Evolving Attackers against Wireless Sensor Networks

Kinga Mrugala
Department of Computer Science
UCL, London
kinga.mrugala.11@ucl.ac.uk

Nilufer Tuptuk
Department of Computer Science
UCL, London
n.tuptuk@cs.ucl.ac.uk

Stephen Hailes
Department of Computer Science
UCL, London
s.hailes@cs.ucl.ac.uk

ABSTRACT

Recent technological improvements in wireless communication and electronics have enabled the development of small, low-cost wireless sensor nodes, capable of monitoring everything from human health to the performance of the electricity grid. A natural consequence is a desire to secure systems containing these nodes. Unfortunately, proving that systems are secure is beyond the current state of the art, and testing for security is problematic: test cases often miss attacks that have never previously been seen. In this paper, we use Genetic Programming (GP) to create attacks against Internet of Things devices, to help identify vulnerabilities before systems are attacked for real. To assess the effectiveness of each attacker, we used it against a wireless sensor network (WSN) with publish-subscribe communications, protected by a literature artificial immune intrusion detection system (IDS). The GP attackers succeeded in suppressing significantly more legitimate messages than a hand-coded attacker, whilst decreasing the likelihood of detection. As a consequence, it was possible to tune the IDS, improving its performance. Whilst these results are preliminary, they demonstrate GP holds significant potential for improving the protection of systems with large attack spaces.

CCS Concepts

•Security and privacy → Artificial immune systems; Mobile and wireless security; •Computing methodologies → Genetic programming; •Networks → *Sensor networks;*

Keywords

Evolutionary Computation, Wireless Network Attacks, Routing Protocols

1. INTRODUCTION

The idea of an Internet of Things (IoT) in which objects of all kinds are networked together, has become so pervasive that it is at the heart of much current academic interest [6] as well as the commercial strategies of large companies [1]. Naturally, this gives rise to security concerns, which become even more important when one realises that IoT technologies are regarded as an exciting near-term prospect for use in industrial control systems - so called Industry 4.0. We present initial results that suggest that evolutionary computing techniques have the potential to revolutionise the way in which we engineer IoT systems by facilitating the automated creation of attackers against which one can test (and so adapt) one's defensive strategies.

Previous studies [4] have used GP to evolve attacks against wired networks, but as far as we know, there are no studies that have investigated using GP to evolve attacks against wireless sensor networks. In this paper, we explore a limited problem to demonstrate the feasibility of evolving attackers. We study a WSN protected with an artificial immune system (AIS) devised by Wallenta et al. [5]. We evolve attackers against it, and demonstrate their effectiveness in comparison to the strategies that are used in evaluating the effectiveness of the original AIS.

2. PRELIMINARIES

In this paper, we use the Sensor Network Based Artificial Immune System (SNAIS) [5], an intrusion detection system, based on a Dendritic Cell (DC) Algorithm [2]. SNAIS is designed to operate in the context of a WSN that uses directed diffusion [3]. Directed diffusion is a routing algorithm used to gather data from a large number of nodes and routing it to nodes that request the data. SNAIS operates in a similar manner to the human immune system, where signals are gathered and then used to decide on future actions.

For this study, we consider attacks based on the interest cache poisoning attack described in [5]. In an interest cache poisoning attack, an attacker attempts to inject bogus interest entries into the interest cache of a node. Under the directed diffusion protocol, the node chooses to where to send a data message using the interest cache entries. If legitimate entries can be forcefully eliminated from the cache by an attacker, nodes will drop legitimate data packets.

3. EXPERIMENTS

Genetic programming was used as the evolutionary approach for attack generation. To explore the problem, two fitness functions were used: the first focuses on suppressing the number of packets delivered from source to sink; the second focuses on trying to lower the precision and recall of SNAIS. Moreover, the attackers were also evolved on two

GECCO'16 Companion July 20-24, 2016, Denver, CO, USA

ACM ISBN 978-1-4503-4323-7/16/07.

DOI: http://dx.doi.org/10.1145/2908961.2908974

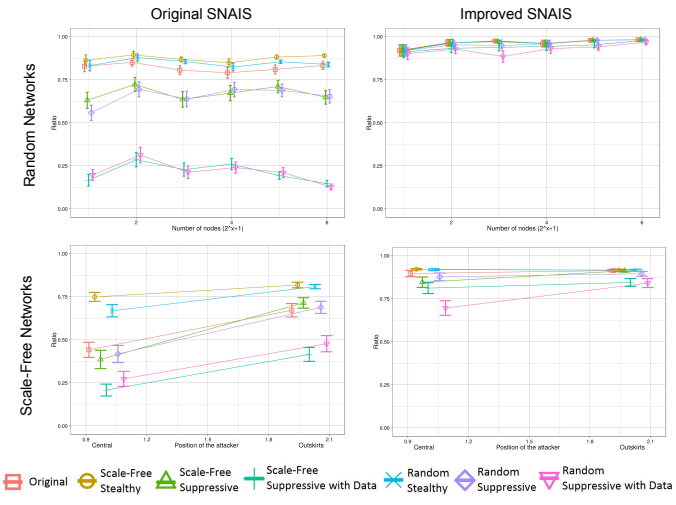

Figure 1: Ratio of Received Packets using original SNAIS and improved SNAIS for Random and Scale-Free Networks

different type of networks (random networks and scale-free networks) to understand the impact of network connectivity. Eight attackers were evolved, using different configurations of fitness functions, function sets and networks. In suppressive attacks, the attacker is attempting to suppress legitimate messages; in stealthy attacks, they are attempting to lower the precision and recall of SNAIS. In attacks with data, they have access to an additional function, not present in the original interest cache poisoning attack, which allows the attacker to also send data packets, not just broadcast interest packets.

4. RESULTS

The results obtained for random networks and scale-free networks are illustrated in Figure 1, illustrating the ratio of received packets using original SNAIS implementation, and the improved SNAIS.

On random networks, the *suppressive* attacker out performs the original paper's attacker. The number of packets that the *suppressive* attacker stopped is on average 17% higher. The *stealthy* attacker performed better than the original in terms of precision and recall. The attacker reduced the precision of the system on average by 17% and the recall by 22%. However, the *stealthy* attacker slightly decreased the suppression rate. The results on scale-free networks were worse, with the *suppressive* attacker not having a real impact on any of the metrics and the *stealthy* attacker producing only a marginal improvement.

Adding a new function to the GP function set did not have any effect on the changes to detection rates. However, for random networks the results show that adding the additional function significantly improved the suppression rate of the packets. The *suppressive-with-data* attacker suppressed on average 44% more packets than the *suppressive* attacker and 60% more packets than the original attacker. At the same time, this change did not increase the detectability of the attacker. Again, it performed slightly worse on scale-free networks; however it produced a 29% improvement in suppressing the flow of packets.

To show if it is possible to use evolved attackers to improve the overall results of the IDS, a simple trustworthiness metric was introduced. The trustworthiness of a source node is used to determine which packets to drop - the less trustworthy the source, the more likely that its packets are dropped. After this change to the AIS, the trustworthiness metric did improve the system's performance. For random networks, all of the attackers failed to suppress a significant number of the packets. The results were not as good for the scale-free networks as for random networks (though still a significant improvement). The number of packets delivered to the sink is around 87%, in comparison to 95% for random networks. For the *stealthy* attacker the number of packets delivered rose by 12% on random networks and by 14% on scale-free networks. Interestingly, in no case did the improved SNAIS result in a significant improvement in precision or recall.

5. CONCLUSION

The results appear to be very promising. The GP attacker performed considerably better than the original handcrafted attacker in both the suppression of legitimate messages and its detectability. As a result of analysing the weaknesses of the original SNAIS that the evolved attackers were seen to be exploiting, a simple trustworthiness metric was introduced, with profound effects, improving packet flow from source to sink by 95% on random networks. The structure of the network does appear to matter: on scale-free networks no attackers performed as well as they did on on random networks, whether or not they were specifically evolved using scale-free networks. Although preliminary, the results obtained are sufficiently encouraging to indicate a need for further work in this area, studying the role of evolutionary approaches in improving the overall security of complex and heterogeneous networks.

6. REFERENCES

[1] D. Evans. The internet of things: How the next evolution of the internet is changing everything. *CISCO Internet Business Solutions Group (IBSG)*, Apr 2011.

[2] J. Greensmith, U. Aickelin, and S. Cayzer. *Artificial Immune Systems: 4th International Conference, ICARIS 2005*, chapter Introducing Dendritic Cells as a Novel Immune-Inspired Algorithm for Anomaly Detection, pages 153–167. 2005.

[3] C. Intanagonwiwat, R. Govindan, D. Estrin, J. Heidemann, and F. Silva. Directed diffusion for wireless sensor networking. *IEEE/ACM Trans. Netw.*, 11(1):2–16, Feb. 2003.

[4] H. G. Kayacık, A. N. Zincir-Heywood, and M. I. Heywood. Evolutionary computation as an artificial attacker: generating evasion attacks for detector vulnerability testing. *Evolutionary Intelligence*, 4(4):243–266, 2011.

[5] C. Wallenta, J. Kim, P. J. Bentley, and S. Hailes. Detecting interest cache poisoning in sensor networks using an artificial immune algorithm. *Applied Intelligence*, 32(1):1–26, 2008.

[6] L. D. Xu, W. He, and S. Li. Internet of things in industries: A survey. *Industrial Informatics, IEEE Transactions on*, 10(4):2233–2243, Nov 2014.