

Efficient Intrusion Detection in Ad-Hoc Networks

Mohammed Al Qurashi
Bournemouth University
Fern Barrow
Dorset, Poole, BH12 5BB
UK
malqurashi@bournemouth.ac.uk

Constantinos Marios Angelopoulos
Bournemouth University
Fern Barrow
Dorset, Poole, BH12 5BB
UK
mangelopoulos@bournemouth.ac.uk

Vasilios Katos
Bournemouth University
Fern Barrow
Dorset, Poole, BH12 5BB
UK
vkatos@bournemouth.ac.uk

We study efficient and lightweight Intrusion Detection Systems for Ad-Hoc networks via the prism of IPv6-enabled Wireless Sensor Networks. These networks consist of highly constrained devices organised in mesh networks following ad-hoc architectures, and as such carry specific characteristics that are not efficiently addressed by current state-of-the-art. In this work we first identify a trade-off between the communication and energy overhead of an IDS (as captured by the number of active IDS agents in the network) and the performance of the system in terms of successfully identifying attacks. In order to fine tune this trade-off, we first model such networks with the use of Random Geometric Graphs as this is a rigorous approach that allows us to capture underlying structural properties of the network. We then introduce a novel architectural approach for IDS by having only a subset of the nodes acting as IDS agents. These nodes are able to efficiently detect attacks at the networking layer in a collaborative manner by monitoring locally available network information provided by IoT routing protocols such as RPL. Our detailed experimental evaluation demonstrates significant performance gains in terms of communication overhead and energy dissipation while maintaining high detection rates.

IDS, IoT, WSN, sinkhole attack

1. INTRODUCTION

Internet of Things represents the next major networking paradigm shift both in qualitative and quantitative terms. Following previous paradigm shifts - such as Web 2.0, Cloud Computing and the Social Media - that enabled computers and people to be interconnected, now *things* and *machines* are able to seamlessly exchange information and data over the Internet.

The profound impact of IoT is highlighted by the number of IoT devices being deployed, as well as the growth of the corresponding market; a number of studies predict that by 2025 more than 55 billion IoT devices will be deployed and around 15USD trillion will be invested in IoT in aggregate between 2017 and 2025. Furthermore, IoT has already commenced its adoption in several verticals, such as healthcare, manufacturing and critical infrastructure Li et al. (2015). The importance and critical role of IoT in the modern economy has naturally made IoT systems targets of malicious activity, such as the infamous Mirai botnet Koliadis et al. (2017) or the case where hackers were able

to affect the steering and braking systems of a Jeep carRing (2015). It is worth noting that IoT networks and systems are not only the subjects but also the means for deploying the attacks (such as in the case of Mirai). Ironically, the deployment of such attacks has also been facilitated - among other factors - by the use of standardised Internet technologies that have enabled the paradigm in the first place. Broadly speaking, security controls can be taxonomised in three layers. At the first layer (also known as the first line of defence) lie preventative countermeasures, such as authentication and access control mechanisms, cryptographic primitives, firewalls, etc. At the second layer (also known as the second line of defence) lie detection countermeasures that are engaged *during* an attack, such as Intrusion Detection Systems. Finally, at the last layer lie recovery measures and processes for post-incident management, such as triaging, security information and event management and digital forensics. Due to the inherent and particular characteristics of IoT (i.e. highly constrained, deployed in mass numbers and their ephemeral availability), the corresponding cybersecurity measures need to be revisited.

A considerable amount of research has been carried out in Intrusion Detection Systems (IDS) concerning deployment architectures, detection strategies and algorithms. However, currently available IDS are designed for “traditional” computer networks, thus making strong assumptions about the system the IDS will be deployed in; e.g. that each node of the network is powerful in terms of resources, is always available and that the nodes communicate over a reliable and high-capacity network. As such, there is a need for innovative, lightweight IDS that will efficiently address the IoT paradigm.

Our contribution. A wireless sensor network (WSN) consists of a set of small and cheap autonomous devices deployed over an area of interest. The devices - commonly referred to as motes or simply sensors - are able to wirelessly communicate with their peers and, in spite of their highly constrained nature, to collaboratively carry out complex tasks. WSNs are a key enabling technology for the IoT and as such share several common characteristics. For the same reason, WSNs have provided an ideal R&D platform for several IoT protocols and technologies, such as the CoAP Shelby et al. (2014), 6LoWPAN Shelby and Bormann (2011) and 6TiSCH Dujovne et al. (2014). In this work we study efficient and lightweight Intrusion Detection Systems for the IoT via the prism of IPv6-enabled WSNs.

We first model a WSN with the use of graph-theoretical means - namely Random Geometric Graphs (RGG). The RGG model efficiently captures spatial characteristics of the network that are closely related to network connectivity; e.g. interdependencies on the existence of wireless links among neighbouring nodes. Then, motivated by how IoT networking protocols, such as RPL, manage and operate the network, we identify inherent trade-offs between the communication overhead introduced by an IDS and its detection rate of attacks such as sinkhole attack. We investigate this trade-off via extended emulations and show there exists an underlying threshold behaviour in the efficiency of the IDS that is related to the connectivity threshold of the RGG model. This allows us to conjure that for peer-to-peer IoT networks, the number of IDS agents that need to be deployed in order to achieve a high detection rate is constant and a function of the ratio between the network area size and the communication range of the nodes.

This paper is organised as follows: Section 2 presents the current state-of-the-art with a special emphasis on the most important contributions in Intrusion Detection Systems in WSN. Sections 3 and 4 introduces the proposed network model and

adopted IDS architecture based on Random Geometric Graphs. Section 5 presents the performance of the proposed system and discusses the simulation results and the findings. Finally, our paper concluded in section 6.

2. RELATED WORK

Intrusion Detection Systems for WSNs and the IoT have attracted significant research interest in the past years; in this section we focus on the most important contributions in the area. Coarsely speaking, IDSs can be classified based on their architecture into systems following a centralized, distributed or a hybrid architecture. In centralized IDSs, all relevant monitoring and detection information has to be reported to a centrally located base station where sophisticated detection algorithms are executed. Here, the base station is considered powerful in terms of processing capabilities and available memory and energy. On the other hand, in distributed architectures each individual network node is running an IDS agent and in cooperation with other agents in the network, they collaboratively detect any on-going attacks. Finally, hybrid IDS architectures demonstrate a combination of the centralized and distributed architectures in an effort to exploit and combine the advantages of each individual approach.

2.1. Centralized IDS Architectures

Moon et al. (2014) proposed a detection method for WSNs that aggregates the functions of the IDS with those of the intrusion prevention system. In this approach, symmetric encryption and oneway hash functions are used to establish the routing path between the base station (BS) and other nodes of the network. Their results showed that their approach reduces the total amount of required energy, in some cases quite significantly. However, this comes at a cost of increased computational overhead due to the use of a symmetric key.

Midi et al. (2017) proposed an intrusion detection system for IoT named Kalis. Kalis is placed at the border router to collect features of the network and use these to dynamically configure appropriate detection techniques. The authors claim that this approach can be applied and extended to new protocols as it is a protocol independent method being based on features.

2.2. Distributed Architecture IDSs

Kumarage et al. (2013) proposed a distributed anomaly detection framework for industrial WSNs that uses in-network hierarchical processing. The nodes of the network are first clustered using fuzzy

c-means clustering, and then run an incremental model to score local and global outliers. The proposed method was evaluated both on synthetic and real data and results showed a better trade-off to be achieved between achieving high-accuracy and introducing a smaller computational and communication overhead to the network.

Maleh et al. (2015) proposed a lightweight IDS for WSNs that combines anomaly-based detection (using support vector machine algorithms) with signature based rules. The authors study cluster-based topologies that reduce communication costs thus leading to extending the network lifetime. Simulation results show that their proposed model can detect abnormal events efficiently and has a high detection rate with a lower false-alarm rate.

Zheng et al. (2016) proposed an energy-efficient location-aware clone detection protocol in densely deployed WSNs, which can guarantee successful clone attack detection and maintain satisfactory network lifetime. Their approach is based on Specification Protocol Analysis that exploits location information of sensors to verify the legitimacy of other sensors in the proximity and detect clone attacks. Simulation results showed that this approach yielded high accuracy in detecting clone attacks and relatively lower energy consumption compared to other approaches.

Coppolino et al. (2013) proposed a distributed anomaly detection method for WSNs that uses both signature and anomaly-based detection techniques. Their framework is composed of a central agent and a number of local agents. The central agent employs data mining detection techniques whereas the local agents uses lighter detection techniques. Decision trees have been adopted as classification algorithm in the detection process of the central agent and their behaviour has been analysed in selected attacks scenarios. The empirical results revealed that this method exhibits low detection accuracy and high false positive rates.

2.3. Hybrid Architecture IDSs

Ponomarchuk and Seo (2010) introduced a detection method that uses misuse and anomaly detection techniques to detect sinkhole and deprivation attacks in WSNs. This method provisions local agents that are installed in each wireless sensor node as well as a central detection agent deployed at the base station. The local IDS agents are responsible for analysing traffic flow handled by the local node, gathering control data and sending it to the central agent. With such an arrangement local nodes are able to detect suspicious activities and to collaboratively contribute to the global detection

process. The authors evaluated their approach through simulations and noted that a high detection rate can be achieved. However, the proposed approach does not consider the highly-constraint nature of WSN and the limitations it poses on real-life systems. Raza et al. (2013) introduced an IDS for WSN that follows a hybrid architecture. Their solution focuses on routing attacks and consists of a central IDS module (running computationally intensive processes) that runs on the Sink node and a lightweight distributed agent that is deployed on sensor nodes. The proposed IDS has three main modules: a central module called mapper, a lightweight intrusion detection module and a firewall. The proposed solution shows a good performance in small networks, but it introduces a massive communication overhead in larger networks. This is mainly due to the fact that the lightweight agent is deployed on every single sensor node of the network, thus leading to bottleneck phenomena to emerge around the Sink as the diameter of the network increases.

The current state-of-the-art on IDSs for WSN and IoT networks are still resource-intensive and do not seem to adequately address the highly constrained nature of the corresponding devices. Centralized IDS architectures introduce significant communication overhead to the network as the base station (or Sink) injects and receives large numbers of requests to and from the nodes related to IDS data collection. Furthermore, in the special case of multi-hop peer-to-peer networks, bottleneck effect phenomena emerge in the areas close to the Sink as the corresponding nodes relay data from/to the rest of the network. Distributed IDS architectures largely rely on the cooperation among the sensor nodes, thus increasing the communication load as well as energy dissipation. Lastly, hybrid IDS architectures achieve a better control and global overview of the network, but currently available solutions also introduce a significant communication overhead that increases proportionally to the number of network nodes.

In this work we focus on hybrid IDS architectures but we show that by taking into account the specifics of IoT protocols, such as the ranking mechanism of RPL, as well as the spatial characteristics of such networks, the number of required IDS agents in the network (and therefore the corresponding overhead) can be greatly reduced while maintaining sufficiently high detection rates.

3. THE NETWORK MODEL

The paradigm of Internet of Things envisions the massive and seamless connection of embedded

systems, smart devices and things over the Internet. Wireless Sensor & Actuator Networks (WSANs) apart from being a key enabling technology for IoT, they also carry several characteristics that are typically found in several IoT systems. WSANs comprise of a big number of ultra-small sensor devices (which we also refer to as sensors), whose purpose is to monitor local environmental conditions (e.g. ambient luminance, temperature, etc.) and drive actuators (e.g. switches, valves, etc.). Each sensor is a fully-autonomous computing and communication device, characterized mainly by its constrained nature in terms of available power supply (battery), its transmission range r , the energy cost of data transmission and the (limited) processing and memory capabilities. In this work we focus our study on WSANs where the sensors are static and are deployed over the network area uniformly at random.

As mentioned earlier, there is a special node within the network the so-called sink S , that represents the gateway device located on the edge of the WSAN network. The Sink is assumed to be powerful in terms of computing power, memory and energy supplies. It is also the device that initiates the self-organisation of the network.

We consider that the random uniform placement of the sensors inside the network area is abstracted by a *Random Geometric Graph*. Random Geometric Graphs (RGG) are formed by n vertices that are placed uniformly at random in the $[0,1]^2$ square. An edge (u, v) exists iff the Euclidean distance of vertices u and v is at most r , where r corresponds to the wireless communication radius of the sensors. This holds assuming a disc radio model; two sensors can communicate with each other iff each one lies inside the communication range of the other. Random Geometric Graphs also have an important nice property: unlike other random graphs, like $G_{n,p}$, edges are not statistically independent of each other. That is, the existence of an edge (u, v) is not independent of the existence of edges (u, w) and (w, v) . This property makes RGG a quite realistic model for WSANs as it captures to a great extent the communication structure of real networks (at least their spatial aspects).

More strictly, consider an area $\mathcal{A} \subset \mathbb{R}^2$ in two dimensional space. An instance of the *random geometric graphs model* $\mathcal{G}(\mathcal{X}_n; r)$ is constructed as follows: select n points \mathcal{X}_n uniformly at random in \mathcal{A} . The set $V = \mathcal{X}_n$ is the set of vertices of the graph and we connect two vertices iff their euclidean distance is at most r . For any vertex $v \in V$ we will denote by $N(v)$ the set of neighbours of v and by $\deg(v) = |N(v)|$ its degree. Furthermore, we will

denote by $\|u - v\|$ the Euclidean distance between the points corresponding to vertices v, u

In Gupta and Kumar (1998) Penrose (2003) it is shown that the connectivity threshold for $\mathcal{G}(\mathcal{X}_n; r)$ is

$$r_c = \sqrt{\frac{\ln n}{\pi n}} \quad (1)$$

This way, the RGG model provides us with a formal tool of constructing and characterising networks as "sparse", "dense" or "normal". We also later find that this threshold also indicates the number of IDS agents needed in order to efficiently monitor a peer-to-peer, ad-hoc wireless network.

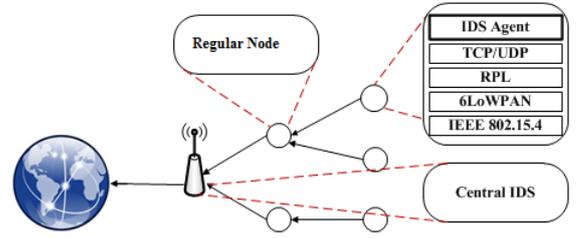


Figure 1: The Proposed Architecture

4. THE PROPOSED IDS

We propose an Intrusion Detection System architecture consisting of a central detection agent located in the base station and a distributed lightweight intrusion detection agent deployed on a subset of the network nodes as shown in Figure 1. The central agent manages the entire detection process and collects relevant data from the distributed agents. Each network node that runs an instance of the distributed agent, monitors and collects data on local network activity from its 1-hop neighbouring nodes. This implies that not all nodes need to run the IDS agent, but only a subset of them such that every node in the network has at least one 1-hop neighbour running the IDS agent. In graph-theoretical terms, such a subset would be a vertex cover of the corresponding RGG graph capturing the structure of the network. This also implies that there exists a minimum set of nodes that are able to efficiently monitor the network without compromising the performance of the IDS. This set corresponds to the minimum vertex cover for the corresponding RGG graph.

We apply the aforementioned approach on the state-of-the-art IDS for WSN by Raza et. al called SVELTE Raza et al. (2013). In their work, the authors consider multi-hop peer-to-peer IPv6-enabled WSNs running the 6LoWPAN stack Shelby and Bormann (2011) on ContikiOS Dunkels et al. (2004). They

Algorithm 1 Detect Sinkhole Attacks

Require: $M \leftarrow$ the list of IDS agents nodes
Require: $N \leftarrow$ the list of Regular nodes

```

for Node in M do
  for Node in N do
    if (Node.Rank+IDSagentNodeRank
    < Node.Parent.rank) then
      Node.fault=Node.fault+1
    end if
  end for
end for
for Node in N do
  if Node.fault>Threshold then
    Alarm
  end if
end for

```

develop an IDS following a hybrid architecture that consists of a centralized module running on the Sink and a distributed agent running on each individual sensor node. The centralized module contains the 6LoWPAN Mapper (6Mapper) which is responsible for gathering information from the sensor nodes on the network topology. In particular, 6Mapper collects information on the rank assigned to each node by the RPL protocol (responsible for constructing and maintaining a global tree-like network structure in a distributed manner) which is closely related to the hop distance of each node from the Sink. This allows a second component - the intrusion detection component - to reconstruct and monitor the network topology for anomalies that indicate an intrusion. For instance, a sinkhole attack could be deployed via a compromised node by having this node falsely announcing to its neighbours a significantly smaller rank. This would have its neighbouring nodes assume that its distance to the Sink is much smaller than the actual one, thus directing all network traffic to go through the compromised node.

As already mentioned, RPL establishes and maintains routing paths between the Sink and the rest of the network nodes by constructing a global tree-like network structure in a distributed way; the Destination Oriented Directed Acyclic Graph (DODAG). The process is initiated by the Sink broadcasting exploratory messages to its immediate neighbouring nodes, which in turn reiterate the process to their neighbouring nodes lying further away in the network. The process is run recursively and eventually results in each node being assigned a rank that depends on its actual hop-distance to the Sink as well as the link quality between neighbouring nodes (as measured by an objective function, such as the ETX metric). In SVELTE, the 6Mapper periodically collects these ranks to reconstruct the DODAG centrally at the Sink in order to monitor the network

against relevant attacks - like sinkhole - by detecting corresponding anomalies as shown in Algorithm 1; for example, if the rank of a node significantly deviates from the rank of its neighbours.

While for each individual node the introduced communication overhead may be small (the messages carrying the 6Mapper requests are 5 bytes long while each response from the nodes is 17 bytes long), engaging each individual node in the process introduces a communication overhead that is proportional to the size of the network. This poses significant scalability issues and adversely affects the connectivity and availability of the network as in multi-hop peer-to-peer networks nodes closer to the Sink also serve traffic coming from the rest of the network.

The key idea behind our approach is that networking protocols designed to address the distributed ad-hoc nature of peer-to-peer IoT networks (such as IPv6-enabled WSNs) make use of network information that is *locally available* to the nodes; as in the case of RPL. This network information can be easily shared with or even be monitored by 1-hop neighbouring nodes. Therefore, for a given set of neighbouring nodes it suffices that only one of them is actively collecting and reporting relevant information to the Sink. This greatly reduces the number of nodes that need to operate as IDS agents, thus mitigating any scalability and performance issues.

In this work we focus on experimentally investigating and evaluating our approach by extending SVELTE (the current state of the art IDS for IPv6-enabled WSNs). In particular, we focus on evaluating the trade-off between the potentially reduced accuracy of the IDS in successfully detecting attacks (due to the smaller number of active IDS agents in the network) versus the reduced communication overhead and increased energy efficiency of the network.

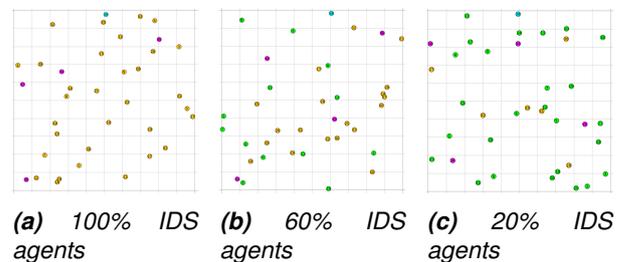


Figure 2: Sparse network

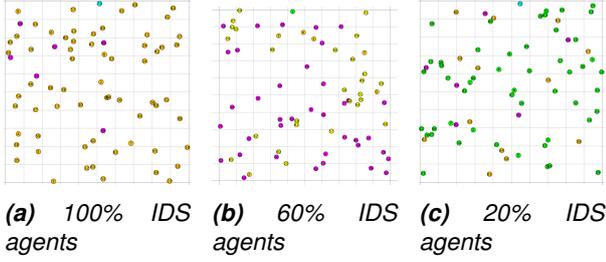


Figure 3: Normal density network

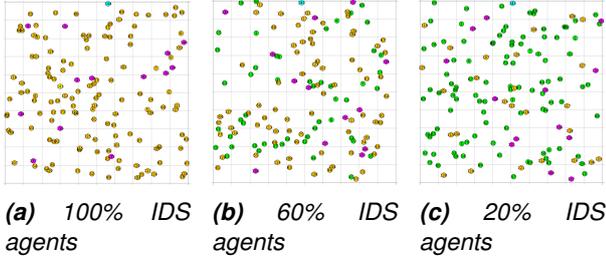


Figure 4: Dense network

5. PERFORMANCE EVALUATION

5.1. Simulation Set-Up

We ran our experiments using the Cooja Osterlind et al. (2006) emulator, which provides a detailed cross-layer simulation for WSNs running the 6LoW-PAN stack. We consider three qualitatively distinct network densities as these are indicated by the Random Geometric Graph model. In particular, we consider a network area $\mathcal{A} = [0, 100]^2$ where n sensor nodes are deployed uniformly at random, for $n \in \{32, 64, 128\}$. Following from equation 1, for each value of n , the corresponding network connectivity threshold is $r_c : \{18.5; 14.3; 11\}$ respectively. Therefore, by setting the sensors' communication range to be $r = 20$, we get three network set ups where r is (a) almost equal to; (b) $\times 1.5$ and (c) $\times 2$ the connectivity threshold, thus resulting in (a) *sparse*, (b) *normal* and (c) *dense* networks. Figures 2, 3 and 4 provide a visual representation of the various network densities.

For each network density, we consider five scenarios where the percentage of nodes acting as IDS agents is 100%, 80%, 60%, 40% and 20% of the total population (yellow nodes in the corresponding figures). Furthermore, in each case we set 10% of the node population to act as malicious nodes (nodes in purple) deploying sinkhole attacks by exploiting the rank mechanism of RPL. Any remaining nodes are regular nodes (nodes in green). For each network configuration we also run a scenario with no nodes operating as IDS nodes. For each scenario we create 10 random instances of the network; this allows us to effectively mitigate

in our simulations any issues that might occur due to the random network topology (in other words we sample the space of RGG instances). For each instance we run 10 iterations of simulating the network operation for a simulation time of 3600 seconds where nodes generate and transmit data approximately every second. For each scenario and each performance metric we compute average values and 95% confidence intervals.

Our findings demonstrate strong concentration around the mean and are therefore deemed statistically significant.

5.2. Evaluation Metrics

In the following subsections we define the metrics that are used to evaluate our proposed IDS architecture.

5.2.1. Detection Rate

We define the detection rate as the number of true positive detections of malicious nodes over the total number of malicious nodes in the network.

$$\text{Detection rate} = \frac{\text{number of true positive detections}}{\text{total number of malicious nodes}} \quad (2)$$

5.2.2. Communication Overhead

We define the communication overhead as the additional volume of data communication introduced in the network as a result of the operation of the IDS. We follow the practice of Raza et al. (2013) and monitor this metric only to the 1-hop neighbouring nodes of the Sink (the rationale is that any network traffic will have to go through these nodes prior to reaching the Sink). We denote by E_{IDS} the energy consumption of the said nodes with the IDS running and with $E_{\overline{\text{IDS}}}$ the energy consumption of the said nodes with no IDS running in the network. Then,

$$\text{Communication overhead} = \frac{E_{\text{IDS}} - E_{\overline{\text{IDS}}}}{E_{\overline{\text{IDS}}}} \quad (3)$$

5.2.3. Total Energy Consumption in the Network

We measure the total energy consumption ΔE_{total} in the network as the difference between the total available energy in the network at the beginning of a simulation and at the end. We denote initial available energy for sensor i by E_{init}^i and the initial available energy for sensor i by E_{final}^i . Then,

$$\Delta E_{\text{total}} = \sum_{i \in n} (E_{\text{init}}^i - E_{\text{final}}^i) \quad (4)$$

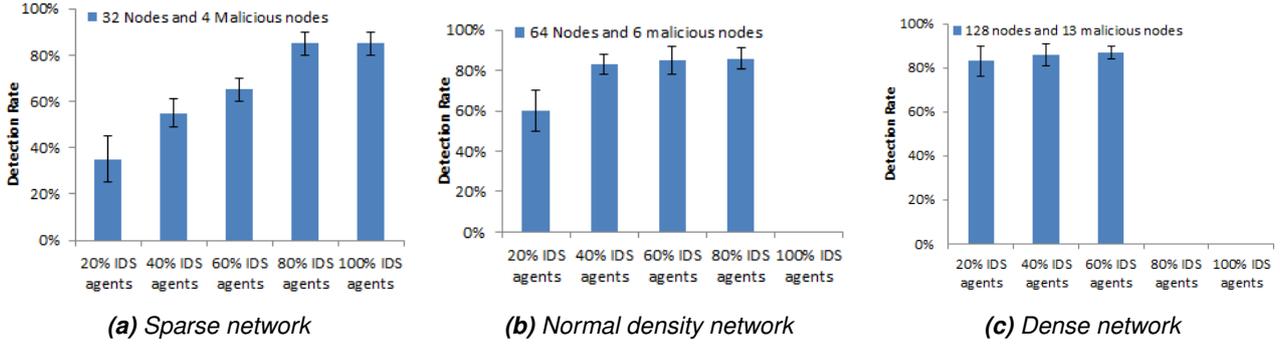


Figure 5: IDS detection rate against the number of active IDS agents in the network as percentage of the node population. Notice that in normal density and dense networks the overhead induced by the IDS is that high that the network gets disconnected prematurely.

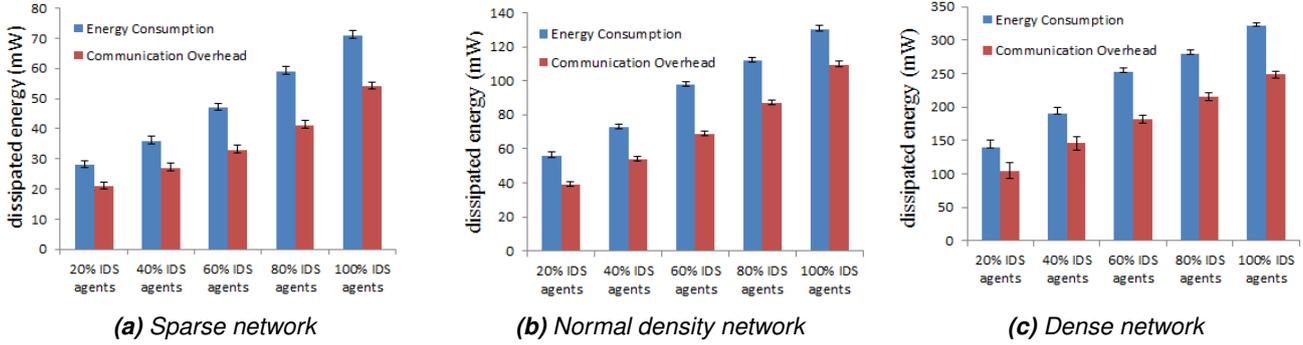


Figure 6: Energy consumption and communication overhead for the entire network

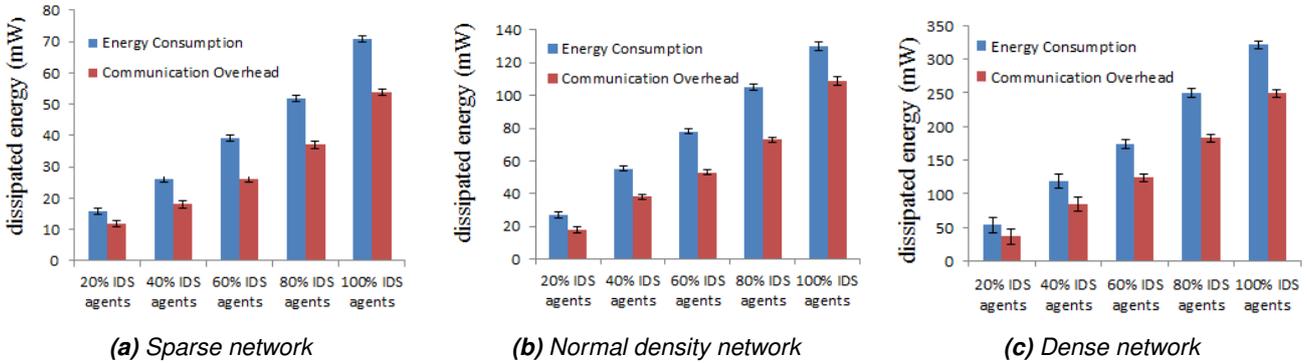


Figure 7: Energy consumption and communication overhead introduced by the IDS

5.3. Simulation Findings

Figure 5a shows that in sparse networks the detection rate remains as high as 85% for the scenarios where 100% and 85% of the node population operates as an IDS agent. However, the detection rate drops at 60% for 60% of the population as IDS agents, and continues to drop further as the percentage of the agents is reduced. This demonstrates that the IDS performance in sparse networks quickly drops due to the fact that areas of the network remain un-monitored. We note, however,

that there is a certain level of resilience for high percentages of IDS agents.

Figure 5b shows the findings for networks of normal density. We note two points. Firstly, the IDS demonstrates a greater degree of resilience as it achieves high detection rates even for percentages of IDS agents as low as 40% of the node population. Second, we note that for 100% of nodes as IDS agents the simulation was not completed due to the fact that the network was rendered disconnected as the nodes lying close to the Sink were not able to handle the increased network traffic. This

highlights the network strain that even light-weight IDSs introduce.

This is also the reason why other works in the literature on IoT and WSN IDS limit their simulation studies in networks with small populations. Figure 5c further highlights these findings as the simulations failed to complete for scenarios considering big numbers of IDS agents (percentages of 100% and 80%). Also, in dense networks the detection rate of the IDS remained at very high levels (circa 80%-85%).

At this point we make another important observation. For all three network densities, the detection rate of the IDS starts to deteriorate significantly (and as shown in 5a, proportionally to the reduction in IDS agents percentage) once the absolute number of IDS agents in the network drops below a constant threshold; in this case below 25 IDS nodes (corresponding to $32 \cdot 80\%$; $64 \cdot 40\%$; $128 \cdot 20\%$). This implies that only *a constant number of IDS agents* is needed to effectively and efficiently monitor the network.

This is a very strong indication that the efficiency of hybrid/distributed IDS for peer-to-peer ad-hoc networks is independent of the number of nodes but related to *underlying fundamental properties of the network*. Following our network modelling with the use of Random Geometric Graphs, we conjecture that this property is the size of the minimum vertex cover of the corresponding RGG instance. We intent to investigate this in our future work employing more formal and rigorous methods from graph theory.

Figures 6 and 7 show that the energy consumption and the communication overhead introduced to the network by the IDS is proportional to the number of nodes operating as IDS agents. This shows the massive gains that can be achieved with respect to fine-tuning the trade-off between energy efficiency and the achieved high detection rate as a result of using a constant number of IDS agents.

6. CONCLUSIONS AND FUTURE WORK

In this work we study efficient and lightweight Intrusion Detection Systems for ad-hoc networks via the prism of IPv6-enabled Wireless Sensor Networks. We first provide a formal model for WSNs with the use of Random Geometric Graphs, a graph-theoretical model to capture the spatial characteristics of WSNs such as interdependencies on the existence of wireless links among neighbouring nodes. Then, motivated from the operation of IoT-specific networking protocols such as RPL, we focus on

network attacks such as the sinkhole or man-in-the-middle. We identify and try to optimise the trade-off between energy efficiency and communication overhead on one hand, and the IDS detection rate on the other. By leveraging upon the distributed nature of such protocols and locally available network information, we propose a method for IDS that requires only a subset of the nodes to operate as IDS agents.

We extend the state of the art on IDS for WSNs by integrating our method and conduct our performance evaluation via extensive emulations. We consider various network densities (as these are formally defined via the RGG model) and show that 1) indeed the IDS detection rates remain at very high levels (around 85%) even with a subset of the nodes as IDS agents; 2) that the required number of IDS agents in the network in order to achieve these levels is independent from the network population and in fact *constant*; 3) that the energy consumption and communication overhead introduced by the IDS is proportional to the number of IDS agents, therefore our method allows for massive energy gains while not affecting the detection rate of the IDS.

In our future work we will employ rigorous graph-theoretical methods and tools to formally prove the conjectures made in this paper. We will also work in providing efficient algorithms for choosing in a distributed way which nodes should operate as IDS agents as well as balancing this role among all the nodes.

ACKNOWLEDGMENT

This work has received funding from the EU's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 778229 (Ideal-Cities).

REFERENCES

- Coppolino, L., D'Antonio, S., Garofalo, A. and Romano, L. (2013), Applying data mining techniques to intrusion detection in wireless sensor networks, *in* 'P2P, Parallel, Grid, Cloud and Internet Computing (3PGCIC), 2013 Eighth International Conference on', IEEE, pp. 247–254.
- Dujovne, D., Watteyne, T., Vilajosana, X. and Thubert, P. (2014), '6tisch: deterministic ip-enabled industrial internet (of things)', *IEEE Communications Magazine* **52**(12), 36–41.
- Dunkels, A., Gronvall, B. and Voigt, T. (2004), Contiki-a lightweight and flexible operating system for tiny networked sensors, *in* 'Local Computer Networks, 2004. 29th Annual IEEE International Conference on', IEEE, pp. 455–462.

- Gupta, P. and Kumar, P. (1998), *Critical power for asymptotic connectivity in wireless networks*, Stochastic Analysis, Control, Optimization and Applications, Boston.
- Kolias, C., Kambourakis, G., Stavrou, A. and Voas, J. (2017), 'Ddos in the iot: Mirai and other botnets', *Computer* **50**(7), 80–84.
- Kumarage, H., Khalil, I., Tari, Z. and Zomaya, A. (2013), 'Distributed anomaly detection for industrial wireless sensor networks based on fuzzy data modelling', *Journal of Parallel and Distributed Computing* **73**(6), 790–806.
- Li, S., Da Xu, L. and Zhao, S. (2015), 'The internet of things: a survey', *Information Systems Frontiers* **17**(2), 243–259.
- Maleh, Y., Ezzati, A., Qasmaoui, Y. and Mbida, M. (2015), 'A global hybrid intrusion detection system for wireless sensor networks', *Procedia Computer Science* **52**, 1047–1052.
- Midi, D., Rullo, A., Mudgerikar, A. and Bertino, E. (2017), Kalisa system for knowledge-driven adaptable intrusion detection for the internet of things, in 'Distributed Computing Systems (ICDCS), 2017 IEEE 37th International Conference on', IEEE, pp. 656–666.
- Moon, S. Y., Kim, J. W. and Cho, T. H. (2014), An energy-efficient routing method with intrusion detection and prevention for wireless sensor networks, in 'Advanced Communication Technology (ICACT), 2014 16th International Conference on', IEEE, pp. 467–470.
- Osterlind, F., Dunkels, A., Eriksson, J., Finne, N. and Voigt, T. (2006), Cross-level sensor network simulation with cooja, in 'Local computer networks, proceedings 2006 31st IEEE conference on', IEEE, pp. 641–648.
- Penrose, M. (2003), *Random Geometric Graphs*, Oxford University Press.
- Ponomarchuk, Y. and Seo, D.-W. (2010), Intrusion detection based on traffic analysis in wireless sensor networks, in 'Wireless and Optical Communications Conference (WOCC), 2010 19th Annual', IEEE, pp. 1–7.
- Raza, S., Wallgren, L. and Voigt, T. (2013), 'SVELTE: Real-time intrusion detection in the Internet of Things', *Ad hoc networks* **11**(8), 2661–2674.
- Ring, T. (2015), 'Connected cars—the next target for hackers', *Network Security* **2015**(11), 11–16.
- Shelby, Z. and Bormann, C. (2011), *6LoWPAN: The wireless embedded Internet*, Vol. 43, John Wiley & Sons.
- Shelby, Z., Hartke, K. and Bormann, C. (2014), The constrained application protocol (coap), Technical report.
- Zheng, Z., Liu, A., Cai, L. X., Chen, Z. and Shen, X. S. (2016), 'Energy and memory efficient clone detection in wireless sensor networks', *IEEE Transactions on Mobile Computing* **15**(5), 1130–1143.