Novel anomaly detection and classification schemes for Machine-to-Machine uplink

Akshay Kumar, Ahmed Abdelhadi and Charles Clancy
Hume Center, Virginia Tech
Email:{akshay2, aabdelhadi, tcc}@vt.edu

***Abstract*—Machine-to-Machine (M2M) networks being connected to the internet at large, inherit all the cyber-vulnerabilities of the standard Information Technology (IT) systems. Since perfect cyber-security and robustness is an idealistic construct, it is worthwhile to design intrusion detection schemes to quickly detect and mitigate the harmful consequences of cyber-attacks. Volumetric anomaly detection have been popularized due to their low-complexity, but they cannot detect low-volume sophisticated attacks and also suffer from high false-alarm rate. To overcome these limitations, feature-based detection schemes have been studied for IT networks. However these schemes cannot be easily adapted to M2M systems due to the fundamental architectural and functional differences between the M2M and IT systems. In this paper, we propose novel feature-based detection schemes for a general M2M uplink to detect distributed Denial-of-Service (DDoS) attacks, emergency scenarios and terminal device failures. The detection for DDoS attack and emergency scenarios involves building up a database of legitimate M2M connections during a training phase and then flagging the new M2M connections as anomalies during the evaluation phase. To distinguish between DDoS attack and emergency scenarios that yield similar signatures for anomaly detection schemes, we propose a modified Canberra distance metric. It basically measures the similarity or differences in the characteristics of inter-arrival time epochs for any two anomalous streams. We detect device failures by inspecting for the decrease in active M2M connections over a reasonably large time interval. Lastly using Monte-Carlo simulations, we show that the proposed anomaly detection schemes have high detection performance and low-false alarm rate.**

***Index Terms*—M2M, Anomaly detection, DDoS attack, M2M emergency, Device failures**

I. Introduction

The global Machine-to-Machine (M2M) market is rapidly expanding with a plethora of booming applications such as smart grid, industrial control system and building automation [1]. M2M subsystems are increasingly using wireless protocols (such as Zigbee, WirelessHART, GSM etc.) and thus prone to (passive) eavesdropping attack and (active) jamming attacks [2]. Also they are vulnerable to cyber-attacks due to the use of *networked* TCP/IP architecture along with the adoption of Internet-of-Things technology to reduce its infrastructure and operational costs [3]. Most of the crucial applications operate on commercial off-the-shelf hardware and Windows/Unix operating systems. Thus, M2M systems inherit all the known cyber-threats and vulnerabilities for standard Information and Communications Technology (IT) systems.

Since designing a completely secure and robust M2M system would be expensive and impractical; there exists a non-zero probability of network intrusion by the malicious actors [4]. Therefore, it is important to design intrusion detection schemes (IDS) that can quickly detect a vast array of cyber threats to mitigate their negative consequences [5]. A large number of host-based and network-based IDS methods have been developed for IT systems [6]. Host-based systems detect attacks by processing large amount of internal accounting audit trail data from the users of the system (see [7] and references therein). Network-based systems collect network traffic from a router or switch (see [8] and references therein). These systems can be either signature-based if they scan network traffic for known attacks with specific bit signatures. Alternatively, they may use machine learning techniques to mine for anomalous traffic patterns resulting from non-standard process behavior; which is also the focus of this work but in the context of M2M systems.

The IT systems are confronted with a vast spectrum of traffic anomalies such as distributed denial of service (DDoS) attack, port-scans, equipment failures, flash crowds etc. This makes automatic anomaly detection and classification particularly challenging. Anomaly detection schemes proposed in literature are either volumetric schemes built upon measuring volume of network traffic or feature-based schemes that search for change in distributional aspects of certain traffic features. Some of the fundamental contribution in this field was made by Lakhina et. al. in [9] wherein they presented a general feature-based anomaly detection framework for IP networks using subspace and machine learning tools. However, these schemes cannot be easily extended to M2M systems due to fundamental differences in their architecture and operational requirements. The field devices such as sensors, are usually the most vulnerable component of a M2M system due to their limited on-board resources which preclude the implementation of current cyber-security methods.

Anomaly detection using machine learning tools such as neural networks and k-means has been studied for wireless sensor networks in [10], [11]. Kumar in [12] proposed a cluster-analysis based anomaly detection framework for a M2M network that uses SS7 signaling. Recently, Sedjelmaci et. al. [13] studied anomaly detection in a IoT network. They

This research is based upon work supported by the National Science Foundation under Grant No. 1134843.

used a game-theoretic approach to intelligently activate the anomaly detection techniques so as to strike a balance between energy consumption and the accuracy of anomaly detection. In this work, we propose a novel anomaly detection and classification framework for a general M2M uplink system. The M2M uplink consists of multiple M2M aggregators (MAs) that aggregate the local sensory traffic, which is eventually processed upstream at a M2M application server (AS). In a M2M uplink, the sensors are usually the most vulnerable to cyber-attacks or internal failures, which manifest themselves as anomalies in the local sensor traffic at each MA or in the global traffic at M2M AS. Therefore, we consider local (global) anomaly detection engine at MAs (AS), the detected anomalies are classified using an anomaly classification engine and eventually passed on to a post-processing operation that generates an appropriate response to the detected anomalies. The local/global anomaly detection engines employ volumetric and feature-based detection techniques for detecting different anomaly patterns resulting from distributed Denial-of-Service (DDoS) attacks, emergency scenarios and sensor failures.

To overcome the performance limitations of volumetric detection schemes, we propose novel feature-based detection schemes for detecting DDoS attacks, emergency scenarios and terminal device failures in M2M applications. The detection for DDoS attack and emergency scenarios involves building up a database of legitimate M2M connections during a training phase and then flagging the new M2M connections as anomalies during the evaluation phase. To distinguish between DDoS attack and emergency scenarios that yield similar signatures for anomaly detection schemes, we propose a modified Canberra distance metric. It basically measures the similarity or differences in the characteristics of inter-arrival time epochs for any two anomalous streams. We detect device failures by inspecting for the decrease in active M2M connections over a reasonably large time interval. Lastly using Monte-Carlo simulations, we show that the proposed anomaly detection schemes have high detection performance and low-false alarm rate.

The rest of the paper is organized as follows. Section II introduces the proposed M2M anomaly detection and classification framework. Then in Section III, we introduce a set of volumetric and feature-based anomaly detection metrics. We use these metrics to detect and classify different anomalies as described in Section IV. Using Monte-Carlo simulations we study the performance of proposed anomaly detection schemes in Section V. Finally Section VI draws some conclusions.

II. System Model

Fig. 1 shows the anomaly detection framework for a general M2M system. The uplink traffic from each local group of terminal devices is first aggregated at a MA and then the aggregated traffic from all MAs is processed at the M2M AS. We assume the aggregated traffic at each MA follows a Poisson arrival process. Similar assumptions have been made in earlier work [14]. While the aggregated data at MAs usually consists of legitimate packets from the sensors, occasionally it may include fraudulent or anomalous traffic that needs to be quickly detected to avoid disastrous consequences to the system. Anomalies in sensor traffic may be arise from events internal to the system such as sensor failures or emergency scenarios (such as fire in a building) which are usually benign with non-malicious intent. Or they may arise from malicious external actors involving dedicated efforts to cause significant damage to the system. A common example of such attack is the Distributed Denial of Service attack.

To detect these anomalies, we propose to implement a local anomaly detection engine at each MA. It stores the specific procedures for detecting each kind of anomaly. The anomaly detection at MAs are basically volumetric and/or feature based detection methods. Volumetric techniques depend on the number of packets arriving at MAs and in general have low implementation complexity. However, they are not good at detecting sophisticated attacks that are not accompanied with appreciable change in traffic volume. Additionally, they suffer from relatively high false alarm rate for high anomaly detection probabilities. On the contrary, feature-based technique utilize some measurable traffic characteristics such as entropy of source or destination address. They can detect a very wide spectrum of anomalies but suffer from high complexity.

It is important to point out here that there may be anomalies that cannot be detected by either volumetric or feature based detection methods at individual MA. This may occur due to minuscule changes in traffic volume at each MA or the target feature for detecting the anomaly requires analysis of the traffic across MAs. For detecting these sophisticated attacks, we propose to implement a global anomaly detection engine at the M2M AS. Again, the global anomaly detection engine can be volumetric and/or feature-based depending on the particular anomaly.

The anomalies detected at MAs or AS are forwarded to a anomaly classification engine that contains rules for classifying anomalies that generate similar signature at the anomaly detection engine. Lastly, the detected and classified anomalies along with their meta-data are passed to the post-processing step to take appropriate defensive action to protect the system.

III. Anomaly Detection Metrics

We now propose a set of volumetric and feature-based anomaly detection metrics used in this work. For this, we first divide the time horizon into time-slots, each of duration T sec. Then let k ($\forall\ k \in \mathbb{Z}^{+}$) denote the discrete time-index mapped to the time-slot, $[(k-1)T, kT]$ and let the set $\mathcal{K} = \{1, 2, \cdots, k\}$.

Volumetric metric: The volumetric metric considered in this paper is number of packets arrived at i^{th} MA or AS in k^{th} time-slot denoted by $N_i(k)$ and $N(k)$ respectively.

Feature-based metrics: We first define a *M2M connection* between a sensor and its corresponding MA if the MA receives at least one packet from that particular sensor. The first metric we consider, involves building up a database of unique M2M

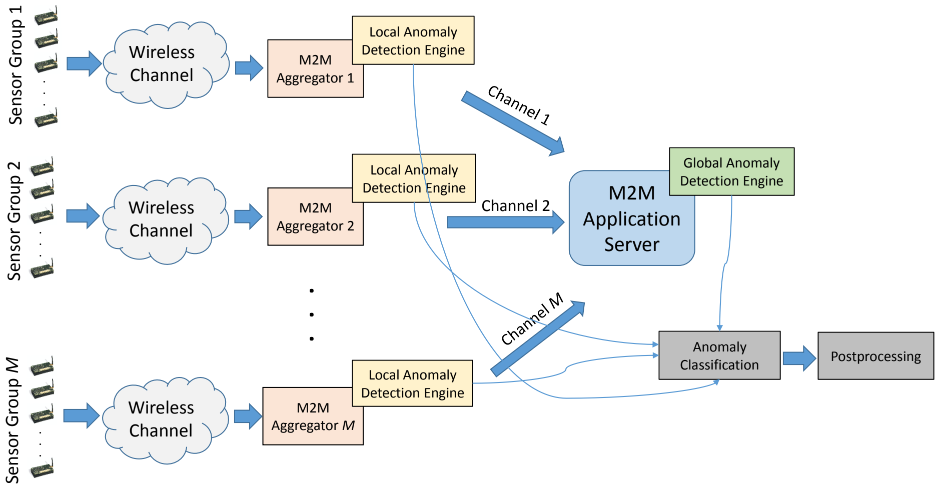

Fig. 1: Anomaly detection framework for M2M system.

connections made at each MA during a training period, $t \in [0, k_o T]$. Mathematically, we have,

$$H_i(k) = \cup_{m \in \mathcal{K}}\ h_i(m)\ \forall\ k \le k_o, \tag{1}$$

where $H_i(k)$ are the contents of the database at i^{th} MA at the end of k^{th} time-slot and $h_i(m)$ is the list of M2M connections made during m^{th} time-slot at i^{th} MA.

After the training period, the database at each MA is used to search for anomalies in each time-slot by determining the number of new M2M connections that are not present in the database. Mathematically, we have,

$$C_i(k) = h_i(k) \setminus H_i(k_o)\ \forall\ k > k_o, \tag{2}$$

where $C_i(k)$ denotes the number of new M2M connections at i^{th} MA during the k^{th} time-slot.

Another feature-based metric we consider in this work is the total number of M2M connections made at i^{th} MA over a sliding window of k_a time-slots. So at the k^{th} time-slot we have, $A_i(k) = |\cup_{m \in \mathcal{K}_a} h_i(m)|\ \forall\ k \ge k_a$. Here the set $\mathcal{K}_a = \{k - k_a + 1, k - k_a + 2, \cdots, k\}$ and the set operator $|.|$ denotes the cardinality of the set.

IV. Anomaly Detection and Classification

In this section, we detect the anomalies using the volumetric or/and feature-based metrics and attribute them to either a DDoS attack, an emergency scenario or device failures.

A. DDoS attack or Emergency scenario

The DDoS attack in a M2M application employs a botnet comprising of sensors that have been compromised using malicious software to send spam messages to the MAs. This leads to depletion of communication and computational resources at MAs and AS and thus prevents the legitimate sensors from communicating to MAs. DDoS attack are marked by increase in traffic volume. However, due to the distributed nature of attack, the increase in traffic at individual MAs may not be significant. Another reason for this is to evade detection at the MAs. Therefore, the global volumetric metric at AS, $N(k)$, may be more useful in this case. Alternatively, the DDoS attacks can be detected by searching for new M2M connections made during the attack, as measured by the metric $C_i(k)$ for i^{th} MA.

To safeguard a M2M system in the event of a disastrous situation (such as fire in a building), a set of sensors actively transmit alarm messages to the MA when they sense the calamity i.e, the measured physical quantity exceeds a predefined threshold. This again is marked by significant increase in the traffic volume at MAs that lie in the disaster zone and thus can be detected using the local volumetric metric $N_i(k)$ at the i^{th} MA. Alternatively, the emergency can be detected by searching for new M2M connections made by the emergency sensors, as measured by the metric $C_i(k)$ for i^{th} MA.

We note that both DDoS and emergency scenario result in similar signatures for the feature-based detection metric $C_i(k)$. However, it is important to distinguish between the two anomalies as they warrant different set of responses from the system operators. To differentiate between the two cases, we compute a modified Canberra distance metric (MCD) for the anomalous packet streams using the following procedure,

1) Select any two anomalous packet streams, P and Q, arriving at a given MA.
2) Select subset of arrival epochs P and Q by defining $p = P[r : r + n_c]$ and $q = Q[1 : n_c]$. Here r is the offset between P and Q and n_c is the number of arrivals used for computation.
3) Synchronize the start of streams p and q by setting, $p = p - (p[1] - q[1]), \text{if}\ p[1] > q[1]$ or $q = q - (q[1] - p[1]), \text{if}\ p[1] < q[1]$.
4) Compute the MCD as,

$$d(p, q) = \sum_{j=1}^{n} \frac{|p[j] - q[j]|}{|p[j]| + |q[j]|}. \tag{3}$$

Since the emergency detecting sensors have similar packet

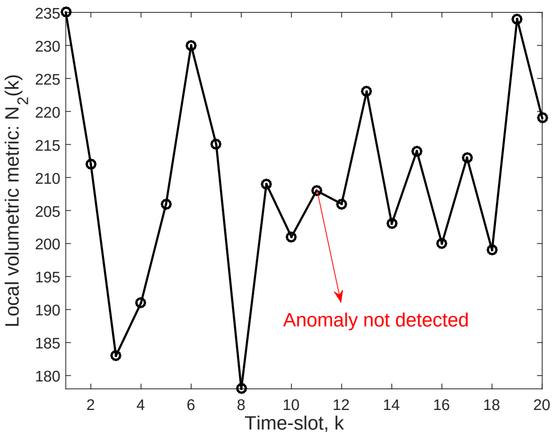

Fig. 2: Local volumetric detection for DDoS attack at MA 2.

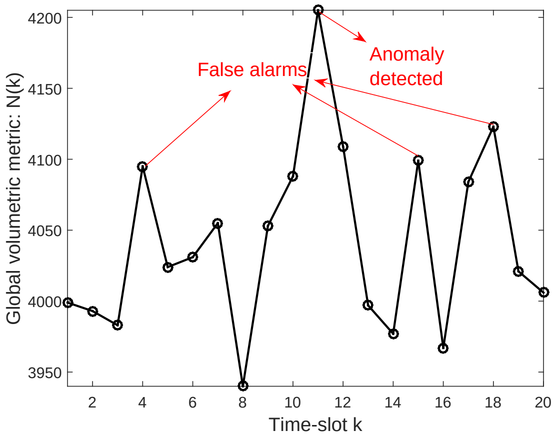

Fig. 3: Global volumetric detection for DDoS attack at M2M AS.

inter-arrival time epochs, we expect the MCD for packet streams from any two sensors to be 0. On the contrary, the packet arrivals from DDoS compromised sensors are usually independent of each other. This is done to obfuscate the anomaly detection engine that is searching for any patterns in traffic characteristics to detect anomalies.

B. Device failures

The terminal devices or sensors in a M2M application might malfunction due to hardware/software problems or be down due to battery exhaustion. It is important to detect and repair these devices as soon as possible, otherwise it may lead to a system failure. The local volumetric metric $N_i(k)$ proposed earlier, can be used to detect these anomalies by searching for dip in traffic volume at the i^{th} MA. Alternatively, we can use the feature-based detection metric $A_i(k)$ at time-slot k by searching for dip in total number of M2M connections at the i^{th} MA over the time duration of k_a previous time-slots. Now under normal system operation, the total number of connections in a sufficiently large window $\mathcal{K}_a$ is equal to the number of devices registered at the MA. Therefore, the reduction in connections can be used to infer the number of devices that are not functioning.

V. Simulation Results

In this section, we use Monte-Carlo simulations to evaluate the performance of proposed anomaly detection and classification framework. The time horizon for the simulation is set to 20 sec. The number of MAs is set to 20 and number of sensors per MA is 100. Since different sensors have potentially different packet arrival rate, we assume it to be uniformly distributed between 1 and 3 packets/sec. The number of sensors for detecting emergencies is set to 5 per MA. Lastly, the duration of each time-slot is $T = 1$ sec.

A. DDoS attack

We consider a botnet comprising of 32 compromised devices uniformly distributed across the M2M network. The attacker increases the packet arrival rate for each compromised sensor negligibly by 0.0032 packet/sec. For simplicity of exposition, we assume the attack lasts from 10 sec to 11 sec. Fig. 2 shows that the local volumetric detection at a MA fails to detect the attack due to relatively low attack volume. However, on analyzing the aggregated traffic volume at M2M AS, the attack may be detected as shown in Fig. 3. But it can also generates several false alarms, if the detection threshold (i.e. target packet volume) is low.

To overcome this limitation and get a clean anomaly detection, we resort to the feature-based detection approach as shown in Fig. 4. We first create a database for known M2M connections at any MA (say MA 2) until $k_o = 4$ time-slots and then use it to detect anomalies by determining any new M2M connections. If the training period is too small, we have a noise floor due to the legitimate M2M connections that were not identified in the training phase. On the other hand, we can set training period quite large to detect all legitimate M2M connections. But it leads to higher resource consumption at MAs and also has a opportunity cost in terms of time missing out on anomalies during that training phase.

B. Emergency scenario

Assume a disaster situation occurs in the coverage area of MA 2 from 13 sec to 14 sec. On detecting the emergency, the emergency sensors transmit alarm messages to MA 2 at a high rate of 25 packets/sec. The significant increase in traffic volume corresponding to the emergency can be easily detected as shown in Fig. 5. Furthermore, the feature-based detection method can also be used to detect the emergency, similar to that for DDoS attack. Fig 6 illustrates this by plotting the number of new M2M connections during the evaluation phase after the database is trained for $k_o = 3$ time-slots.

Since the feature-based metric yields the same signature for both DDoS and emergency scenario, we apply the Modified Canberra distance (described in Section IV-A) to distinguish between the two anomaly classes. Fig. 7 shows the plot

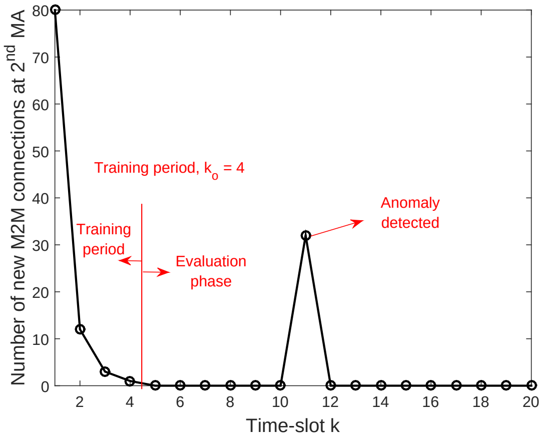

Fig. 4: Feature-based DDoS detection at MA 2.

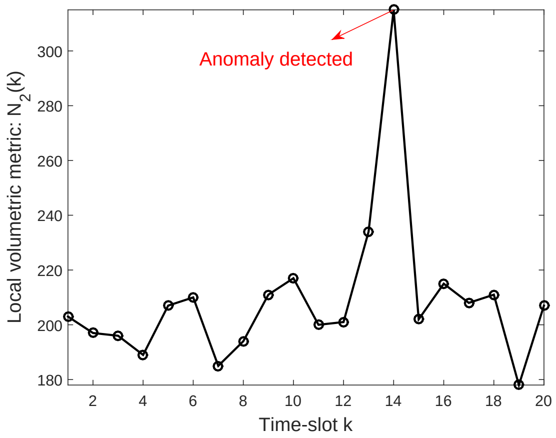

Fig. 5: Local volumetric detection for emergency situation at MA 2.

of Modified Canberra distance (MCD) applied to any two anomalous DDoS or emergency packet streams. As expected, we find that the MCD is 0 for emergency packet streams whereas it is relatively large for DDoS attack. This is true for all values of the packet offset r.

C. *Device failures*

We now consider the detection of device failures. We assume two devices in coverage area of MA 2 are malfunctioning and do not transmit to MA 2 from 10 sec to 16 sec. We first use local volumetric analysis to detect device failures at MA 2 as shown in Fig. 8. We note that it is not tough to detect device failures due to randomness in the number of packets arriving at MAs per time-slot. The problem is worsened when number of failures is small compared to total number of devices. This is true for our simulation as just 2 out of 100 devices are not working.

In order to improve the detection, especially at low failure rates, we resort to the proposed feature-based technique.

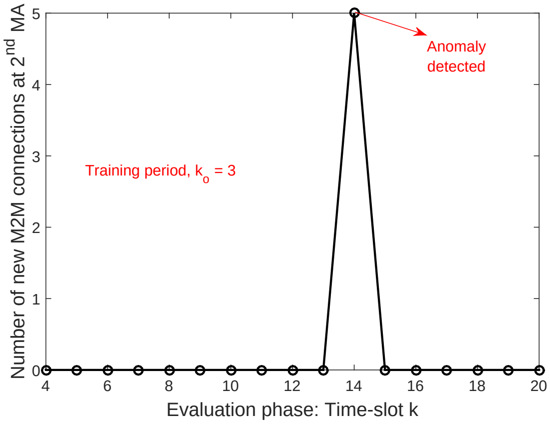

Fig. 6: Feature-based emergency scenario detection at MA 2.

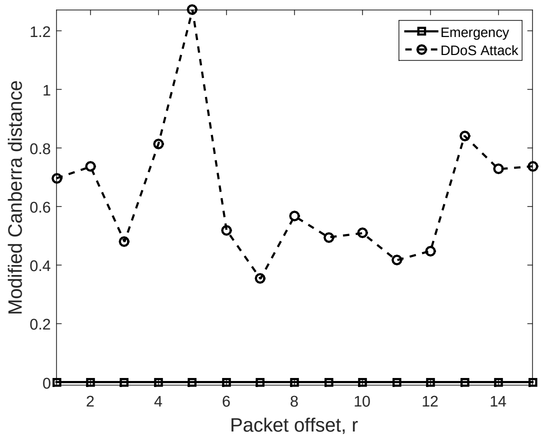

Fig. 7: Modified Canberra metric to distinguish between DDoS attack and emergency scenario. The metric is computed over a window of $n_c = 10$ packets.

Specifically, we compute the total number of M2M connections made over a sliding time-window of k_a time-slots. Fig. 9 plots the metric $A_2(k)$ as a function of time-slots with $k_a = 5$. The value of k_a is set large enough to ensure atleast one packet from each functioning device. We note a clear reduction in number of active devices or M2M connections from 11 sec to 16 sec.

Table I summarizes the performance of volumetric and feature-based anomaly detection schemes for various anomalies considered in this papers. For each anomaly type, the feature-based technique always performs better than its volumetric counterpart because it is designed to capture the unique variations in traffic characteristics induced by the anomaly of interest. The global volumetric techniques are not very relevant for detecting emergency scenarios and device failures. This is because these events manifest primarily as anomalies in the aggregated traffic at the MA to which the anomalous devices are connected. In fact, if global volumetric techniques are used, then it would be harder to detect these anomalies due to the

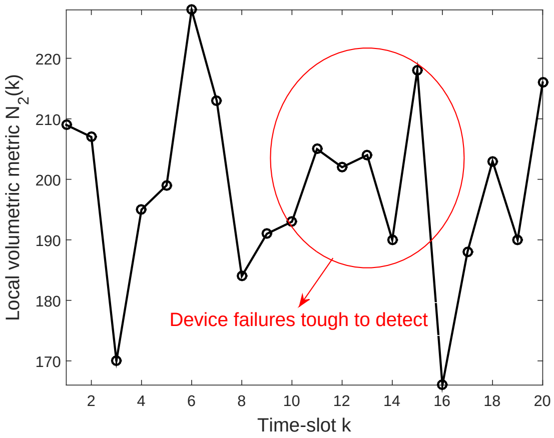

Fig. 8: Local volumetric detection for failure of devices connected to MA 2.

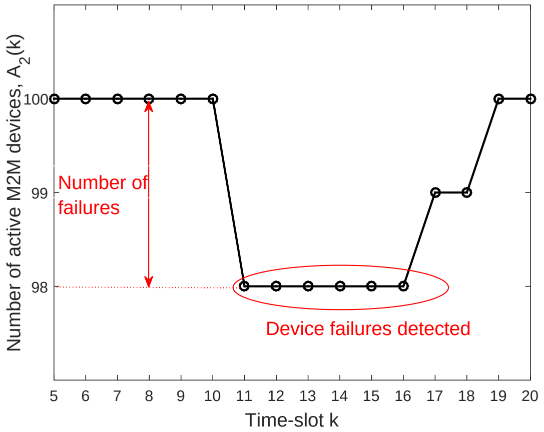

Fig. 9: Feature-based detection of device failures with $k_a = 5$.

increase in the non-anomalous traffic at the AS, arriving from other MAs.

VI. Conclusions

In this paper, we presented novel anomaly detection and classification schemes for M2M uplink. The proposed intrusion detection framework consists of local and global anomaly detection engine at M2M aggregators and M2M application server respectively. The detected anomalies are then forwarded to an anomaly classification engine and finally sent to the post-processing operation. To overcome the generally poor detection performance of volumetric anomaly detection techniques, we proposed novel feature-based detection schemes for detecting DDoS attack, emergency scenarios and device failures. The proposed feature-based schemes are built on predetermining the legitimate M2M connections and then flags the new M2M connections as anomalies during the evaluation phase. We also proposed a modified Canberra distance metric to distinguish between the DDoS attack and emergency scenarios which have similar signatures for proposed volumetric and feature-based detection schemes. It basically measures the similarity or differences in the characteristics of inter-arrival time epochs for any two anomalous streams. Device failures are detected by inspecting for the decrease in active M2M connections over a reasonably large time interval. Using extensive Monte-Carlo simulations for the M2M uplink system, we showed that the proposed anomaly detection techniques accurately detect different anomalies with very low false-alarms.

TABLE I: Comparative performance analysis of proposed anomaly detection schemes.

Anomaly Type	Local Volumetric	Global Volumetric	Feature-based
DDoS	Poor	Fair	Good
Emergency	Fair	Poor	Good
Device failures	Fair	Poor	Good

References

[1] G. Intelligence, "From concept to delivery: the M2M market today." https://goo.gl/yFmi5s, Feb. 2014.

[2] F. Ennesser, "Security in Machine-to-Machine Communication: The role of the Telecommunication Operator," tech. rep., Cinterion - Gemalto, 2012.

[3] J. Latvakoski, M. B. Alaya, H. Ganem, B. Jubeh, A. Iivari, J. Leguay, J. M. Bosch, and N. Granqvist, "Towards Horizontal Architecture for Autonomic M2M Service Networks," *Future Internet*, vol. 6, no. 2, p. 261, 2014.

[4] E. J. M. Colbert and S. Hutchinson, "Intrusion Detection in Industrial Control Systems," in *Cyber-security of SCADA and Other Industrial Control Systems* (E. J. M. Colbert and A. Kott, eds.), ch. 11, Springer, 2016.

[5] M. Bishop, S. Cheung, and C. Wee, "The threat from the net [Internet security]," *IEEE Spectrum*, vol. 34, pp. 56–63, Aug 1997.

[6] R. G. Bace, *Intrusion Detection*. Sams Publishing, 2000.

[7] P. T. Vit Bukac and M. Deutsch, "Advances and Challenges in Standalone Host-Based Intrusion Detection Systems," in *Trust, Privacy and Security in Digital Business* (S. Fischer-Hbner, S. Katsikas, and G. Quirchmayr, eds.), Springer, 2012.

[8] S. Sanyal, N. Das, and T. Sarkar, "SURVEY ON HOST AND NETWORK BASED INTRUSION DETECTION SYSTEM," *Acta Technica Corviniensis - Bulletin of Engineering*, vol. 8, pp. 17–20, Jan 2015.

[9] A. Lakhina, M. Crovella, and C. Diot, "Mining Anomalies Using Traffic Feature Distributions," *SIGCOMM Comput. Commun. Rev.*, vol. 35, pp. 217–228, Oct. 2005.

[10] S. Rajasegarar, C. Leckie, M. Palaniswami, and J. C. Bezdek, "Distributed Anomaly Detection in Wireless Sensor Networks," in *IEEE Singapore International Conference on Communication Systems*, pp. 1–5, Oct 2006.

[11] H. B. Wang, Z. Yuan, and C. D. Wang, "Intrusion Detection for Wireless Sensor Networks Based on Multi-agent and Refined Clustering," in *International Conference on Communications and Mobile Computing*, vol. 3, pp. 450–454, Jan 2009.

[12] P. D. Kumar, "A Cloud Based Automated Anomaly Detection Framework," Master's thesis, The University of Texas at Arlington, Texas, U.S.A,, 2014.

[13] H. Sedjelmaci, S. M. Senouci, and M. Al-Bahri, "A lightweight anomaly detection technique for low-resource IoT devices: A game-theoretic methodology," in *IEEE International Conference on Communications (ICC)*, pp. 1–6, May 2016.

[14] H. S. Dhillon, H. Huang, H. Viswanathan, and R. A. Valenzuela, "Fundamentals of Throughput Maximization With Random Arrivals for M2M Communications," *IEEE Transactions on Communications*, vol. 62, pp. 4094–4109, Nov 2014.