Event Handoff Unobservability in WSN

Stefano Ortolani¹, Mauro Conti¹, Bruno Crispo², and Roberto Di Pietro³

¹ Vrije Universiteit, Amsterdam, The Netherlands

`{ortolani,mconti}@few.vu.nl`

² University of Trento, Trento, Italy

`crispo@disi.unitn.it`

³ Università di Roma Tre, Rome, Italy

`dipietro@mat.uniroma3.it`

Abstract. The open nature of communications in Wireless Sensor Networks (WSNs) makes it easy for an adversary to trace all the communications within the network. If techniques such as encryption may be employed to protect data privacy (i.e. the content of a message), countermeasures to deceive context privacy (e.g. the source of a message) are much less straightforward. In recent years, the research community addressed the problem of context privacy. Some work aimed to hide the position of the collecting node. Other work investigated on hiding the position of an event—sensed by the WSN. However, the solutions proposed for events hiding either: (i) considered only static events; (ii) are not efficient. In this work, we describe open issues that we identified in the current research. In particular, we consider the problem of efficiently hiding mobile events.

1 Introduction

Due to the open nature of communications in Wireless Sensor Networks (WSNs), it is fairly easy for an adversary to trace all the communications within the network. If techniques such as encryption may be employed to protect data privacy (i.e. the content of a message), countermeasures to deceive context privacy (e.g. the source of a message) are much less straightforward. The research community has recently started addressing this issue in the unique context of WSNs [1]. The resource constrained environment, together with the enhanced adversary capabilities—e.g. the adversary can be mobile and eavesdrop all the communications—, have no correspondence in the wired setting, hence calling for novel solutions to address context privacy issues in WSNs [2].

In many applications (e.g., sensing and reporting the location of a convoy) the source of a message itself reveals the events a WSN is sensing. In order to protect these events, and thus assuring context privacy, we need to conceal that an event took place.

In this paper we identify the following open issues for the current research in this area:

- Open Issue 1. If bogus traffic is used to hide the real one, the adversary success (in capturing a node that routed a real event) might not be just linear with the amount of real event.
 - Open Issue 2. Solving the Open Issue 1 might imply losing the unobservability property.
- Open Issue 3. An enhanced privacy property, $\hat{k}$ -anonymity, could be defined to describe the fact that an event cannot be distinguished between $k-1$ other events, but also requiring the other $k-1$ event to be not real events.
- Open Issue 4. Hiding the trajectory of an event should take into consideration the anatomy of the trajectory itself.
- Open Issue 5. The need for a common metric for privacy and energy consumption.

Organization. In Section 2 we revise the state of the art in the context of privacy preservation in WSN. In Section 3 we introduce the system model and the adversary model considered in this paper. In sections 4, 4.1, 5, 6, and 7 we present different open issues we identified in current research. We conclude our work in Section 8.

2 Related Work

Providing privacy to WSNs would allow a wider application of this technology. WSNs privacy has been first addressed from the point of view of communication confidentiality. Recently, also the context privacy has been investigated—a survey can be found in [2]. WSNs privacy also depends on the specific use of the network. E.g. [3,4] addressed the problem of *privacy preserving* data aggregation [5].

The problem of guaranteeing the privacy of a node sending data to the *BS* (not throughout the aggregation process) has been initially addressed leveraging how the messages are routed. In [6,7], the authors aim at hiding the source of a message, forwarding messages to the *BS* using random walks and dummy traffic. In [8,9] the *BS* location is protected by letting the nodes send messages to a random node, instead of the *BS*. This random node will then aggregate the traffic and communicate to the *BS*. This work consider a local adversary. In [10], the authors propose a partial k -anonymity solution for event source. While they also consider mobile events, the solution is quite energy demanding—property not desirable in WSNs.

In [11] the energy aspect has been taken more into consideration. The authors used *carefully chosen* dummy traffic to conceal the event source location and formalized the concept of unobservability for wireless communication. Nodes acting as aggregator proxies are used to reduce the communication overhead. Another solution involving dummy traffic but not proxies has been proposed in [12]. In [1] the authors demonstrate that to achieve perfect global privacy performance benefit must be sacrificed. They also introduce the notion of strong source anonymity. We observe that the solutions in [1,11] introduce a delay in

the message reaching the *BS*. The solution proposed in [13] switches on demand from a statistically-strong source anonymity scheme (i.e. [1]) to a k -anonymity scheme (i.e. [10]). How to solve the handoff problem in a secure and distributed manner is left as future work.

Finally, randomizing the node ID has also been proposed [14]. However, the adversary model considered does not leverage techniques such as traffic and rate analysis.

3 System and Adversary Model

While previous solutions consider mainly static events, in this work we deal with mobile events. A mobile event involves a set of sensing nodes, in a way that is dependent upon time and location: the set of nodes sensing a mobile events define a handoff trajectory. Hiding the handoff trajectory of mobile events is not yet well investigated [13]. In particular, we deal with a specific type of mobile events: the ones originating on the WSN perimeter (i.e. its border) and eventually ending inside the WSN area.

The adversary is assumed to be global, passive, and external. Furthermore, we take into account the possibility that, once the traffic is analyzed, the adversary is willing to verify the gathered information. This means that the adversary physically checks the locations of both real and dummy events. Constrained by a given time interval T_a , the adversary inspects a subset of the nodes believed accountable for the traffic he previously gathered—the adversary’s revenue is proportional to how many checked locations previously corresponded to a real event.

4 Open Issue 1: A Non-linear Adversary Gain

We want to guarantee the privacy of a real event being sensed by the WSN, considering the models described in Section 3. In particular, we aim to conceal which nodes sensed a real event. In order to this, we constrain other nodes to act as they sensed a real event too, thus sensing a dummy event. The property we want to guarantee can be summarized by the following definition.

Definition 1. (*Real event (T,k) -unobservability*). We define a real event unobservability over the variables T and k . In particular, consider the observation O for a time interval T . The probability of a real event e entering the network with a rate following a Poisson distribution of parameter l , $0 \leq l \leq k$, is equal to the probability of e given O . If this holds for each possible choice of l , then e is called (T,k) -unobservable. Formally:

$$\text{if } \forall O, P(e) = P(e|O) \text{ then } e \text{ is } (T,k)\text{-unobservable.}$$

In other words, just observing the network does not give any information on the Poisson parameter l , $0 \leq l \leq k$ at which real events actually enter into the network.

Furthermore, we have the following aim. Let us assume that the adversary becomes active in tampering nodes—to check if these nodes sensed real events. We want to keep constant the adversary’s success probability. In particular, we want to reach this target even if the rate of real events l varies.

A possible approach. In the literature real events are modeled as Poisson processes of a given ratio l [15]. Given this, a possible solution could leverage a self-adaptive scheme that, given real events taking place with rate l , carefully produces dummy events with rate $k - l$. The parameter k is the rate of the overall events (i.e. real and dummy ones) we assume the WSN can deal with. Hence, we select k such that the rate of real events l is $\leq k$.

If the overall rate k is fixed, the adversary is not able anymore to distinguish whenever a message corresponded to a real event. However, since the scheme would need to adapt to the ratio of real events, an additional parameter T is needed: this parameter defines the time interval in which every node “learns” the actual amount of real events. According to that, each node tunes the rate of dummy events. We call the so defined privacy property (T, k) -unobservability.

In this setting, with the growth of l , the success ratio of the adversary, that checks the event source location, grows as well since there are less dummy events. Given $\phi = k - l$ dummy events and l real events, a naïve solution would be to increase ϕ accordingly. In other words, we may want to keep a fixed ratio between ϕ and l . However, such a solution presents two different concerns: (1) the amount of events a WSN can generate is limited; (2) given a fixed displacement of the network, the more events there are, the more likely is that the adversary may discover a real event.

Fig. 1. The same WSN with two different events densities. The light circle represents the area the attacker may inspect in a given time interval T_a .

The former concern represents also the upper bound of real events the WSN is able to cope with. Since concealing real events requires us to generate additional dummy events, it is our interest to keep this quantity to a minimum. The latter concern is, in turn, depicted in Figure 1 (the legend is reported in Figure 1a).

Figure 1b shows a WSN where two dummy events (dark arrows) and one real event (light arrow) takes place, i.e. $\phi = 2$ and $l = 1$ (light arrow). The adversary, identified by the star, is able to inspect some of the nodes generating dummy traffic within the light circle. We remind that the adversary is constrained by a time interval T_a , therefore he may miss to find the depicted real event.

Let us analyze the following case now: we want to increase the dummy events ($\phi = 4$) in order to deal with a growth of real events ($l = 2$). Figure 1c shows this scenario: the area available to the adversary includes two different real events now. The probability that the adversary chooses the real event location intuitively increases: previously only one node out of eight would have been a successful choice (the node sending a real event). Merely increasing the amount of dummy events in a proportional manner is not the desired solution.

To address these concerns, we have to take into account the density d_e of the nodes relaying these events. Less events trivially correspond to a lower d_e . If we increase the events, the attacker has more nodes to check (i.e. higher d_e).

4.1 Open Issue 2: Proportional Adversary Gain Means Losing Unobservability

The solution reported in the previous section strongly relied on the concept of k -anonymity; since the amount of events was kept constant (k), l real events were concealed among $k - l$ dummy events. As long as k was not changing, real events hitting the network were made unobservable.

However, the more the real events were increasing, the higher was the probability of success of our adversary. To cope with that behavior we proposed to increase the rate of dummy events in a more than proportional manner. Unfortunately this countermeasure also enjoys a side-effect: a passive and global adversary has all the means to infer whether the overall amount of events k changes. This piece of information trivially discloses whether a WSN sensed more or less real events, hence real events do not enjoy the unobservability property anymore. Is it then possible to keep the adversary probability of success proportional while keeping real events unobservable? Is losing unobservability in favor of k -anonymity the right path to provide context privacy w.r.t. a WSN?

5 Open Issue 3: An Enhanced Privacy Property $\hat{k}$ -anonymity

A subject is considered k -anonymous whenever it is concealed among $k - 1$ other subjects. This concept has often been applied to solve the problem of releasing sensitive data-sets [16]: a record was appointed as k -anonymous if the information for each person contained in the release (i.e. quasi-identifier) cannot be distinguished from at least $k - 1$ individuals whose information also appears in the release. Other work (e.g. [17]) applied the same idea to anonymous communication networks: a sender was considered k -anonymous in case an external observer was not able to distinguish which of the k peers was the actual sender.

However, since the adversary physically checks the locations of both dummy and real events, one might desire that not only one event cannot be identified within k possible events, but also that all the others $k-1$ events are bogus ones. We call this privacy property $\hat{k}$ -anonymity. It holds if and only if a real event is concealed among $k - 1$ dummy events.

6 Open Issue 4: Trajectory's Anatomy

In order to conceal to an adversary the location of nodes sensing an event, we outlined in Section 4 an approach that generates dummy events. These dummy events are generated by sending to the sink node the same type of message used whether a real event is sensed. A global adversary, in fact, can not distinguish between messages corresponding to real events and those referred to dummy events. The context privacy of a real event is therefore assured.

Fig. 2. Two different type of trajectories: a panda and a vehicle

However, since we deal with mobile events, we need to take into account which nodes are cooperating to build a dummy mobile event. In other words, since a mobile event is supposed to span across different nodes, a real event is expected to exhibit a trajectory. Therefore, whenever we generate a dummy event, a proper set of nodes must be chosen in order to produce a trajectory that resembles a real one.

This problem can be exemplified in Figure 2 by analyzing the trajectories produced by an endangered panda and by a vehicle. We do not expect a vehicle to visit multiple times the same locations; a panda instead is more prone to visit locations, such as source of water, which have already been visited. Trying

to anonymize a vehicle with trajectories produced by a panda splits the set of sensed events in two different subsets: the subset of dummy events and the one of real events. Once the adversary obtains the opportunity to correctly label these two subsets, the privacy of real events can not be any longer assured. The type of events a WSN is supposed to deal with is therefore an important piece of information for any sound and secure solution.

7 Open Issue 5: A Metric for Privacy and Energy Consumption

Wireless Sensor Network is a novel field where the consumption of energy is often a system requirement. Moreover, since all the nodes have a physical location, the initiator and the recipient of any communication may become sensitive assets. A typical scenario is a WSN sensing a panda threatened by poachers: the node that initiates the communication trivially discloses the location of the panda; likewise, the recipient discloses where all the information is eventually collected. In both cases an adversary can be interested in the disclosed information.

A Privacy Enhancing Technology (PET) has an immediate impact on the battery life of any type of device. Data privacy, for instance, is often provided by means of some security primitives to encrypt and decrypt the exchanged data. No matter which primitives are chosen, the node's CPU becomes accountable of any additional computation. This means that any node sending or receiving encrypted messages will suffer from a reduced battery life.

However, as mentioned before, applications relying on WSNs advise for novel solution in the field of contextual privacy. Since what has to be concealed are the nodes taking part in a communication, any PET has to generate additional traffic to somehow anonymize the real communication. In case of an active adversary, we point out the following trade-off: the more dummy traffic is generated, the more private the real communication is; consequently the more energy must be consumed in the process.

Providing any sort of PET is therefore a rather expensive task. Existing works [18] proposed approaches to model the consumption of energy in terms of the adopted security primitive. What has not been yet proposed is a general model able to provide some reasoning to evaluate the proposed PETs. In particular, given a PET, we shall be able to assess the level of privacy the proposed solution provides, and the overhead produced in terms of energy consumption.

The latter is a rather interesting problem: since all the messages are eventually delivered to the sink node, we shall not expect an evenly distributed energy consumption. Instead, we may expect nodes close to the sink node to handle a higher rate of messages if compared to nodes lying on the perimeter. Figure 3 depicts exactly this behavior: nodes in the red circle are supposed to send more messages than nodes lying in a more external area. This kind of behavior creates what are known to be hot spots, i.e. areas where the consumption of energy is expected to be higher. We believe a PET shall take into the existence of these special areas, and mitigate their rate of occurrence.

Fig. 3. Hot spot induced by the sink node being the ultimate recipient of any message

8 Conclusion

In this paper we described different open issues in context privacy of Wireless Sensor Networks (WSNs). In particular, we considered the problem of hiding mobile events—e.g. presence of animal or vehicles—that are sensed by the network itself. We observed how current solutions designed to hide static events either (i) are not able to hide mobile events, or (ii) are not efficient. We think that the solution of the open issues presented in this paper would remove barriers for a wide adoption of WSNs. Our future works aim to solve these issues.

References

1. Shao, M., Yang, Y., Zhu, S., Cao, G.: Towards statistically strong source anonymity for sensor networks. In: Proceedings of the 27th IEEE International Conference on Computer Communications (INFOCOM 2008), pp. 51–55 (2008)
2. Li, N., Zhang, N., Das, S., Thuraisingham, B.: Privacy preservation in wireless sensor networks: A state-of-the-art survey. *Ad Hoc Networks* 7, 1501–1514 (2009)
3. He, W., Liu, X., Nguyen, H., Nahrstedt, K., Abdelzaher, T.: Pda: Privacy-preserving data aggregation in wireless sensor networks. In: Proceedings of the 26th IEEE International Conference on Computer Communications (INFOCOM 2007), pp. 2045–2053 (2007)
4. Conti, M., Zhang, L., Roy, S., Di Pietro, R., Jajodia, S., Mancini, L.V.: Privacy-preserving robust data aggregation in wireless sensor networks. *Security and Communication Networks* 2, 195–213 (2009)
5. Xi, Y., Schwiebert, L., Shi, W.: Preserving source location privacy in monitoring-based wireless sensor networks. In: 20th International Parallel and Distributed Processing Symposium (IPDPS 2006), pp. 77–88 (2006)

6. Kamat, P., Zhang, Y., Trappe, W., Ozturk, C.: Enhancing source-location privacy in sensor network routing. In: Proceedings of 25th IEEE International Conference on Distributed Computing Systems (ICDCS 2005), pp. 599–608 (2005)
7. Ozturk, C., Zhang, Y., Trappe, W.: Source-location privacy in energy-constrained sensor network routing. In: Proceedings of the 2nd ACM workshop on Security of Ad hoc and Sensor Networks, pp. 88–93 (2004)
8. Conner, W., Abdelzaher, T., Nahrstedt, K.: Using data aggregation to prevent traffic analysis in wireless sensor networks. In: Gibbons, P.B., Abdelzaher, T., Aspnes, J., Rao, R. (eds.) DCOSS 2006. LNCS, vol. 4026, p. 202. Springer, Heidelberg (2006)
9. Jian, Y., Chen, S., Zhang, Z., Zhang, L.: Protecting receiver-location privacy in wireless sensor networks. In: Proceedings of the 26th IEEE International Conference on Computer Communications (INFOCOM 2007), pp. 1955–1963 (2007)
10. Mehta, K., Liu, D., Wright, M.: Location privacy in sensor networks against a global eavesdropper. In: Proceedings of the IEEE International Conference on Network Protocols (ICNP 2007), pp. 314–323 (2007)
11. Yang, Y., Shao, M., Zhu, S., Urgaonkar, B., Cao, G.: Towards event source unobservability with minimum network traffic in sensor networks. In: Proceedings of the 1st ACM Conference on Wireless Network Security (WiSec 2008), pp. 77–88 (2008)
12. Ouyang, Y., Le, Z., Liu, D., Ford, J., Makedon, F.: Source location privacy against laptop-class attacks in sensor networks. In: Proceedings of the 4th International Conference on Security and Privacy in Communication Networks (SecureComm 2008), pp. 1–10 (2008)
13. Yang, Y., Zhu, S., Cao, G., LaPorta, T.: An active global attack model for sensor source location privacy: Analysis and countermeasures. In: Chen, Y., Dimitriou, T.D., Zhou, J. (eds.) SecureComm 2009. Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering, vol. 19, pp. 373–393. Springer, Heidelberg (2009)
14. Ouyang, Y., Le, Z., Xu, Y., Triandopoulos, N., Zhang, S., Ford, J., Makedon, F.: Providing anonymity in wireless sensor networks. In: IEEE International Conference on Pervasive Services, pp. 145–148 (2007)
15. Diaz, C., Seys, S., Claessens, J., Preneel, B.: Towards measuring anonymity. In: Dingledine, R. (ed.) PET 2003. LNCS, vol. 2760, pp. 184–188. Springer, Heidelberg (2003)
16. Sweeney, L.: k-anonymity: A model for protecting privacy. *International Journal of Uncertainty Fuzziness and Knowledge Based Systems* 10, 557–570 (2002)
17. von Ahn, L., Bortz, A., Hopper, N.J.: k-anonymous message transmission. In: Proceedings of the 10th ACM Conference on Computer and Communications Security (CCS 2003), pp. 122–130 (2003)
18. Potlapally, N., Ravi, S., Raghunathan, A., Jha, N.: A study of the energy consumption characteristics of cryptographic algorithms and security protocols. *IEEE Transactions on Mobile Computing* 5, 128–143 (2006)