security*, vol. pp. 7-13, 2016.

[11]   K. Hodgson. 2015. The Internet of [Security] Things. *SDM Magazine*, Available: http://www.sdmmag.com/articles/91564-the-internet-of-security-things.

[12]   D. Jiang and C. ShiWei. 2010. A study of information security for m2m of iot. *In Advanced Computer Theory and Engineering (ICACTE)*, 3rd International Conference on, vol. 3. IEEE, pp. V3–576.

[13]   B. Schneier. 2011. Secrets and lies: digital security in a networked world. John Wiley & Sons.

[14]   H. Tschofenig, J. Arkko, D. Thaler and D. McPherson. 2015. Architectural Considerations in Smart Object Networking. RFC 7452, DOI 10.17487/RFC7452.

[15]   D. Thaler, H. Tschofenig and M. Barnes. 2015. Architectural Considerations in Smart Object Networking.

[16]   Int Area Wiki - Internet-of-Things Directorate. IOTDirWiki. IETF, n.d. Web. 06 Sept. 2015. http://trac.tools.ietf.org/area/int/trac/wiki/IOTDirWik

[17]   O. Elloumi et al. 2015. IoT/M2M from research to standards: The next steps (part I) [guest editorial], IEEE Communications Magazine, vol. 53, pp. 8-9.

[18]   S. Meinrath and V. Pickard. 2008. Transcending net neutrality: Ten steps toward an open Internet, Education Week Commentary, vol. 12, pp. 1.

[19]   ISO JTC-1, Internet of things preliminary report, [online] Available: http://www.iso.org/iso/internetofthingsreport-jtcl.pdf.

[20]   Rose, K., S. Eldridge, and L. Chapin. 2015. The Internet of Things: An Overview– Understanding the Issues and Challenges of a More Connected World. The Internet Society (ISOC).

[21]   P. Kasinathan et al. 2013. Denial-of-service detection in 6LoWPAN based internet of things, *in Wireless and Mobile Computing, Networking and Communications (WiMob)*, IEEE 9th International Conference, pp. 600-607.

[22]   S. Raza, L. Wallgren and T. Voigt. 2013. SVELTE: Real-time intrusion detection in the Internet of Things, *Ad Hoc Networks,* vol. 11, pp. 2661-2674.

[23]   Y. Zhang, Z. Ge, A. Greenberg, and M. Roughan. 2005. Network anomography, *in Proceedings of the 5th ACM SIGCOMM conference on Internet Measurement*, IMC '05, (Berkeley, CA, USA), pp. 30–30, USENIX Association.

[24]   G. Cormode and S. M. Muthukrishnan. 2005. What's new: finding significant differences in network data streams, *Networking*, IEEE/ACM Transactions on, vol. 13, pp. 1219 – 1232, Dec.

[25]   A. Lakhina, M. Crovella, and C. Diot. 2004. Diagnosing network-wide traffic anomalies, in Proceedings of the 2004 conference on Applications, technologies, architectures, and protocols for computer communications, SIGCOMM '04, (New York,NY, USA), pp. 219–230, ACM.

[26]   A. Lakhina, M. Crovella, and C. Diot. 2004. Characterization of network-wide anomalies in traffic flows, *in Proceedings of the 4th ACM SIGCOMM conference on Internet measurement*, IMC '04, (New York, NY, USA), pp. 201–206, ACM.

[27]   A. Lakhina, M. Crovella, and C. Diot. 2005. Mining anomalies using traffic feature distributions, *in Proceedings of the ACM SIGCOMM 2005 conference*, SIGCOMM '05, (New York, NY, USA), pp. 217–228, ACM.

[28]   P. Bereziński, B. Jasiul and M. Szpyrka. 2015. An entropy-based network anomaly detection method, *Entropy*, vol. 17, pp. 2367-2408.

[29]   P. Casas, P. Fiadino and A. D'Alconzo. 2016. Machine-learning based approaches for anomaly detection and classification in cellular networks, *in Proceedings of the 8th International Workshop on Traffic Monitoring and Analysis*, pp. 1-8.