

Asynchronous Training in Wireless Sensor Networks

F. Barsi¹, A.A. Bertossi², F. Betti Sorbelli¹, R. Ciotti¹, S. Olariu³, and M.C. Pinotti¹

¹ Department of Computer Science and Mathematics, University of Perugia,
06123 Perugia, Italy, {barsi,pinotti}@unipg.it

² Department of Computer Science, University of Bologna, Mura Anteo Zamboni 7,
40127 Bologna, Italy, bertossi@cs.unibo.it

³ Department of Computer Science, Old Dominion University, Norfolk,
VA 23529-0162, USA, olariu@cs.odu.edu

Abstract. A scalable energy-efficient training protocol is proposed for massively-deployed sensor networks, where sensors are initially anonymous and unaware of their location. The protocol is based on an intuitive coordinate system imposed onto the deployment area which partitions the anonymous sensors into clusters. The protocol is asynchronous, in the sense that the sensors wake up for the first time at random, then alternate between sleep and awake periods both of fixed length, and no explicit synchronization is performed between them and the sink. Theoretical properties are stated under which the training of all the sensors is possible. Moreover, a worst-case analysis as well as an experimental evaluation of the performance is presented, showing that the protocol is lightweight and flexible.

1 Introduction

Ultra-high integration and low-power electronics have enabled the development of miniaturized, low-cost, battery-operated sensor nodes (*sensors*, for short) that integrate signal processing and wireless communications capabilities [1, 14]. Many applications require the aggregation of massively deployed sensors into sophisticated infrastructures, called *sensor networks*. Recently, it has been recognized that it would be beneficial to augment the sensor networks by more powerful entities, called *sinks*. While the sensors are tasked mainly to sense their immediate neighbourhood, the sinks collect, aggregate and fuse the data harvested by the sensors in order to act on the environment in a meaningful way. The typical mode of operation of a sink is to task the sensors in a portion of a disk of radius ρ centered at itself to produce data relevant to the mission at hand. Once this data has been aggregated, the sink has a good idea of what action to take. For instance, Figure 1(a) and Figure 1(b) illustrate, respectively, a disk around a sink and a disk subdivided in portions.

The massive deployment of tiny sensors results in sensors initially unaware of their location. However, many probable applications as well as the sink assignment of tasks to the sensors require that individual sensors have to determine either their exact geographic location or else a coarse-grain approximation thereof.




Fig. 1. (a) *The sensors in a disk centered at a sink.* (b) *The disk subdivided in portions.*

The former task is referred to as *localization* and has been extensively studied in the literature [8, 10]. The latter task, referred to as *training*, has been considered in several recent papers by Olariu *et al.* [3, 11–13]. In particular, they devised some training protocols for sensor networks, which differ on whether or not sensors and sink need some kind of explicit synchronization. Such training protocols have different performance, measured in terms of total time for training, overall sensor awake time, and number of sensor sleep/wake transitions.

The main contribution of this paper is to further study the task of training, assuming the same asynchronous model as that originally defined in [13]. In particular, the model in [13] assumes that the sink and the sensors are asynchronous, in the sense that the sensors wake up for the first time at random and then alternate between sleep and awake periods both of fixed length, while no explicit synchronization is performed between them and the sink. The present paper completes the work of [13], by stating novel theoretical properties of the parameters of the training protocol under which the training of all the sensors in the network is possible. Moreover, improvements of the protocol are presented which are lightweight in terms of both the number of wake/sleep transitions and the overall sensor awake time for training.

The remainder of this paper is organized as follows. Section 2 discusses the wireless sensor network model and introduces the task of training. Training imposes a coordinate system which divides the sensor network area into equiangular wedges and concentric coronas centered at the sink, as first suggested in [12]. Section 3 is the backbone of the entire paper, presenting the theoretical underpinnings of the training protocol, along with its worst-case performance analysis. Section 4 presents an experimental evaluation of the performance, tested on randomly generated instances, showing that the protocol behaves much better in the average case than in the worst case. Finally, Section 5 offers concluding remarks.

2 The network model

In this work a wireless sensor network is assumed that consists of a sink and a set of sensors randomly deployed in its broadcast range as illustrated in Figure 1(a). For simplicity, the sink is centrally placed, although this is not really necessary.

A sensor is a device that possesses three basic capabilities: sensory, computation, and wireless communication, and operates subject to the following fundamental constraints:

- a. Each sensor has a modest non-renewable energy budget and a transmission range of r ;
- b. In order to save energy, each sensor alternates between *sleep* periods and *awake* periods – the sensor sleep-awake cycle is of total length L out of which the sensor is in sleep mode for $L - d$ time and in awake mode for d time;
- c. Each sensor is asynchronous – it wakes up for the first time according to its internal clock and is not engaging in an explicit synchronization protocol with either the sink or the other sensors;
- d. Each sensor has no global information about the network topology, but can hear transmissions from the sink;
- e. Sensors are *anonymous* – to assume the simplest sensor model, sensors do not need individually unique IDs;
- f. Individual sensors must work *unattended* – once deployed it is either infeasible or impractical to devote attention to individual sensors.

The task of training is essential in several applications. One example is clustering where the set of sensors deployed in an area is partitioned into clusters [1, 2, 5]. As a result of training, we impose a coordinate system onto the sensor network in such a way that each sensor belongs to exactly one cluster. The coordinate system involves establishing [12]:

1. *Coronas*: The deployment area is covered by k coronas c_0, c_1, \dots, c_{k-1} determined by k concentric circles, centered at the sink, whose radii are $0 < r_0 < r_1 < \dots < r_{k-1} = \rho$;
2. *Wedges*: The deployment area is ruled into a number of equiangular wedges, centered at the sink, which are established by directional transmission [11].

For the sake of simplicity, in this paper, it is assumed that the corona width is equal to the sensor transmission range r , and hence the (outer) radius r_i of corona c_i is equal to $(i + 1)r$. As illustrated in Figure 1(b), at the end of the training period each sensor has acquired two coordinates: the identity of the corona in which it lies, as well as the identity of the wedge to which it belongs. In particular, a cluster is the locus of all nodes having the same coordinates in the coordinate systems [11].

3 The corona training protocol

The main goal of this section is to present the details of the corona training protocol (the wedge training protocol is similar and will not be discussed), where each individual sensor has to learn the identity of the corona to which it belongs, regardless of the moment when it wakes up for the first time. To see how this is done, it is useful to assume the time ruled into slots. The sensors and the sink use equally long, in phase slots, but they do not necessarily start counting the time from the same slot.

The idea of the protocol, called Flat-, is as follows. Immediately after deployment the sink cyclically repeats a transmission cycle which involves k broadcasts at successively lower power levels. Each broadcast lasts for a slot and transmits a beacon equal to the identity of the outmost corona reached. Precisely, the sink starts out by transmitting the beacon $k - 1$ at the highest power, sufficient to reach the sensors up to the outmost corona c_{k-1} ; next, the sink transmits the beacon $k - 2$ at a power level that can be received up to corona c_{k-2} , but not by the sensors in corona c_{k-1} . For the subsequent $k - 2$ slots, the sink continues to transmit at decreasing power levels until it concludes its transmission cycle with a broadcast that can be received only by the sensors in corona c_0 . In general, at time slot τ , with $\tau \geq 0$, the sink transmits the beacon $k - 1 - |\tau|_k$ with a power level that can reach all the sensors up to corona $c_{k-1-|\tau|_k}$, where $|a|_b$ stands for the non negative remainder of the integer division between a and b (i.e. $|a|_b$ is the same as a modulo b). The sink transmission cycle is repeated for a time sufficient to accomplish the entire corona training protocol.

In order to describe the Flat- protocol for sensors, it is crucial to point out that each sensor is aware of the sink behaviour and of the total number k of coronas. Immediately after deployment, each sensor wakes up at random within the 0-th and the $(k - 1)$ -th time slot and starts listening to the sink for d time slots (that is, its awake period). Then, the sensor goes back to sleep for $L - d$ time slots (that is, its sleep period). Such a sleep/wake transition will be repeated until the sensor will learn the identity of the corona to which belongs, that is, until the sensor will be trained. Each sensor, during the training process, uses a k -bit register R to keep track of the beacons, i.e. corona identities, transmitted by the sink while the sensor is awake. As soon as the sensor hears a sink transmission for the first time, it starts to fill the register R and it is able to learn the sink global time t within the current sink transmission cycle, that is $t = |\tau|_k$. From now on, such a time will regularly increase so that the sensor can derive from t the beacon $|k - 1 - t|_k$ that the sink is transmitting. Then, in each time slot when the sensor is awake, one entry of R can be always set either to 0 or to 1. In fact, if the sensor hears beacon c , then it sets $R_c = 1$, while if the sensor hears nothing, it sets $R_{|k-1-t|_k} = 0$. Note that the awake sensors which belong to corona c , with $c > 0$, are able to receive any transmitted beacon from c up to $k - 1$, whereas they cannot hear the beacons from 0 up to $c - 1$. Hence, if a sensor sets $R_c = 0$ (resp., $R_c = 1$) then it belongs to a corona whose identity is higher than (resp., smaller than or equal to) c . Note that only the sensors in corona 0 can hear beacon 0 and thus they are the only ones which can set $R_0 = 1$. From the above discussion, the following *training condition* holds:

Lemma 1. *A sensor which belongs to corona c , with $c > 0$, is trained as soon as the entries R_c and R_{c-1} of its register R are set to 1 and 0, respectively. A sensor which is in corona 0 is trained as soon as R_0 is set to 1.*

In the following, some conditions on the parameters k, L , and d will be investigated which guarantee that all the sensors are trained, independent of their first wakeup time and from the corona c they belong to. Hereafter, let (a, b)

denote the *greatest common divisor* between a and b . Moreover, if $(a, b) = 1$, let $\left|\frac{1}{a}\right|_b$ be the multiplicative inverse of a modulo b (e.g. see [7]).

Consider a sensor that wakes up for the first time at the global time slot $\tau = x$, while the sink is transmitting the beacon $c_x = |k - 1 - \tau|_k = |k - 1 - x|_k$. The i -th sleep-awake cycle of such a sensor starts at time $x + iL$ while the sink is transmitting the beacon $|k - 1 - x - iL|_k = |c_x - i|L|_k|_k$, with $i \geq 0$. Observe that L and k can be rewritten as $L = gL'$ and $k = gk'$, where $g = (L, k)$. Since $|L|_k = |gL'|_k = g|L'|_{k'}$, one has $|c_x - i|L|_k|_k = |c_x - ig|L'|_{k'}$. Hence, there are exactly k' different coronas which can be transmitted by the sink when the sensor starts its awake period, independent of how long the training process will be. Indeed, since $|c_x - (i + k')|L|_k|_k = |c_x - (i + k')g|L'|_{k'}$ and $|c_x - (ig + k'g)|L'|_{k'}$ are the same corona is transmitted again at the beginning of any two awake periods of the sensor which are k' apart. Moreover, for any two awake periods, say the i -th and the j -th ones, such that $i > j$ and $i - j < k'$, the coronas c_{x+iL} and c_{x+jL} are distinct and differ by a multiple of g . Such overall k' corona identities can be rearranged so that in the new order two consecutive coronas differ exactly by g . Indeed the s -th corona in the new order, that is $|c_x - sg|_k$, corresponds to the first beacon transmitted in the $\left|s\left|\frac{1}{L'}\right|_{k'}\right|_{k'}$ -th awake period, with $0 \leq s \leq k' - 1$. Therefore, after exactly k' sleep-awake cycles, and hence $k'L$ time slots, the sink has performed $\frac{k'L}{k} = \frac{k'L}{gk'} = \frac{L}{g} = L'$ transmission cycles. From now on, the behaviour of the sensor and the sink is cyclically repeated with a period of $k'L$ time slots. In other words, in the k' -th awake period, the sensor and the sink are in the same reciprocal state as in the 0-th one, the only difference being that the sensor has heard the sink at least once. Summarizing:

Lemma 2. *Fixed L, d , and k , there are exactly $k' = \frac{k}{(L, k)}$ different corona identities that can be transmitted by the sink when the sensor starts any awake period. Assuming that the sensor wakes up for the first time at slot x , $0 \leq x \leq k - 1$, then the corona identity transmitted when the sensor starts its i -th awake period is $|c_x - i(L, k)|L'|_{k'}$. Such k' coronas identities can be reindexed as $|c_x - s(L, k)|_k$, for $0 \leq s \leq k' - 1$.*

Thus, since during an awake period of the sensor the sink transmits d distinct beacons, overall the sink transmits no more than $k'd$ different corona identities during the first k' awake periods of the sensor, and such coronas will be repeatedly transmitted again. Recalling that a sensor starts to fill R only after it heard the sink for the first time and observing that all the entries that the sensor can fill are set in further k' sleep-awake cycles, it follows:

Lemma 3. *Fixed L, d , and k , all the entries of R the sensor can fill are set within the first $\frac{2k}{(L, k)}$ sleep-awake cycles, or equivalently, $\frac{2L}{(L, k)}$ sink transmission cycles.*

In other words, if a sensor has not been trained after $\frac{2kL}{(L, k)}$ time slots, it will never be trained, independent of how long the training process will continue.

Theorem 1. *The training condition is satisfied for all the sensors after at most $2k' = 2\frac{k}{(L,k)}$ sleep/wake cycles if and only if $d \geq (L, k)$.*

Proof. For brevity let $g = (L, k)$. By contradiction, suppose that all the sensors have been trained and let $d < g$. By Lemmas 2 and 3, after at most $\frac{2k}{g}$ sleep-awake periods, each sensor has filled at most $k'd$ entries of R . Since $d < g$, each sensor has filled less than k entries of R . Such filled entries depend on the time slot x when the sensor woke up for the first time. Consider now all the sensors that woke up at the same time x . Note that they have filled, although with different configurations, the same positions of R independent of the corona they belong. Let c be one unfilled entry of R . By the hypothesis of massive random deployment, there is at least one sensor that woke up at time x in each corona, and hence at least one sensor in corona c . Clearly, such a sensor will not be trained because the training condition in Lemma 1 will be never satisfied.

Conversely, if $d \geq g$, by Lemma 2, in k' consecutive sleep-awake cycles, the beacons transmitted by the sink in the first slot of such k' cycles are exactly g apart. Since an awake period lasts $d \geq g$ slots, at least g new corona identities are transmitted by the sink during an awake period of the sensor. Hence, after the first k' awake periods, the sensor fills at least g entries of R in each awake period and completely fills R in at most other k' awake periods. Therefore, the sensor is trained in at most $2k'$ consecutive awake periods by Lemma 3. Note that this happens for all the sensors, independent of their first wake-up time x and of the corona c to which they belong.

In the following, some properties of the training protocol are analyzed starting from a couple of particular cases, namely, when $d = (L, k)$ and $d = |L|_k$. Note that, since $d = |L|_k = (L, k)|L'|_{k'}$, Theorem 1 holds in both cases.

First, the maximum number of sleep/wake transitions required to train a sensor is discussed. Precisely, the following lemma specifies when a sensor, that wakes up for the first time at slot x , is awake while the sink is transmitting c .

Lemma 4. *Let c be any corona identity and assume $d = (L, k)$. The sink transmits the beacon c during the $i_{c,x}$ -th awake period of a sensor that wakes up for the first time at slot x , where $i_{c,x} = \left\lfloor \left\lfloor \frac{|c_x - c|_k}{d} \right\rfloor \left| \frac{1}{L'} \right|_{k'} \right\rfloor$, $L' = \frac{L}{d}$, and $k' = \frac{k}{d}$.*

Proof. When the sensor wakes up at time x the sink is transmitting the beacon c_x . Moreover, the beacon values decrease within a sink transmission cycle. Thus, the beacon c will be transmitted, starting from c_x , during the j -th group of d consecutive corona identities such that $j = \left\lfloor \frac{|c_x - c|_k}{d} \right\rfloor$. Such a j -th group of d consecutive corona identities will be transmitted during the $i_{c,x}$ -th sensor awake period in which the sink transmits $\left\lfloor c_x - \left\lfloor \frac{|c_x - c|_k}{d} \right\rfloor d \right\rfloor_k$ as the first beacon. Hence, by Lemma 2, $i_{c,x}$ is derived by solving the equation $\left\lfloor c_x - \left\lfloor \frac{|c_x - c|_k}{d} \right\rfloor d \right\rfloor_k = \left\lfloor c_x - i_{c,x}(L, k) \left| \frac{1}{L'} \right|_{k'} \right\rfloor_k$. Recalling that $d = (L, k)$, the solution of the equation is $i_{c,x} = \left\lfloor \left\lfloor \frac{|c_x - c|_k}{d} \right\rfloor \left| \frac{1}{L'} \right|_{k'} \right\rfloor$.

Lemma 5. *Let c be any corona identity and assume $d = \lfloor L \rfloor_k$. The sink transmits the beacon c during the $i_{c,x}$ -th awake period of a sensor which wakes up for the first time at slot x , where $i_{c,x} = \left\lfloor \frac{|c_x - c|_k}{d} \right\rfloor$.*

Theorem 2. *Let $(L, k) \leq d < \lfloor L \rfloor_k$. A sensor which wakes up for the first time at slot x and belongs to corona c , with $c > 0$, is trained during the i -th awake period where $i = i_{c-1,x}$, if $i_{c,x} \leq i_{c-1,x}$, or $i \leq i_{c,x} + \left\lfloor \frac{1}{L'} \right\rfloor_{k'}$, if $i_{c,x} > i_{c-1,x}$. If $c = 0$, then $i = i_{0,x}$.*

Proof. Consider first the case $d = (L, k)$. If $i_{c,x} \leq i_{c-1,x}$, during the $i_{c,x}$ awake period the sensor hears the beacon c and hence it sets $R_c = 1$. Moreover, during the $i_{c-1,x}$ awake period, the sensor sets $R_{c-1} = 0$ because it does not hear $c-1$ but, having already heard c , it knows what the sink is transmitting. If $i_{c,x} > i_{c-1,x}$, in the worst case the sensor hears for the first time during the $i_{c,x}$ -th awake period and sets $R_c = 1$. Then, the beacon $c-1$ will be transmitted at the i -th awake period such that $|c_x - i(L, k)|_{L'}|_{k'}|_k = |c_x - (j+1)d|_k$, where $j = \left\lfloor \frac{|c_x - c|_k}{d} \right\rfloor$. Solving the above equation, one has $i = (j+1) \left\lfloor \frac{1}{L'} \right\rfloor_{k'}|_{k'} = i_{c,x} + \left\lfloor \frac{1}{L'} \right\rfloor_{k'}$. When $d > (L, k)$, since by Lemma 2 the k' coronas transmitted by the sink when the sensor wakes up do not depend on d , the sensor cannot be trained later than in the case $d = (L, k)$.

Theorem 3. *Let $\lfloor L \rfloor_k \leq d < k$. A sensor which wakes up for the first time at slot x and belongs to corona c , with $c > 0$, is trained during the i -th awake period where $i = i_{c-1,x}$, if $i_{c,x} \leq i_{c-1,x}$, or $i \leq i_{c,x} + 1$, if $i_{c,x} > i_{c-1,x}$. If $c = 0$, then $i = i_{0,x}$.*

In order to analytically evaluate the performance of the Flat- training protocol, let us consider the number ν of sensor sleep/wake transitions, the overall sensor awake time ω , and the total time τ for training. Since a sleep-awake period has length L , and a sensor is awake for d time slots per sleep-awake period, one has $\omega = \nu d$ and $\tau = \nu L$. Thus, the worst case performance for the Flat- protocol can be summarized as follows:

Corollary 1. *Fixed L , d , and k , if $d < (L, k)$ then there are sensors which cannot be trained by the Flat- protocol; otherwise all the sensors are trained, and:*

1. *If $(L, k) \leq d < \lfloor L \rfloor_k$, then $\nu \leq \frac{k}{(L, k)} + \left\lfloor \frac{1}{L'} \right\rfloor_{k'}$, where $k' = \frac{k}{(L, k)}$ and $L' = \frac{L}{(L, k)}$;*
2. *If $\lfloor L \rfloor_k \leq d < k$, then $\nu \leq \left\lfloor \frac{k}{\lfloor L \rfloor_k} \right\rfloor + 1$,*
3. *If $d=k$, then $\nu \leq 2$.*

3.1 Improvements

The Flat- protocol can be improved so as to reduce the number ν of sleep/wake transitions, and hence also the overall sensor awake time as well as the total time for training.

In fact, as soon as a sensor hears the sink transmission for the first time, it learns from the beacon the sink global time modulo the sink transmission cycle. Therefore, it can immediately retrieve backwards the coronas which it did not hear and which were transmitted by the sink during its previous awake periods, setting to 0 the corresponding entries of R . The resulting improved protocol is called *Flat*. Observed that the sensor behaviour is the same as it would have set the entries of R since its first wake up, Lemma 3 and Theorem 1 can be restated as follows:

Lemma 6. *Fixed L, d , and k , all the entries of R the sensor can fill are set within the first $\frac{k}{(L,k)}$ sleep-awake cycles, or equivalently, $\frac{L}{(L,k)}$ sink transmission cycles.*

Theorem 4. *The training condition is satisfied for all the sensors after at most $k' = \frac{k}{(L,k)}$ sleep/wake cycles if and only if $d \geq (L, k)$.*

In other words, after at most $k'L$ time slots the training process is completed. Such a bound is tight in the particular case that $d = (L, k)$, while it can be lowered when $d = |L|_k$. Indeed, Theorems 2 and 3 become:

Theorem 5. *A sensor which wakes up for the first time at slot x and belongs to corona c is trained during the i -th awake period where $i = \max\{i_{c-1,x}, i_{c,x}\}$, if $c > 0$, or $i = i_{0,x}$, if $c = 0$.*

Note that i varies between 0 and $\frac{k}{|L|_k}$ when $d \geq |L|_k$, whereas it varies between 0 and $\frac{k}{(L,k)}$ otherwise. Hence, the worst case performance for the Flat protocol is summarized below:

Corollary 2. *Fixed L, d , and k , if $d < (L, k)$ then there are sensors which cannot be trained by the Flat protocol; otherwise all the sensors are trained, and:*

1. *If $(L, k) \leq d < |L|_k$, then $\nu \leq \frac{k}{(L,k)}$;*
2. *If $|L|_k \leq d < k$, then $\nu \leq \lceil \frac{k}{|L|_k} \rceil$;*
3. *If $d = k$, then $\nu = 1$.*

Note that, when $d = (L, k)$ or $d = |L|_k$, each of the k distinct beacons is transmitted exactly once in the $\lceil \frac{k}{d} \rceil$ awake periods during which each sensor is trained.

A further improvement to the Flat protocol exploits the fact that when a sensor hears a beacon c , it knows that it will also hear all the beacons greater than c , and thus it can immediately set to 1 the entries from R_c up to R_{k-1} . Similarly, when a sensor sets an entry R_c to 0, it knows that it cannot hear any beacon smaller than c , and thus it can immediately set to 0 the entries from R_{c-1} down to R_0 , too. In contrast to the previous protocols, the sensor now fills entries of R relative to beacons not yet transmitted during its awake periods. Therefore, it can look ahead to decide whether it is worthy or not to wake up in the next awake period. If the d entries of R that will be transmitted by the

sink in the next awake period have already been filled, then the sensor can skip its next awake period, thus saving energy. The sensor repeats the look ahead process above until at least one unfilled entry is detected among the d entries corresponding to a future awake period. The resulting protocol is called *Flat+*. Clearly, the number ν of sleep/wake transitions of *Flat+* cannot be larger than that of *Flat*. Moreover, when $d = |L|_k$ or $d = (L, k)$, one can find bad instances where ν , in the worst case, is the same for both *Flat+* and *Flat*. For example, when $d = |L|_k$, a sensor which belongs to corona c and wakes up when the sink transmits $c_x = c - 1$ requires $\lceil \frac{k}{d} \rceil$ transitions to be trained by both protocols. However, as it will be experimentally checked in the following section, the average behaviour of *Flat+* is much better than that of *Flat*.

4 Experimental tests

In this section, the worst and average performance of the *Flat-*, *Flat*, and *Flat+* protocols are experimentally tested. In the simulation, the number k of coronas is fixed to 64, and each corona has a unit width. There are $N = 10000$ sensors uniformly distributed within a circle, centered at the sink, having radius $\rho = k$. Precisely, the polar coordinates of each sensor are generated choosing at random two real numbers. The first one, uniformly distributed between 0 and k , represents the radial coordinate of the sensor, that is, its distance from the sink. The second number, uniformly distributed between 0 and 2π , represents the angular coordinate of the sensor, that is, the positive angle required to reach the sensor from the polar axis. The length L of the sensor sleep-awake cycle assumes the values 104 and 168. Finally, in all the experiments, the sensor awake period d is an integer that varies, with a step of 4, between the greatest common divisor $(L, k) = 8$ and $k = 64$, thus including $|L|_k = 40$. The results are reported only when all the sensors can be trained, that is for $d \geq 8$, and are averaged over 3 independent experiments. In the experiments, both the worst and average number of transitions, denoted by ν_{\max} and ν_{avg} , as well as both the worst and average overall sensor awake time, ω_{\max} and ω_{avg} , are evaluated. Such average values are obtained by summing up the values for each single sensor and then dividing by the number of sensors. Moreover, the total time τ , which measures the time required to terminate the whole training process, is evaluated.

Figure 2 shows the number ν_{\max} and ν_{avg} of transitions for the different values of d . According to Corollaries 1 and 2, *Flat-* has $\nu_{\max} = 13$ when $d = 8$, while both *Flat* and *Flat+* have $\nu_{\max} = 8$. Similarly, when $d = 40$, all protocols take $\nu_{\max} = 2$ transitions. Except for the extreme values $d = 8$ and $d = 64$, the greatest percentage of gain for ν_{\max} is achieved when $d = 24$, where both *Flat+* and *Flat* employ forty percent less transitions than *Flat-*. As regard to the average performance, one notes that ν_{avg} is considerable better than ν_{\max} for all three protocols. *Flat* and *Flat-* have almost the same average performances, while *Flat+* always behaves better than them. In particular, its greatest percentage of gain for ν_{avg} is obtained in the range $8 \leq d \leq 20$, where *Flat+* improves about twenty/thirty percent upon *Flat-*.

Figure 3 shows the awake times $\omega_{\max} = \nu_{\max}d$ and $\omega_{\text{avg}} = \nu_{\text{avg}}d$, which measure the overall energy spent by each sensor to be trained. Although the




Fig. 2. Number of transitions when $k = 64$, $L = 104$, and $8 \leq d \leq 64$.




Fig. 3. Overall sensor awake time when $k = 64$, $L = 104$, and $8 \leq d \leq 64$.

number of transitions decreases as d increases, Figure 3 suggests to choose a small value of d from the sensor awake time perspective. The minimum ω_{\max} is achieved by Flat and Flat+ for $d = 8$ and $d = 64$, as expected by Corollaries 1 and 2. However, when $d = 8$, ω_{avg} lowers to about two thirds of ω_{\max} for Flat- and Flat, and to about one third for Flat+. Note that Flat+ has the maximum gain when d is small. Indeed, it can fill the same entries of R just listening to the sink for a single slot or for d slots. Hence, small values of d save the same number of transitions as larger values, but allow sensors to reduce their energy consumption because they stay awake for smaller periods.

Figure 4 exhibits the total time τ required to accomplish the entire training task, for both $L = 104$ and $L = 168$. Since $|168|_{64} = |104|_{64} = 40$, by Lemma 2, each protocol maintains the same behaviour with respect to the number of transitions. Thus, the plots for $L = 168$ of ν_{\max} and ν_{avg} , and hence of ω_{\max} and




Fig. 4. Total time for training when $k = 64$, $L = 104$ or $L = 168$, and $8 \leq d \leq 64$.

ω_{avg} , are exactly the same as those shown in Figures 2 and 3. Recalling that $\tau = \nu_{\text{max}}L$, the total time for $L = 168$ scales by a constant $\frac{168}{104}$, as depicted in Figure 4. In general, all values of L such that $|L|_k$ is the same present the properties above, namely, ν and ω are identical, while τ scales. Therefore, the minimum total time τ is achieved for the smallest value of L . However, larger values of L could be also selected in order to increase the longevity of the wireless sensor network. Fixed d , a longer L results in a longer life as the life of a sensor is measured in terms of the overall number of sleep-awake cycles until its energy is exhausted.

5 Concluding remarks

In this work a protocol has been proposed which employs the asynchronous model originally presented in [13] and is lightweight in terms of the number of sleep/wake transitions and overall sensor awake time for training. Among the protocol variants, Flat- is the simplest one from the computational viewpoint because each sensor performs $O(1)$ operations per time slot. In contrast, Flat+ has the best performance for small values of d , but it cannot be used if the sensor is not allowed to skip one or more awake periods.

The results presented in this paper show that the protocol is flexible, in the sense that its parameters can be properly tuned. For instance, fixed the number k of coronas, one can decide the optimal values of d and L so as to minimize the number of sleep/wake transitions and/or the overall awake time per sensor. Conversely, one can fix the desired number of sleep/wake transitions, and then select suitable values of d and L .

However, several questions still remain open. First of all, it would be interesting to provide the analytical average behaviour of the protocol. In addition, a good idea for further work should be that of comparing the performance of the protocol proposed in the present paper with that devised in [3]. Indeed, the synchronous training protocol of [3] presents an irregular toggling between sleep

and wake periods, so as to optimize the overall time for training, but it consumes energy in the explicit synchronization between the sensors and the sink to handle such irregular sleep/wake toggling. In contrast, the protocol proposed in Section 3 may force sensors to be awake for a longer time but avoids irregular toggling because sensors alternate between awake and sleep periods both of fixed length. Moreover, in this paper, a boolean (i.e. on/off) transmission model was assumed. That is, for each sink transmission range ir , all the sensors within a disk of radius ir around the sink hear the transmission, while all the sensors out of such a disk do not. Unfortunately, fading and shadowing impact on the connectivity of the network making such an on/off assumption almost impossible in practice [6]. Thus, the impact of the pseudo-ring corona in a real scenario is a very interesting aspect to be further investigated.

References

1. I. F. Akyildiz, W. Su, Y. Sankarasubramanian, and E. Cayirci. Wireless sensor networks: a survey. *Computer Networks*, 38(4):393–422, 2002.
2. S. Bandyopadhyay and E. Coyle. An efficient hierarchical clustering algorithm for wireless sensor networks. *Proc. IEEE INFOCOM 2003*, San Francisco, CA, April 2003.
3. A. A. Bertossi, S. Olariu, and M.C. Pinotti. Efficient training of sensor networks. *Proc. 2nd International Workshop on Algorithmic Aspects of Wireless Sensor Networks (AlgoSensors 2006)*. Lecture Notes in Computer Science 4240, 1–12, 2006.
4. D. Culler, D. Estrin, and M. Srivastava. Overview of sensor networks. *IEEE Computer*, 37(8):41–49, 2004.
5. S. Ghiasi, A. Srivastava, X. Yang, and M. Sarrafzadeh. Optimal energy-aware clustering in sensor networks. *Sensors*, 2:258–269, 2002.
6. J.M. Gorce, R. Zhang, and H. Parvery. Impact of radio link unreliability on the connectivity of wireless sensor networks. *EURASIP Journal on Wireless Communications and Networking*, 2007.
7. H. Griffin, *Elementary Theory of Numbers*, McGraw Hill, New York, 1954.
8. K. Langendoen and N. Reijers. Distributed localization algorithm. In *Embedded Systems Handbook*, R. Zurawski (Editor), CRC Press, Boca Raton, FL, 2004.
9. J. J. Lee, B. Krishnamachari, and C. C. Jay Kuo. Impact of heterogeneous deployment on lifetime sensing coverage in sensor networks. *Proc. IEEE SECON*, 2004.
10. D. Nicosescu. Positioning in ad-hoc sensor networks. *IEEE Network*, 18(4):24–29, 2004.
11. S. Olariu, A. Waada, L. Wilson, and M. Eltoweissy. Wireless sensor networks leveraging the virtual infrastructure. *IEEE Network*, 18(4):51–56, 2004.
12. A. Waada, S. Olariu, L. Wilson, M. Eltoweissy, and K. Jones. Training a wireless sensor network. *Mobile Networks and Applications*, 10(1):151–168, 2005.
13. Q. Xu, R. Ishak, S. Olariu, and S. Salleh. On asynchronous training in sensor networks. *Proc. 3rd Intl. Conf. on Advances in Mobile Multimedia*, K.Lumpur, September 2005.
14. V.V. Zhirnov and D.J.C. Herr. New frontiers: self-assembly and nano-electronics. *IEEE Computer*, 34(1):34–43, 2001.