Traffic Analysis Attacks and Countermeasures in Wireless Body Area Sensor Networks

Levente Buttyan and Tamas Holczer
Laboratory of Cryptography and Systems Security (CrySyS)
Budapest University of Technology and Economics
Email: {buttyan, holczer}@crysys.hu

***Abstract*—In this paper, we study the problem of traffic analysis attacks in wireless body area sensor networks. When these networks are used in health-care for remote patient monitoring, traffic analysis can reveal the type of medical sensors mounted on the patient, and this information may be used to infer the patient's health problems. We show that simple signal processing methods can be used effectively for performing traffic analysis attacks and identifying the sensor types in a rather weak adversary model. We then investigate possible traffic obfuscation mechanisms aiming at hiding the regular patterns in the observable wireless traffic. Among the investigated countermeasures, traffic shaping, a mechanism that introduces carefully chosen delays for message transmissions, appears to be the best choice, as it achieves close to optimal protection and incurs no overhead.**

I. Introduction

Wireless Body Area Sensor Networks (BASNs) consist of wearable sensors with wireless communication capabilities. Their typical application area is the health-care domain, where they can be used for remote patient monitoring. In this context, the BASN nodes are equipped with different types of medical sensors capable for collecting, for instance, ECG signals, temperature readings, pulse counts, and blood pressure measurements. The BASN nodes may perform some pre-processing of the collected data, however, due to their resource constraints, they cannot accomplish complex processing tasks, neither they can store large amount of data. Therefore, the BASN nodes off-load, using their wireless communication capabilities, their collected data to some gateway node, which is assumed to be a more powerful device such as a smart phone or PDA. The gateway can host various medical applications that use the sensor data collected from the BASN nodes. These applications may interact with the patient himself via well-designed graphical user interfaces, and in addition, they may also provide controlled access to the collected sensor data for remote parties such as a physician.

Wireless communication in case of BASNs typically means radio communication, which is known to be vulnerable to eavesdropping attacks. While message contents can be effectively protected against disclosure by encryption, traffic analysis attacks are still possible even on encrypted traffic. Moreover, in case of remote patient monitoring applications, traffic analysis is a serious threat, because traffic patterns may reveal the type of sensors mounted on the patient's body, which in turn may allow for inferring information about the possible health problems of the patient. Identification of sensor types is made possible by the fact that medical sensors typically perform measurements in a periodic manner. As the BASN nodes have limited storage capabilities, these periodic measurements are typically transferred to the gateway device immediately, resulting in a periodic communication pattern between the BASN nodes and the gateway. A key observation is that different types of sensors perform measurements with different frequencies; for instance, ECG sampling rate can be in the range of 100 to 1000 Hz [1], while temperature sampling usually has a much lower rate of 1 measurement in every 5 minutes. Consequently, the types of sensors on the patient's body can be identified by identifying the different frequency components in the observable traffic. This can potentially work even if the messages are fully encrypted, including their addressing information.

In this paper, our goal is to investigate the feasibility of traffic analysis attacks on BASN traffic, and to propose countermeasures that alleviate the traffic analysis problem. For this, we first introduce a system model and an attacker model in Section II that determine the conditions and the framework of our study. Then, in Section III, we show that simple signal processing techniques provide effective traffic analysis tools that allow identification of different sensors even by a very weak adversary and even if there are perturbations in the regularity of the transmissions of the nodes. In Section IV, we propose countermeasures against traffic analysis and evaluate their effectiveness by means of simulations. Finally, we conclude the paper in Section VI.

II. System and attacker models

The system that we study consists of some BASN nodes and a gateway that use single-hop wireless communication to send messages to each other. The majority of the messages are generated by the BASN nodes that transmit sensor readings to the gateway for storage and further processing. The traffic generated by a BASN node is largely periodic, however, we allow some perturbations in regularity in the form of slight shifts in exact message transmission times (which can represent delays and irregularities due to processing and networking issues) and in the form of a certain level of burstiness (which models the transmission of the result of a measurement campaign where multiple readings of the same sensor are collected and transmitted in multiple messages in a short period of time).

978-1-4673-1239-4/12/$31.00

More specifically, we assume that time is slotted, there are multiple sources, and each source generates traffic independently with the following parameters:

- **Frequency:** This parameter is a constant that determines the frequency of the activity of the given source. Note that activity does not necessarily mean a single message transmission, but it can also mean the transmission of multiple messages in a burst. In the latter case, the frequency parameter determines the frequency of the bursty periods.
- **Deviation:** This parameter is a random variable and it determines a random shift of the exact transmission time (or beginning of the burst) with respect to the regular time determined by the frequency parameter.
- **Burstiness:** This parameter determines the level of burstiness of the traffic generated by the source. We allow for three types of burstiness in our model: no burst, fix burst, and Markov burst. No burst means that the source always transmits a single message when it is active. Fix burst means that a constant number of messages are transmitted when the source is active. Finally, in case of the Markov burst model, the amount of messages in the bursty period is a random value determined by a geometric distribution with some parameter p. In other words, after sending a message, the source continues transmission in the next time slot with some probability p, and stops sending more messages (and terminates the burst) with probability $1-p$.

We assume that the messages are encrypted, including their addressing information[1], and they have a uniform size (message padding can be used to achieve this). In addition, sources of transmissions cannot be distinguished by the attacker based on location information (all BASN nodes are mounted on the body of the patient, and the gateway is carried on the patient too, such that all devices appear to be too close to each other for the attacker), and radio fingerprinting is beyond the capabilities of the attacker (it is not yet mature and reliable enough as a technology). Thus, the attacker cannot distinguish different sources easily by their unique addresses, by their physical characteristics, or by any special properties of the messages themselves (i.e., all messages look a random string of bits with uniform length). These assumptions are needed, since otherwise the attacker could trivially identify different sources and distinguish their transmissions. Thus, in our model, the attacker can only observe the fact of a wireless transmission, and his knowledge, therefore, can be represented by a time series $X_1, X_2, \ldots$, where X_i is a random variable representing the time of the i-th transmission. Note that this is a very weak attacker model in the sense that the capabilities of the attacker are kept at the minimum.

Finally, we assume that the attacker observes the wireless channel for some extended period of time, and his goal is to identify different sources and the characteristics of their traffic, most notably, the frequency of their activities, as this information can reveal the type of the source.

[1]Note that BASN nodes only communicate with the gateway, so each BASN node has a single key that it can use to decrypt messages and determine after decryption if the message was destined to it or not. Furthermore, by appropriate design, this decision can be made without requiring to decrypt the entire message (e.g., the message header can be encrypted separately). The gateway needs to try multiple keys for decrypting incoming messages, but it is a more powerful device. In any case, a careful design of the header encryption scheme is required to prevent Denial-of-Service type attacks where an attacker floods the network with randomly generated messages and forces the nodes to spend precious energy resources to decrypt and verify them.

III. Traffic analysis attacks

As the attacker essentially wants to identify the different periodic components of the observed aggregate traffic, it is natural to interpret the time series representing the transmission events as a discrete signal, and use standard signal processing techniques to transform and analyze it in the spectral domain in order to identify the strongest frequency components. In this section, we introduce the standard Discrete Fourier Transform (DST) and the Welch Averaged Periodogram (WAP) [2] for this purpose, and we show that WAP can be used effectively for traffic analysis. Our selection of these tools has been inspired by [3] where these methods have already been used successfully for traffic analysis in a WiFi based multi-hop ad hoc network.

- **Discrete Fourier Transform (DFT):** Many spectral processing techniques use the standard Discrete Fourier Transform (DFT) to compute the spectrum of a signal in the time domain. The DFT of a uniformly sampled signal $x(n)$ with M samples provides an M-point discrete spectrum $X_M(k)$, where

$$X_M(k) \;=\; \sum_{n=0}^{M-1} x(n) e^{-j\frac{2\pi kn}{M}} \;=\; DFT[x(n)] \qquad (1)$$

  The resulting spectrum $X_M(k)$ is a vector of complex numbers. This spectrum can be efficiently computed by the Fast Fourier Transform (FFT). The peak values in the spectrum correspond to frequencies of message transmission events. Thus, examination of the DFT spectrum can provide a visualization of flows in the form of characteristic peaks.
  However, it is often the case that the spectral content contains many harmonically related peaks. In addition, when the signal is noisy, conventional DFT processing does not provide a good unbiased estimate of the signal power spectrum. A better result in this case can be obtained with the signal periodogram which utilizes averaging in order to reduce the influence of noise.
- **Welch Averaged Periodogram (WAP):** The Welch Averaged Periodogram (WAP) uses windowing to account for the aperiodic nature of the signal. The periodogram $P_x(k)$ of a signal $x(n)$ is generated by averaging the power of K separate spectra $X_L^{(r)}(k)$ computed over K different segments of the data, each of length L:

$$P_x(k) = \frac{1}{KU} \sum_{r=0}^{K-1} |X_L^{(r)}(k)|^2 \qquad (2)$$

Figure 1	Sources	Frequencies	Deviations	Burstiness
(a)	2	$1/23$	2	no burst
		$1/29$	3	no burst
(b)	2	$1/23$	0	Markov $p = 0.4$
		$1/29$	0	Markov $p = 0.4$
(c)	2	$1/23$	2	fix 3
		$1/29$	3	Markov $p = 0.4$
(d)	3	$1/23$	2	Markov $p = 0.4$
		$1/29$	3	Markov $p = 0.4$
		$1/37$	3	Markov $p = 0.3$

TABLE I: Summary of parameter settings for the performance evaluation of DFT and WAP based traffic analysis

where

$$X_L^{(r)}(k) = DFT[\omega(n)x_r(n)] \quad (3)$$

$$U = \frac{1}{L}\sum_{n=0}^{L-1}\omega^2(n) \quad (4)$$

where $x_r(n)$ is the r-th windowed segment of $x(n)$, $\omega(n)$ is a windowing function used to reduce artifacts caused by the abrupt changes at the endpoints of the window, and U is the normalized window power.

We investigated the performance of DFT and WAP by means of simulations. The simulation environment was OMNeT++ (version 4.2.1) with the MiXiM framework (version 2.2.1, 802.15.4, ZigBee protocol stack). We generated the time series of transmission events according to the model described in Section II with different number of sources and different combinations for the parameters of frequency, deviation, and burstiness. Then, we computed the spectrum of the signal obtained from the time series of transmission events by DFT and by WAP. The results are shown in Figure 1. The parameter settings of the figures are summarized in Table I.

As one can see from the figures, both DFT and WAP identify the transmission frequency of the sources despite of the noise caused by the deviation and burstiness parameters, as well as by the characteristics of the wireless channel and the ZigBee protocol used in the simulations. However, WAP better eliminates the harmonic components in most of the cases. Hence, we can conclude that traffic analysis based on WAP works in a robust way, and allows the attacker to identify the sources and the characteristics of their traffic patterns with high confidence.

IV. Countermeasures

In order to make traffic analysis harder for the attacker, the traffic patterns must be obfuscated. This can be achieved by adding some dummy traffic to the original traffic or by altering the timing of transmissions such that they do not follow a regular pattern.

An obvious disadvantage of using dummy traffic to obfuscate the original traffic patterns is that this approach introduces some communication (and computing) overhead, which ultimately results in an increased energy consumption by the BASN nodes. This means that using dummy traffic reduces the lifetime of the BASN nodes, which is undesirable for practical reasons. Note, however, that the energy consumption overhead of the BASN nodes can be reduced, if we let the gateway generate all the dummy messages. This is possible, because in our model, the attacker cannot distinguish between messages generated by the BASN nodes and messages generated by the gateway (they all appear to be random bit strings of some uniform length). In addition, the gateway has more resources and its battery can be recharged easier. But even in this case, the BASN nodes still need to receive and decrypt at least the header of each message, including dummy messages, in order to determine who is the intended destination of the message. Reception and computation also consume energy, therefore, even if the dummy messages are generated by the gateway, the energy consumption of the BASN nodes will be increased.

The advantage of altering the timing of transmissions as a traffic pattern obfuscation mechanism for BASNs is that this approach has no overhead. On the other hand, it introduces some delivery delay for every message and requires some memory on the BASN nodes to store sensor readings until they are transmitted to the gateway. Thus, this approach requires some trade-off between the maximum delay introduced for messages and the level of protection achieved against traffic analysis attacks.

In the following, we introduce and investigate three traffic pattern obfuscation mechanisms:

- **Dummy noise:** This is a dummy traffic generation mechanism, where the gateway injects dummy messages in a completely random manner, hence, creating noise that covers the real traffic.
- **Dummy source:** This is also a dummy traffic generation mechanism, where the gateway simulates the transmissions of a fake sensor, hence, making the attacker believe that the given type of sensor is mounted on the patient, while in reality, it is not. More specifically, if the attacker knows that this mechanism is used for dummy traffic generation, then he cannot be sure that any identified sensor type is actually on the patient or its presence is just simulated by the gateway.
- **Traffic shaping:** This mechanism assumes that the delivery of sensor readings can be delayed up to a maximum delay D, and there are some limited memory on the BASN nodes to store sensor data until they are transmitted to the gateway. When a new sensor reading is obtained, the BASN node schedules its transmission such that the sensor data does not suffer a longer delay than D and the observable inter-transmission times have a close to uniform distribution. This is achieved by continuously tracking (and updating in the memory) the empirical distribution, i.e., the histogram of some granularity, of the inter-transmission times on the wireless interface of the BASN node, and scheduling a new transmission such that the resulting new inter-transmission time falls in the smallest size bin of the histogram.

We evaluated the performance of these countermeasures by means of simulations using OMNeT++ (version 4.2.1) with the

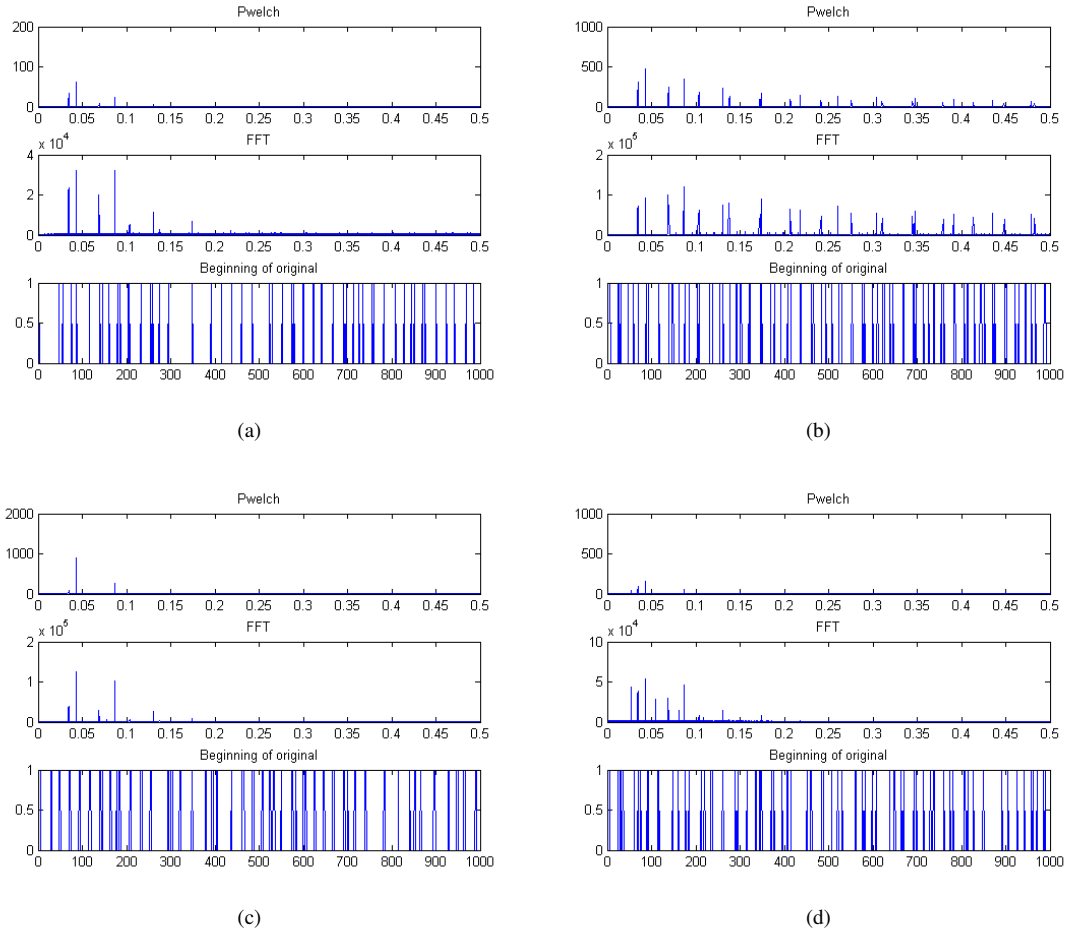

Fig. 1: Traffic analysis with DFT and WAP. Parameter settings are summarized in Table I.

MiXiM framework (version 2.2.1, 802.15.4, ZigBee protocol stack) as the simulation environment. For the purpose of the evaluation, we gave advantage to the attacker, and we used a single real source that generated completely regular traffic with frequency equal to $1/23$. We applied the traffic pattern obfuscation mechanisms described above, and computed the DFT and WAP spectra of the signal representing the time series of transmission events observed by the attacker. The results are shown in Figure 2.

Figure 2(a) shows the spectra of the traffic of the real source without using any countermeasures. The peaks at frequency $1/23$ and its multiples in both the DFT and WAP spectra are well identifiable. Figure 2(b) shows the spectra when the gateway generates dummy transmissions randomly. We can observe the noise in the spectra, although the DFT spectrum preserves the peak at frequency $1/23$, and hence, this method does not properly covers the original traffic pattern. Figure 2(c) shows the spectra when the gateway simulates a fake sensor that generates transmissions with frequency equal to $1/29$. Clearly, for the attacker both frequencies ($1/23$, $1/29$) appear in the spectra just as if there were two sensors transmitting. Finally, Figure 2(d) shows the situation when the BASN node obfuscates its traffic pattern by traffic shaping, using $D = 100$ and 100 bins in the histogram. It can be clearly seen that neither the DFT nor the WAP spectrum preserves any useful information for the attacker.

For the purpose of some quantitative comparison of the different approaches, we define the *entropy of the WAP spectrum*, denoted by H, as a metric to measure the information content in the spectrum of the traffic signal. For this, we compute the WAP spectrum of a given length (in our experiments we used length $L = 100$, which resulted in an $L/2 + 1 = 51$ point periodogram), we omit the component at frequency equal to 0, we normalize the remaining values and interpret them as elements of a discrete probability distribution, and finally, we compute the entropy of this distribution. The entropy values that we obtained for the original traffic signal and the obfuscated signals are shown in Table II, together with the theoretical maximum of the entropy (which was $\log 50 = 5.64$ in our case). Clearly, the closer the entropy of a WAP spectrum is to the theoretical maximum, the larger is its information content, and larger information content of the WAP spectrum means higher uncertainty for the attacker.

As we can see from the table, adding dummy noise and

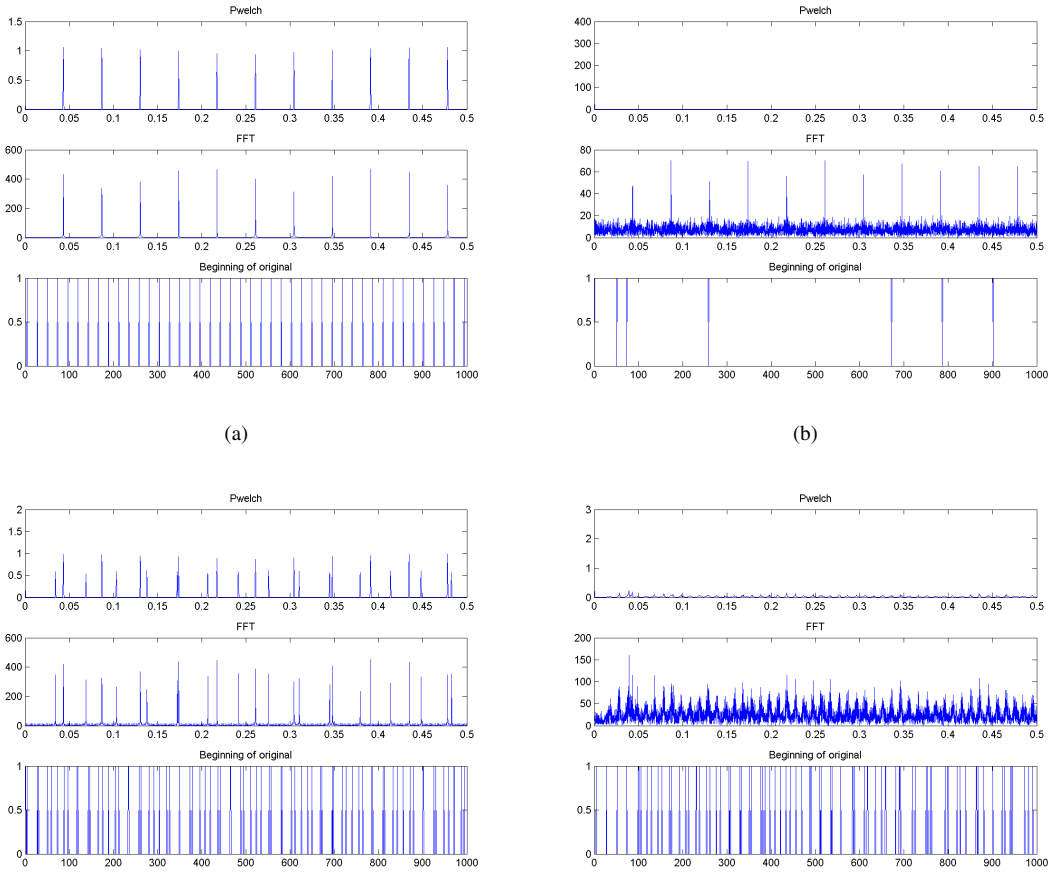

Fig. 2: Traffic analysis of the obfuscated traffic with DFT and WAP.

Obfuscation	Entropy
None	0.0264
Dummy noise	5.5270
Dummy source	1.6078
Traffic shaping	5.2762
Theoretical maximum	5.6439

TABLE II: Entropy values obtained for the original traffic signal and the obfuscated signals. The table also shows the maximum achievable entropy for the spectrum length we used in our experiment.

using traffic shaping both result in an entropy value close to the theoretical optimum. However, as we said before, the dummy noise mechanisms incurs a considerable message overhead, while traffic shaping has no overhead at all. Therefore, among the three traffic pattern obfuscation techniques, traffic shaping seems to be the most promising in our application.

We investigated the performance of traffic shaping a bit further. More specifically, we measured the entropy of the WAP spectrum of a traffic signal obfuscated with traffic shaping using different values for the maximum delay that we allow for the message transmissions and for different granularities (number of bins) of the histogram used to track the empirical distribution of the inter-transmission times. The results of these measurements are shown in Figure 3, where part (a) shows the measured entropy as a function of the maximum delay, and part (b) shows the entropy as the function of the number of the bins. Recall that the frequency of the original traffic signal was $1/23$, meaning that a message was sent in every 23th time slots. From Figure 3(a), we can see that the entropy of the WAP spectrum increases as the maximum delay increases, and the minimum value of the maximum delay where a sufficiently high level of protection is provided is around twice the period length, in our case, around 50. From Figure 3(b), we can see that when the maximum delay as above this threshold, then the number of bins used in the histogram, does not really influence the performance.

V. Related work

The paper with the strongest relation to our work is [3], in which the authors study traffic analysis attacks in wireless networks and show that signal processing techniques such as DFT and WAP can be very effective in practice. Indeed, the work in [3] inspired us to use DFT and WAP in the special

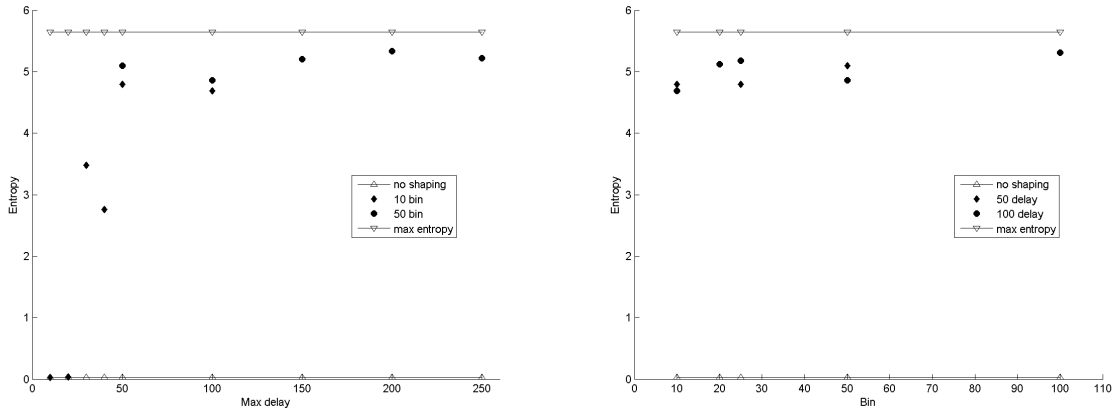

(a) (b)

Fig. 3: Entropy of WAP as a function of (a) the maximum delay and (b) the number of the bins.

case of traffic analysis in BASNs. Besides the difference in the application domains, and hence, in the related wireless technologies, a major difference between [3] and our work is that the former does not propose and investigate possible countermeasures at all. Another related paper is [4], where the authors use statistical and structural content models for separating different traffic types that use different application-layer protocols. However, their approach is based on the flow content, whereas in our model, the attacker cannot access the content of the messages.

Countermeasures against traffic analysis attacks are proposed in [5], [6], [7], [8]. However, [5], [6] are limited to traffic padding with dummy messages. In addition, in [7], [8], the authors consider a problem somewhat different from ours, namely, the identification of the sink in a multi-hop wireless sensor network using the spatial patterns of the traffic. In our work, we focus on traffic shaping that does not require dummy messages, and we are not concerned with the spatial properties of the traffic.

VI. Conclusion

In this paper, we studied traffic analysis attacks in BASNs and we proposed some countermeasures. One of our main conclusions is that simple signal processing methods can be used effectively for performing traffic analysis attacks and identifying the sensor types mounted on the patients. The attacks work in a weak attacker model. Clearly, in a stronger model (e.g., if header encryption is not supported), the traffic analysis attacks become easier; or even trivial. The other main conclusion is that among the possible countermeasures that we investigated, traffic shaping by introducing carefully chosen inter-transmission delays appears to be the best solution, because it has no overhead and it obfuscates traffic patterns effectively. Furthermore, in order to achieve close to optimal protection, the maximum delay used by traffic shaping should be at least twice of the period length of the traffic signal. We believe that in practice, this amount of delay for sensor readings is acceptable as it does not really affect the applications that use the collected sensor data.

Acknowledgment: The work described in this paper is based on results of the CHIRON project[2], which receives research funding from the EU ARTEMIS Joint Undertaking. Apart from this, ARTEMIS has no responsibility for the content of this paper. The information in this document is provided as is and no guarantee or warranty is given that the information is fit for any particular purpose. The user thereof uses the information at its sole risk and liability.

References

[1] R. Fensli, E. Gunnarson, and T. Gundersen, "A wearable ECG-recording system for continuous arrhythmia monitoring in a wireless tele-home-care situation," in *Proceedings of the 18th IEEE Symposium on Computer-Based Medical Systems*, June 2005.

[2] P. D. Welch, "The use of fast fourier transform for estimation of power spectra: A method based on time averaging over short, modified periodograms," *IEEE Transactions on Audio Electroacoustics*, vol. AU-15, pp. 70–73, 1967.

[3] C. Partridge, D. Cousins, A. W. Jackson, R. Krishnan, T. Saxena, and W. Strayer, "Using signal processing to analyze wireless data traffic," in *Proceedings of the 1st ACM workshop on Wireless security*. ACM, 2002, pp. 67–76.

[4] J. Ma, K. Levchenko, C. Kreibich, S. Savage, and G. M. Voelker, "Unexpected means of protocol inference," in *Proceedings of the 6th ACM SIGCOMM conference on Internet measurement*, ser. IMC '06, 2006.

[5] X. Fu, B. Graham, R. Bettati, and W. Zhao, "On effectiveness of link padding for statistical traffic analysis attacks," in *Proceedings of the 23rd International Conference on Distributed Computing Systems*, ser. ICDCS '03, 2003.

[6] R. E. Newman, I. S. Moskowitz, P. Syverson, and A. Serjantov, "Metrics for traffic analysis prevention," in *Proceedings of Privacy Enhancing Technologies workshop*, ser. PET '03, 2003.

[7] J. Deng, R. Han, and S. Mishra, "Countermeasures against traffic analysis attacks in wireless sensor networks," in *Proceedings of the First International Conference on Security and Privacy for Emerging Areas in Communications Networks*, ser. SECURECOMM '05, 2005.

[8] Y. Fan, Y. Jiang, H. Zhu, J. Chen, and X. Shen, "Network coding based privacy preservation against traffic analysis in multi-hop wireless networks," *IEEE Transactions on Wireless Communications*, vol. 10, no. 3, March 2011.

[2] http://www.chiron-project.eu/