Confidentiality In Sensor Networks: Transactional Information

Sameer Pai*, Marci Meingast†, Tanya Roosta†, Sergio Bermudez*
Stephen Wicker*, Deirdre K. Mulligan‡, Shankar Sastry†
*School of Electrical and Computer Engineering
Cornell University, Ithaca, NY 14853
{sameer.pai, sab222}@cornell.edu & wicker@ece.cornell.edu
†Department of Electrical Engineering and Computer Science
University of California, Berkeley, CA 94704
{marci, roosta, sastry}@eecs.berkeley.edu
‡Boalt Hall School of Law
University of California, Berkeley, CA 94704
dmulligan@law.berkeley.edu

Abstract

In a sensor network environment, elements such as message rate, message size, mote frequency, and message routing can reveal information about the sensors deployed, frequency of events monitored, network topology, parties deploying the network, and location of subjects and objects moving through the networked space. Collectively, we refer to these elements as transactional data. Where the confidentiality of the content of the networks communications has been secured through encryption and authentication techniques, the ability of network outsiders and insiders to observe elements or the totality of this transactional data can also compromise network confidentiality. This paper describes four types of transactional data typically observable in sensor networks and discusses the information that can be derived through its observation and analysis. The paper argues that measures to limit the availability and utility of transactional data are essential to preserving confidentiality in sensor networks.

I. Introduction

Sensor networks consist of numerous tiny autonomous wireless devices with limited energy and memory which are used for monitoring physical phenomena. The data they gather can be analyzed to extract important information regarding events, objects and individuals. Sensor networks, an emerging technology, are destined to play an important role in monitoring people, objects and infrastructure for purposes ranging from environmental assessment and research, in-home patient care, disaster mitigation, energy demand and response, inventory monitoring, surveillance and law enforcement.

According to the International Organization for Standardization (ISO), confidentiality is defined as the assurance that information is only accessible to those authorized to have access. Confidentiality is provided through policies and practices that ensure that information flows only to those who are entitled to access it. In a network, such as a sensor network, confidentiality requires the development of rules governing access to transactional data and technical measures that enforce those rules and prevent adversaries[1] from violating them. In many cases, preserving network confidentiality is a first step in preserving the privacy of the users and deployers of the sensor network, and, where individuals are present in the network space, their privacy and safety as well[2]. Transactional data can provide an adversary insight into the likely content of the transmissions, as well as divulging information about the network's purpose, its deployers, and the objects and subjects of observation that authorized users of the network

This work was supported in part by TRUST (The Team for Research in Ubiquitous Secure Technology), which receives support from the National Science Foundation (NSF award number CCF-0424422) and the following organizations: Cisco, ESCHER, HP, IBM, Intel, Microsoft, ORNL, Qualcomm, Pirelli, Sun and Symantec. The first author is also supported by a National Science Foundation IGERT Fellowship.

[1]Adversaries are those entities which do not have authorization to access the network.

[2]In this paper, the sensor network users, deployers, and individuals that are present in the network space are collectively referred to as the stakeholders of the sensor network.

may wish to keep confidential from one another and outsiders. As the applications for sensor networks continue to grow, it is of increasing importance that risks to confidentiality posed by the availability of transactional data be identified, and mitigations and solutions devised.

It is widely understood that traffic analysis can provide information about the users and usage of a network, separate and above what may be provided through the content of the data packets. The risks posed by the availability of transactional data have been studied in other communication networks including traditional telecommunications and packet switched networks such as the Internet, and solutions have been proposed. For example, in [7], the authors propose a general approach to prevent the traffic pattern of an IP-based network from being analyzed by an intruder. Privacy law protecting electronic communications recognizes, and provides protections for, transactional data, for example IP-headers and signaling and routing information, although the protections are substantially weaker than those afforded to the content of communications [9], [6]. While providing a taxonomy and framework useful for considering the privacy implications of transactional data, existing law is unlikely to apply to the majority of sensor networks. While there is a wide body of engineering literature regarding content confidentiality in sensor networks, there has been little discussion of transactional confidentiality.

II. Content Confidentiality in Sensor Networks

Content in a sensor network is defined as the meaning of the information contained within the data that are being communicated in the network. Content confidentiality entails disallowing those external to the network from inferring the meaning of the messages being sent.

Content confidentiality in sensor networks has been extensively discussed in current literature[3]. In the engineering community there is a wide body of research dealing with the protection of content confidentiality in sensor networks. A nice summary of available literature discussing this topic is given in [10].

Two of the important methods developed for protecting content confidentiality in sensor networks are authentication and data encryption. For example, in order to perform data encryption, motes[4] in a sensor network can use pre-distritbuted keys or use some method to generate the keys dynamically. This is accomplished by utilizing a *key management* protocol. A key management protocol falls into one of three categories: deterministic, probabilistic, and hybrid[5].

III. Transactional Confidentiality in Sensor Networks

In contrast to content, transactions in a sensor network relate to the information gathered as a consequence of generation, transmission, and routing of data messages within the network (e.g. by an adversary conducting traffic analysis). Therefore, although the sensor network could have different methods in place for content confidentiality protection (e.g. encryption and authentication), the transactional confidentiality of the network can be jeopardized. As a result, transactional confidentiality protection within a sensor network involves preventing adversaries from learning information based on the creation and flow of the messages within the network.

Existing research on transactional confidentiality examines the risks posed by message rate and message routing analysis. In [5] and [4], the main focus is on the transactional confidentiality associated with the routing of messages and the message transmission rate within a sensor network. These papers discuss methods by which an adversary can determine the location of a source by violating transactional confidentiality. The papers also discuss possible solutions to these problems using altered routing and rate-control algorithms. However, they do not explore all the aspects of how transactional confidentiality can be violated. In [1], the problem of confidentiality for information about the location of a sensor network sink is discussed. As in [5] and [4], the authors of [1] present two examples where transactional confidentiality can be breached in a sensor network and countermeasures. These examples involve an adversary acquiring transactional information from message routing and the message transmission rate in the sensor network. These papers provide an incomplete analysis of the information that can be gleaned through analysis of available transactional data.

[3]This is referred to as *content privacy* in some engineering literature.

[4]A mote is a small wireless electronic communication platform combined with a set of transducers to sense phenomena in the region surrounding the platform.

[5]In the deterministic protocol the key chains are selected deterministically and in the probabilistic case the keys are chosen randomly from a given key pool and distributed among the motes.

The literature from law, policy and engineering has identified the identity of communicators as a particularly important piece of transactional data. To promote anonymity within networks, the identity and source of information must be protected against disclosure[6]. Research into technical means of protecting anonymity in digital networks has been conducted over the past two decades. One example of this research is Onion Routing which aims to provide anonymity to those sending and receiving packets on the Internet (See tor.eff.org). Along the same lines, there is research pertaining to anonymity in wireless sensor networks [14]. This work develops a solution to protect a source's identity in the network. The information protected by this solution only relates to a small subset of the problems pertaining to transactional confidentiality. We note that in sensor networks, the problem of anonymous communication has still not been examined in depth and solutions such as onion routing which rely on layers of encryption may demand processing power that is unavailable in many sensor networks.

The privacy implications of the availability of location information in cellular and global positioning satellite based networks and through providers operating on those networks has been addressed through federal law. Technologists have also identified location information as particularly sensitive and have undertaken efforts to consider technical and policy mechanisms to provide users of networked devices with control over its disclosure and use (See http://www.ietf.org/html.charters/geopriv-charter).

In this paper we provide a more complete overview of the transactional data available in sensor networks and a richer picture of the information that can be derived through its analysis. We then use deployed sensor networks to illustrate the sorts of interests, including privacy of subjects operating in the networked space, that can be compromised through analysis of transactional data.

IV. Potential Transactional Confidentiality Risks in Sensor Networks

The issues involved with transactional confidentiality are inherently more complex than those of content confidentiality given that an adversary can learn information just by the mere presence of a message being transmitted. For example, transactional confidentiality issues arise when sensors generate messages as a result of monitoring the physical infrastructure of buildings. These sensor networks may be designed for use in monitoring room temperature or light conditions. However, by using the transactional information from the network, an adversary could also locate people in the building.

The complexity of protecting transactional confidentiality is due to the protocols designed to operate as means of communication among networked motes. In order to maintain transactional confidentiality, these protocols must have transactional data protections integrated into their design from the very beginning. Thus, for many protocols, there is a necessity for re-design in order to protect transactional confidentiality. In this section we present four ways in which transactional confidentiality can be compromised: carrier frequency, message rate, message size, and routing. Some aspects of message rate and routing have been discussed in previous research (e.g. see [1]). However, in this section we describe additional aspects of message rate and routing which can give transactional information to an adversary. We also examine how carrier frequency and message size can be used to breach transactional confidentiality.

When relating transactional confidentiality and sensor networks, one important question arises: is it possible to relate particular confidential information to a stakeholder or a group of stakeholders? The answer is more obvious when talking about the Internet, where Internet Protocol (IP) addresses can be related back to a specific stakeholder(s). Currently, sensor networks are moving towards using IP addresses in network layer routing. For example, Arch Rock Corporation recently announced an IP-based wireless sensor network[7]. In sensor networks that do not rely on IP addresses, machine learning algorithms, using computational and statistical methods, could be used to extract information that might allow for the identification of a stakeholder.

Using the sensor network message rate, message size, mote frequency, and message routing, an adversary can infer a great deal of transactional information. Next, we will discuss in greater detail these ways in which an adversary can acquire information about the sensors that are deployed, the frequency of events being monitored, the network topology, the parties which deployed the sensor network, the parties and objects that the network is monitoring, and their location.

[6]Network anonymity differs in two ways from network confidentiality: 1) it is concerned with the anonymization of the network source and 2) it is particularly concerned with not allowing entities within the network from learning source identifying information from network data.

[7]$www.archrock.com/news_events/press_releases/2007.03.19.php$

(a)

Commonly Used Name	Frequency Range (MHz)
MICA or MICA1	902 to 928 433.1 to 434.9
MICA2	868 to 870 902 to 928 433.1 to 434.8 313.9 to 316.1
MICA2DOT	868 to 870 902 to 928 433.1 to 434.8 313.9 to 316.1
MICAz	2400 to 2483.5
Cricket	433.1 to 434.8
IRIS	2400 to 2483.5
TelosB	2400 to 2483.5

(b)

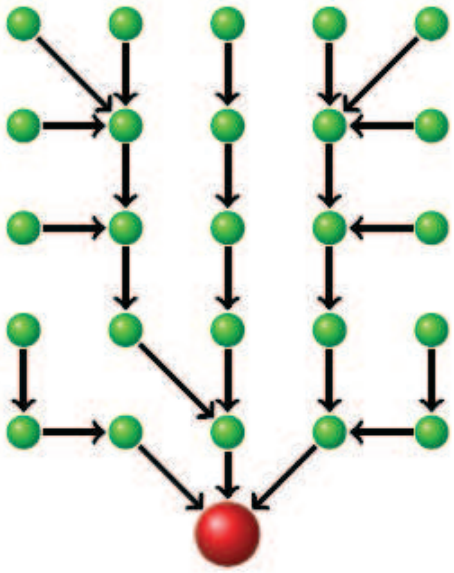

Fig. 1. Figures showing methods by which an adversary can access transactional information (a) The radio specification of different sensor platforms. (b) An example of a traffic pattern in sensor networks. Green dots represent sensor motes and the red dot represents the base station.

A. Carrier Frequency Examples

Carrier frequency is one area that can be used to compromise transactional confidentiality. It has been shown in [13] that frequency-based side channel attacks can break cryptographic schemes, violating content confidentiality. However, this type of attack can also be used to breach transactional confidentiality in sensor networks. It is conceivable for an adversary that uses a spectrum analyzer to detect the carrier frequency of the data communicated by the motes. This compromises the confidentiality of the network since different sensor platforms use different radio frequencies for communication. There are many mote platforms, both commercial and non-commercial, which each communicate on different fixed carrier frequencies. For example, the table in Figure 1(a) shows the carrier frequency of different commercial mote platforms from Crossbow technologies. Once the adversary finds out which carrier frequency these motes are operating at, it can figure out the hardware platform used. Knowing the hardware platform enables the adversary to exploit the software vulnerabilities of the particular version of the software that runs on that mote platform, which allows the possibility of a confidentiality attack.

In addition, an adversary looking at the carrier frequency can help it determine what sort of devices are in operation and who might be using the network. By definition, acquisition of information about the sensor network architecture and the network stakeholders by an adversary constitutes a breach in transactional confidentiality. For example, in the United States, the 300-420 MHz band is designated for government use, including meteorology and military use. Motes operating in this band would be clear indicators of a government network. The 2400-2483.5 MHz bands are used for Industrial, Scientific, and Medical bands (ISM); IEEE 802.11a, 802.11b, and 802.11g Wireless LAN; and IEEE 802.15.4. In the last example, an adversary can determine both the type of mote platforms being used as well which group of individuals that deployed these platforms by detecting the network carrier frequency of operation. Therefore, by looking at the carrier frequency, transactional confidentiality of the network could potentially be violated.

An adversary can acquire confidential information about the sensors that are deployed or the parties which deployed the sensor network using knowledge of the carrier frequency. In use with additional information, such as routing schemes or message size, the carrier frequency can be used to further violate transactional confidentiality. For example, if from the carrier frequency an adversary sees that it is a government network, then looking at the

routing scheme the adversary can determine where the data is being transmitted to, thus violating transactional confidentiality.

B. Message Rate Examples

The rate at which messages are generated by motes in a sensor network can relay information about the nature of the event they are monitoring. Specifically the message rate can reveal information about the frequency of events monitored and the location of subjects or objects that are in the networked space. Sensor networks can use fixed rate or event driven communication protocols. In both cases, transactional information can be gathered by an adversary.

Since conserving battery power is an important consideration in sensor networks, event-driven protocols are used to conserve the energy of the motes. Event-driven sensor network applications are such that mote data transmission is triggered by specific events (e.g. a large change in room temperature). The rate at which these messages are generated can give an adversary specific information about the regularity of certain events being monitored by the sensor network [4]. This constitutes a breach in network confidentiality. For particular applications of sensor networks (e.g. health monitoring or in-house temperature monitoring for electric power demand-response), this breach can also lead to a breach in privacy for the stakeholders of the sensor network.

Fixed rate communication protocols are also sometimes used in sensor networks. In particular applications and designs of sensor networks, the motes in the network report information periodically with a fixed period. This type of communication also allows an adversary to breach transactional confidentiality.

As a specific example, in some tracking schemes, as an object moves from one location to another, the rate at which the motes generate their messages will decrease in one region and increase in another region. An adversary can observe the rate at which sensor messages are being generated by the motes in its neighborhood, and then move toward the mote that has a higher message rate. Moving towards the mote that has a high message rate will lead the adversary to the source, i.e. the mote closest to the geographic location of the object [5]. In this case the moving object can be located by the adversary, creating a confidentiality breach and breach of the privacy of the sensor network stakeholders.

When monitoring a mote, an adversary can examine the rate at which it receives messages to determine its radial distance from the mote [4]. An adversary with knowledge of the message rate, for a particular application of sensor networks, can compare this with the rate at which it receives messages. In many cases the message reception rate decreases with increasing distance between the adversary and the mote reporting the information. This allows the adversary to breach network confidentiality. For particular applications of sensor networks, motes are attached to an object that is moving in the network. For example, in the health monitoring applications, body sensors are used on a patient's body to monitor vital signs, such as heart beat. In many designs within this context, the networked motes send messages at a fixed rate. Thus, an adversary can determine the distance between this mote and itself. This can allow the adversary to determine the location of the object which has the motes attached. Therefore, this breach can also constitute a breach in privacy for the stakeholders of the sensor network.

We note that in either of these examples, if the adversary has access to multiple transceivers in different locations within the sensor network, then it can determine the location of a monitored object or individual through triangulation. For instance, using its set of transceivers, the adversary can determine the rate at which it receives messages from fixed or mobile motes. From the different rates that it observes and based on the locations of its transceivers, the adversary can calculate the position of objects or individuals being monitored by the sensor network [3].

To validate this idea, we conducted an experiment using two Crossbow MICAz motes. The MICAz motes have a maximum data rate of 250kbps and a maximum range of about 100 meters. In our experiment, one of the motes acted as a transmitter and another mote was used to intercept the transmitted messages. The motes were placed on top of an empty parking garage that was 150 meters long and 20 meters wide. The transmitting mote was placed at a variable distance away from the intercepting mote. The transmitting mote sent a message every 8 seconds or 7.5 messages per minute. The message reception rate was calculated at the intercepting mote at distances of 10 meters, 30 meters and 50 meters allowing for a period of 5 minutes at each distance. The results of this experiment, as can be seen in Figure 2, show that message reception rate falls with increasing distance. Therefore, an adversary using a similar intercepting device can determine the radial distance of a mote transmitting at a fixed rate.

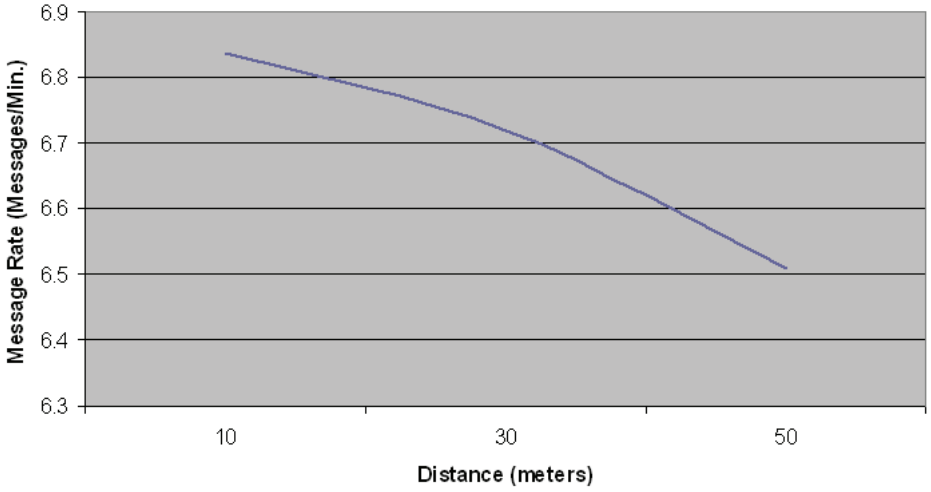

Fig. 2. The message reception rate versus the distance of the receiver from a transmitting mote.

C. Message Size Examples

The size of messages being routed through a sensor network can give information about the network topology, the events being monitored, as well as what the network is monitoring and location of subjects or objects within the network space. For instance, routing protocols where data gets concatenated or aggregated as it moves from a source mote to a sink mote can provide transactional information for an adversary. Assuming a smart adversary, we expect that it can infer multiple things with the changing size of data messages being transmitted through the sensor network. Protocol designs for automatically changing message size based on the amount of data reported by motes in the sensor network are suggested in [2]. In these protocols, there is an assumption that messages are routed with varying sizes in order to decrease energy spent in communication. However, there are two immediate impacts that these protocols can have on confidentiality:

1) An adversary can infer information about the type of data being communicated and as a result what the network is monitoring.
2) An adversary can infer information about the location of a central sink or fusion center

The first of these confidentiality breaches is due to the nature of the data itself. If the amount of data being collected is small (post-compression), and if the message size is small, due to protocol design, the adversary could infer that the network is collecting slow-varying data (e.g. coarse daily human temperature readings). If, however, data size is large, then message size is large, and the adversary could infer that the network is collecting fast-varying data (e.g. acceleration) [1], [2], [12]. This information can be used by an adversary to determine what is being monitored by the network [1], [2]. This form of data analysis by an adversary breaches the confidentiality of the network. Once again, as with the message rate examples, this breach of confidentiality could result in privacy problems for the stakeholders of the sensor network.

The second of these confidentiality breaches is due to the process of data aggregation in the sensor network in conjunction with varying message size. In-network aggregation of data is suggested as an energy saving method in sensor networks [2]. Certain methods of data aggregation could lead to confidentiality breaches. For example, if the network employs shortest-path aggregation, motes furthest from a fusion center, or sink mote, will rarely participate in aggregating data. However, motes closest to the fusion center or the sink will almost always participate in data aggregation. If message size grows with the amount of data aggregation at a mote, an adversary could gain information about the location of the source motes, fusion centers, or network sinks based on the size of the messages generated around these motes. This could lead to potential security and confidentiality breaches for the network. In particular applications of sensor networks, for instance in tracking, knowing the location of an important mote in the network (e.g. the source mote) could lead to a breach in privacy for the stakeholders of the sensor network. This connection will be made more clear in the routing examples given in subsection IV-D as well as the real-world examples given in Section V.

D. Routing Examples

As mentioned earlier, sensor networks have a special topology. When an event occurs, the motes send their messages to a base station or a number of base stations. This means that the routing pattern is usually many-to-one or many-to-few. Figure 1(b) shows an example of the many-to-one traffic pattern in sensor networks.

Using this knowledge, an adversary can therefore observe the traffic pattern through the sensor network and deduce the location of the base station or other strategically located motes. Furthermore, this can reveal information about the location of subjects or objects moving through the networked space. For example, if shortest path routing is used, the adversary can monitor the motes in its neighborhood and move from one mote to another by following the progression of the forwarded messages (e.g. using angle of message arrival). As a result, the location of important motes in the sensor network are compromised [5]. In this case, the adversary can learn the topology of the network. Furthermore, as a direct consequence of the adversary learning location of important networked motes, in tracking applications of sensor networks, an adversary can also learn the location of the object or person being tracked. In particular, due to the multi-hop communication protocols used by many sensor networks, an adversary could trace a stream of messages, one hop at a time, back to the source of information (i.e. a sensor mote or motes reporting the location of an object or person). Using this method, the adversary is able to determine the location of the object or person being tracked by the sensor network. Location based information can be used to uniquely identify an object or a person, which violates the privacy of sensor network stakeholders.

It is important to note that the extent to which an adversary can gain transactional information about the structure of a sensor network has direct correlation to the routing algorithm used [1], [5]. Therefore, when deciding on what routing algorithm to use, we need to take into account transactional confidentiality of the sensor network, and choose the routing algorithm wisely. Due to the potential of transactional confidentiality breaches, the privacy of network stakeholders could be compromised (e.g. if the network is used to track individuals).

There have been some routing algorithms suggested in the literature that are specifically designed to mitigate the problem of traffic analysis attacks. In [1], the authors use different methods to decorrelate the traffic pattern. One of the solutions proposed is based on the idea of a Random Walk. When a mote receives a message, it forwards the message to one of its parents with probability p, and it uses a random forwarding algorithm with probability $1-p$. In random forwarding algorithm, the mote selects one of its neighbors with equal probability to forward the message.

A second solution proposed in [1] is to use *Fractal Propagation*. In this method, several fake messages are generated and propagated throughout the network. The objective of having these fake messages is to disguise the true messages. When a mote hears one of its neighbors forwarding a *real* message, it generates a fake message with probability p_f and sends it to one of its neighbors. The combination of Random Walk and Fractal Propagation creates more randomness in the traffic pattern of the network. This reduces the amount of transactional information an adversary can obtain from message routing.

One drawback of using these schemes is the amount of energy overhead they introduce which results in a reduced lifetime of the sensor network. Therefore, designing energy-efficient routing methods that maintain transactional confidentiality in the sensor network are an important future area of research.

V. Examples of Transactional Confidentiality Infringement

Previous research has not discussed the real world examples of sensor network applications which might suffer from violations of transactional confidentiality. In this section, we analyze the potential for transactional confidentiality breaches in some real world applications of sensor networks.

A. UVA's ALARM-NET Wireless Sensor System for Assisted-Living and Residence Monitoring

The University of Virginia's (UVA's) ALARM-NET system, described in [15], is based on the use of a multi-tiered heterogeneous sensor network designed to monitor people. The application of ALARM-NET is that of having an automatic real-time monitoring system, deployed in an assisted-living facility or residence, to report on the residents' health, activity, and environment. Specifically, body sensors and motes collect human vital sign data. Stationary sensors and motes are used to collect data pertaining to resident motion and environmental conditions. Although content confidentiality has been considered, there remain multiple concerns with regard to transactional confidentiality that have not been considered in ALARM-NET.

The ALARM-NET system employs three network tiers. The first tier consists of a set of mobile body networks made up of Crossbow MICAz motes, each designated for a particular resident in the assisted living facility. The second tier consists of stationary MICAz motes in different rooms of the residence. The last tier is an IP based network comprised of Crossbow Stargate processor nodes running Linux. The IP-based networked nodes are connected to a personal computer and a personal digital assistant for user-interface and a back-end database for information storage. The data are accessible to a health care provider, which monitors the residence.

The data from the body sensors is broadcast, single-hop to the nearest stationary sensor motes in the second tier of the ALARM-NET system. The second tier of motes use multi-hop communication based on the standard TinyOS 1.1.15 configuration to send and forward data from the body sensor network to a node in the IP based network [15]. These second tier MICAz motes employ a shortest-path-first routing algorithm to a single root mote. In this algorithm, any mote chooses the closest parent using bi-directional link estimation. The overall shortest-path route to the destination is then determined based on these choices. This routing algorithm faces the same drawbacks as multiple other schemes in terms of transactional confidentiality protection (see discussion of transactional confidentiality and routing above in Section IV-D).

The ALARM-NET architecture is such that body-sensor mote data reporting is event-driven. Thus, the rate at which messages are transmitted is directly dependent on the rate at which a particular event is generated. This leaves the ALARM-NET system susceptible to adversarial confidentiality attacks that would try to acquire the location of the resident wearing the body sensors or the type and nature of the event being monitored by the sensors (e.g. if an event is triggered due to a large fluctuation in heart-rate, see Section IV-B).

Another drawback of ALARM-NET is the use of open-system architecture and communications. The system is comprised of commercially available MICAz radio-processor boards [15]. These systems are susceptible to carrier-frequency based transactional confidentiality attacks by an adversary, such as those described in Section IV-A.

From the above analysis, there seem to be at least three types of transactional data a network adversary can acquire when ALARM-NET is employed. This could lead to the adversary infringing upon the privacy rights of those individuals whose health and vital signs are being monitored by ALARM-NET, the users who are acquiring the data, as well as the health-care provider. This could allow the adversary to acquire private information about the resident being monitored and perhaps the health-care provider. In particular, information regarding the health of a resident can be assessed by an adversary based on a combination of location information (from the message rate, routing, and mote carrier frequency based data) in combination with the regularity or irregularity of particular events being monitored by the body sensor network (from the message rate). The information that can be acquired could potentially be valuable for insurance companies which insure a particular resident or even the employer(s) of the resident. The adversarial acquisition of this type of information could hold the health care provider, that is employing the ALARM-NET system, liable for breach of the Health Insurance Portability and Accountability Act (HIPAA).

B. Great Duck Island, habitat monitoring

An example of sensor networks for habitat monitoring purpose is the Great Duck Island, Maine, deployment, which is used to monitor a species of seabirds. Avian scientists needed to collect various measurements pertaining to the nesting environment and behavior of these seabirds. Specifically, the motes measured: light, temperature, relative humidity, and barometric pressure. The system was comprised of wireless motes, gateways, and a base station. The motes were carefully placed into the nests of the seabirds and their surroundings. The base station, connected to the Internet, collected all data from the gateways and uploaded it to several remote databases, [8].

In the design of this system, confidentiality was not considered. The scientists that deployed the sensor network expected the information they collected to be kept private. An adversary (e.g. a poacher) could deduce, for example, the presence of a bird in a particular place. With the transactional data generated by the sensor network, an adversary could obtain confidential information about activities of the birds. If this same technology is modified for use on monitoring people in public or private spaces, breaches in transactional confidentiality will lead to additional privacy issues [11].

This project utilized Mica motes, which are subject to potential transactional confidentiality breaches as a result of frequency strikes, as described in IV-A. An adversary with appropriate measurement equipment can scan the radio frequency spectrum to find out the frequency in use by the network. Then it can deduce the hardware being

used, due to the dependency between frequency and hardware. In addition, it might even be able to discover the TinyOS software version and, further, exploit a bug in that specific version.

In this system, the messages for specific types of sensed data are always the same size. Taking advantage of this protocol design and knowledge of the timescale over which different phenomena are likely to change (e.g. temperature changes occur slowly), an adversary can make estimates of certain variables of the monitored phenomena by making comparisons between different message sizes, as explained in Section IV-C. For example, slow changes with small message sizes over time might indicate barometric pressure, the temperature, or relative humidity being measured by the network. Fast changes through the day with large message sizes might imply solar radiation.

For instance, an adversary might use the temperature, to discover the presence of the birds under surveillance, in this case, even where the animal is located, allowing the adversary to capture it. Furthermore, an adversary can take advantage of the routing tree formed by the routing protocol used in this system. The adversary can do a route traceback to a source mote sending data and thereby determine the location of a particular bird and its nest, as explained in Section IV-D. This again, violates not only the confidentiality of the network, but also jeopardizes the information that the scientists want to keep private.

VI. Conclusions

While content confidentiality is often considered in sensor network design, it is also important to consider transactional confidentiality. Information learned by an adversary as a result of the data flowing through the network, regardless of the content of the data, is a concern. We have defined what transactional confidentiality encompasses and analyzed how it can be used for breaching confidentiality in sensor networks. Routing paths, message size, message rate, and carrier frequency are all forms of transactional data that we have shown can lead to privacy violations for users of the network.

Currently, it is possible to provide some safeguards against content confidentiality breaches in sensor networks. However, it is not straightforward how to protect transactional confidentiality in sensor networks. As we have shown, this is a complex issue. There remains much work to be done in developing schemes to limit the amount of transactional information that can be gathered by an adversary.

References

[1] J. Deng, R. Han, and S. Mishra. Countermeasures against traffic analysis attacks in wireless sensor networks. *The Third ACM Workshop on Security of Ad Hoc and Sensor Networks (SASN)*, 2005.

[2] T. He, B. Blum, J. Stankovic, and T. Abdelzaher. Aida: Adaptive application-independent data aggregation in wireless sensor networks. *ACM Transactions on Embedded Computing System, Special issue on Dynamically Adaptable Embedded Systems,*, 2004.

[3] J. Hightower and G. Borriella. Location systems for ubiquitous computing. *IEEE Computer*, 34(8):57–66, 2001.

[4] P. Kamat, W. Xu, W. Trappe, and Y. Zhang. Temporal privacy in wireless sensor networks. *ICDCS '07: Proceedings of the 27th International Conference on Distributed Computing Systems*, 2007.

[5] P. Kamat, Y. Zhang, and W. Trappe. Enhancing source-location privacy in sensor network routing. *Proceedings of the* 25^{th} *IEEE International Conference on Distributed Computing Systems*, 2005.

[6] O. S. Kerr. A user's guide to the stored communications act - and a legislator's guide to amending it. *George Washington Law Review*, 72, 2004.

[7] D. Liu, C.-H. Chi, and M. Li. Normalizing traffic pattern with anonymity for mission critical applications. *37th Annual Simulation Symposium, 2004. Proceedings.*, April 2004.

[8] A. Mainwaring, D. Culler, J. Polastre, R. Szewczyk, and J. Anderson. Wireless sensor networks for habitat monitoring. *WSNA '02, Proceedings of the 1st ACM international workshop on Wireless sensor networks and applications*, 2002.

[9] D. K. Mulligan. Reasonable expectations in electronic communications: A critical perspective on the electronic communications privacy act. *George Washington Law Review*, 72, 2004.

[10] A. Perrig, R. Szewczyk, J. Tygar, V. Wen, and D. Culler. Spins: security protocol for sensor networks. *Wireless Networks*, 2002.

[11] R. Szewczyk, E. Osterweil, J. Polastre, M. Hamilton, A. Mainwaring, and D. Estrin. Habitat monitoring with sensor networks. *Communications of the ACM*, 2004.

[12] P. Thomson, J. M. Casas, J. M. Arbelaez, and J. Caicedo. Real-time health monitoring of civil infrastructure systems in colombia. *Proceedings of the 6th Annual International Symposium on NDE for Health Monitoring and Diagnostics*, March 4-8 2001.

[13] C. C. Tiu. A new frequency-based side channel attack for embedded systems. *M.S. thesis*, 2005.

[14] A. Wadaa, S. Olariu, L. Wilson, M. Eltoweissy, and K. Jones. On providing anonymity in wireless sensor networks. *ICPADS '04: Proceedings of the Parallel and Distributed Systems, Tenth International Conference on*, 2004.

[15] A. Wood, G. Virone, T. Doan, Q. Cao, L. Selavo, Y. Wu, L. Fang, Z. He, S. Lin, and J. Stankovic. Alarm-net: Wireless sensor networks for assisted-living and residential monitoring. Technical Report CS-2006-11, Department of Computer Science, University of Virginia, 2006.