Chapter 12
Probabilistic Data Propagation in Wireless Sensor Networks

Sotiris Nikoletseas and Paul G. Spirakis

Abstract We study the problem of data propagation in distributed wireless sensor networks. We present two characteristic methods for data propagation: the first one performs a local, greedy optimization to minimize the number of data transmissions needed, while the second creates probabilistically optimized redundant data transmissions to trade off energy efficiency with fault tolerance. Both approaches make use of randomization at both the algorithmic design and analysis; this demonstrates the suitability for distributed sensor network algorithms of probabilistic techniques, because of their simplicity, locality, efficiency, and load-balancing features.

12.1 Introduction

12.1.1 A Brief Overview of Wireless Sensor Networks

Recent dramatic developments in micro-electro-mechanical (MEMS) systems, wireless communications, and digital electronics have already led to the development of small in size, low-power, low-cost sensor devices. Such extremely small devices integrate sensing, data processing, and wireless communication capabilities. Also, they are equipped with a small but effective operating system and are able to switch between "sleeping" and "awake" modes to save energy. Pioneering groups (like the "Smart Dust" Project at Berkeley, the "Wireless Integrated Network Sensors" Project at UCLA, and the "Ultra low Wireless Sensor" Project at MIT) pursue diverse important goals, like a total volume of a few cubic millimeters and extremely low energy consumption, by using alternative technologies, based on radio frequency (RF) or optical (laser) transmission.

Examining each such device individually might appear to have small utility, however, the effective *distributed co-ordination* of large numbers of such devices can lead to the efficient accomplishment of large sensing tasks. Large numbers of sensor nodes can be deployed (usually in an ad hoc manner) in areas of interest

S. Nikoletseas (✉)
Research Academic Computer Technology Institute (RACTI) and Department of Computer Engineering and Informatics, University of Patras, Patras, Greece
e-mail: nikole@cti.gr

S. Nikoletseas, J.D.P. Rolim (eds.), *Theoretical Aspects of Distributed Computing in Sensor Networks,* Monographs in Theoretical Computer Science. An EATCS Series,
DOI 10.1007/978-3-642-14849-1_12,

(such as inaccessible terrains, disaster locations, and ambient intelligence settings) and use *self-organization and collaborative methods* to spontaneously form a sensor network.

Their wide range of applications is based on the possible use of various sensor types (i.e., thermal, visual, seismic, acoustic, radar, magnetic) in order to monitor a wide variety of conditions (e.g., temperature, object presence and motion, humidity, pressure, noise levels). Thus, sensor networks can be used for continuous sensing, reactive event detection, location sensing as well as micro-sensing. Hence, sensor networks have important applications, including (a) environmental applications (such as fire detection, flood detection, precision agriculture), (b) health applications (like telemonitoring of human physiological data), (c) security applications (nuclear, biological, and chemical attack detection), and (d) home applications (e.g., smart environments and home automation). For an excellent survey of wireless sensor networks see [1] and also [7, 14].

We note that a single, universal, and technology independent model is missing in the state of the art, so in each protocol we present we precisely state the particular modeling assumptions (weaker or stronger) needed. Also, in Sect. 12.1.3 we make an attempt to define a hierarchy of models and relations between them.

12.1.2 Critical Challenges

The efficient and robust realization of such large, highly dynamic, complex, non-conventional networking environments is *a challenging algorithmic and technological task*. Features including the huge number of sensor devices involved, the severe power, computational and memory limitations, their dense deployment and frequent failures, pose *new design, analysis, and implementation aspects* which are essentially different not only with respect to distributed computing and systems approaches but also to ad hoc (such as mobile or radio) networking techniques.

We emphasize the following characteristic differences between sensor networks and ad hoc networks:

- The number of sensor particles in a sensor network is extremely large compared to that in a typical ad hoc network.
- Sensor networks are typically prone to frequent faults.
- Because of faults as well as energy limitations, sensor nodes may (permanently or temporarily) join or leave the network. This leads to highly dynamic network topology changes.
- The density of deployed devices in sensor networks is much higher than in ad hoc networks.
- The limitations in energy, computational power, and memory are much more severe.

Because of the above rather unique characteristics of sensor networks, efficient and robust distributed protocols and algorithms should exhibit the following critical properties:

12.1.2.1 Scalability

Distributed protocols for sensor networks should be highly scalable, in the sense that they should operate efficiently in extremely large networks composed of huge numbers of nodes. This feature calls for an urgent need to prove by analytical means and also validate (by large scale simulations and experimental testbeds) certain efficiency and robustness (and their trade-offs) guarantees for asymptotic network sizes. A protocol may work well for a few sensors, but its performance (or even correctness) may dramatically deteriorate as the deployment scale increases.

12.1.2.2 Efficiency

Because of the severe energy limitations of sensor networks and also because of their time-critical application scenarios, protocols for sensor networks should be very efficient, with respect to both energy and time.

12.1.2.3 Fault Tolerance

Sensor particles are prone to several types of faults and unavailability and may become inoperative (permanently or temporarily). Various reasons for such faults include physical damage during either the deployment or the operation phase, permanent (or temporary) cease of operation in the case of power exhaustion (or because of energy saving schemes, respectively). The sensor network should be able to continue its proper operation for as long as possible despite the fact that certain nodes in it may fail. Self-stabilization of the protocol is an aspect of fault tolerance which is especially relevant in this network setting.

12.1.3 Models and Relations Between Them

We note that a single, universal, technology- and cost-independent abstract model is missing in the state of the art, so for each protocol and problem we precisely state the particular modeling assumptions (weaker or stronger) needed. Also, we provide the following hierarchy of models and relations between them. Clearly this is just a starting point and a partial contribution to this matter of great importance.

12.1.3.1 Basic Model M_0

A *sensor cloud* (a set of sensors, which may be called grain particles) is spread in a region of interest (for a graphical presentation, see Fig. 12.1). The deployment is random uniform. Two important network parameters (crucially affecting topology and connectivity) are the particle density d (usually measured in numbers of $particles/m^2$) and the maximum transmission range $\mathcal{R}$ of each grain particle. The sensors do not move. A two-dimensional framework is adopted.

There is a single point in the network, which is called the “sink” S, that represents the fixed control center where data should be reported to. The sink is very powerful,

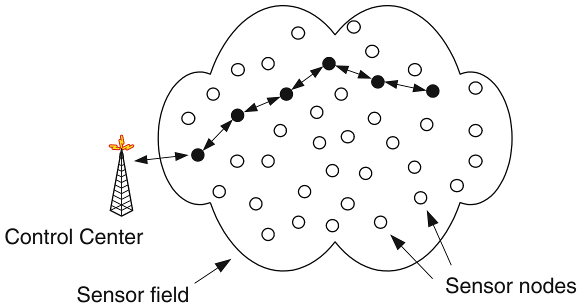

Fig. 12.1 A sensor cloud

in terms of both energy and computing power. Sensors can be aware of the direction (and position) of the sink, as well as of their distance to it. Such information can be, e.g., obtained during a setup phase, by having the (powerful) sink broadcast tiny control messages to the entire network region.

We assume that events, that should be reported to the control center, occur in the network region. Furthermore, we assume that these events are happening at random uniform positions.

As far as energy dissipation is concerned, we basically assume that the energy spent at a sensor when transmitting data is proportional to the square of the transmitting distance. Note, however, that the energy dissipation for receiving is not always negligible with respect to the energy when sending such as in case when transmissions are too short and/or radio electronics energy is high (see [10]).

A variation of this basic model includes multiple and/or mobile sinks. The following "wall" notion generalizes that of the sink and may correspond to multiple (and/or moving) sinks. A *receiving wall* $\mathcal{W}$ is defined to be an infinite (or at least quite large) line in the sensor cloud plane. Any particle transmission within range $\mathcal{R}$ from the wall $\mathcal{W}$ is received by $\mathcal{W}$. $\mathcal{W}$ is assumed to have very strong computing power, is able to collect and analyze received data, and has a constant power supply, and so it has no energy constraints. The wall represents in fact the authorities (the fixed control center) which the realization of a crucial event and collected data should be reported to. Note that a wall of appropriately big (finite) length may suffice. The notion of multiple sinks which may be static or moving has also been studied in [19], where Triantafilloy, Ntarmos, Nikoletseas, and Spirakis introduce "NanoPeer Worlds," merging notions from Peer-to-Peer Computing and Sensor Clouds.

The M_0 model is assumed for the LTP protocol.

12.1.3.2 A Stronger Model M_1

This model is derived by augmenting the model M_0 with additional (stronger) assumptions about the abilities of the sensors and the sensor deployment. In particular, the network topology can be a lattice (or grid) deployment of sensors. This structured placement of grain particles is motivated by certain applications, where it is possible to have a pre-deployed sensor network, where sensors are put (possibly by a human or a robot) in a way that they form a *2-dimensional lattice*. Note indeed that such sensor networks, deployed in a structured way, might be useful, e.g., in precise agriculture or smart buildings, where humans or robots may want to deploy the sensors in a lattice structure to monitor in a rather homogenous and uniform way certain conditions in the spatial area of interest. Also, exact terrain monitoring in military applications may also need some sort of a grid-like shaped sensor network. Note that Akyildiz et al. in a fundamental state of the art survey [1] refer to the pre-deployment possibility. Also, Intanagonwiwat et al. [11] explicitly refers to the lattice case. Moreover, as the authors of [11] state in an extended version of their work [12], they consider, for reasons of "analytic tractability," a square grid topology.

We further assume that each grain particle has the following abilities:

(i) It can estimate the direction of a received transmission (e.g., via the technology of direction-sensing antennae). We note that at least a gross sense of direction may be technologically possible by smart and/or multiple antennas. Regarding the extra cost of the antennas and circuitry, we note that certainly such a cost is introduced but as technology advances this cost may lower.
(ii) It can estimate the distance from a nearby particle that did the transmission (e.g., via estimation of the attenuation of the received signal).
(iii) It knows the direction toward the sink S. This can be implemented during a setup phase, where the (very powerful in energy) sink broadcasts the information about itself to all particles.
(iv) All particles have a common co-ordinates system.

Notice that GPS information is not assumed by this model. Also, the global structure of the network is not assumed known.

This model is used in the PFR protocol.

12.1.4 The Energy Efficiency Challenge in Routing

Since one of the most severe limitations of sensor devices is their limited energy supply (battery), one of the most crucial goals in designing efficient protocols for wireless sensor networks is minimizing the energy consumption in the network. This goal has various aspects, including: (a) minimizing the total energy spent in the network, (b) minimizing the number (or the range) of data transmissions, (c) combining energy efficiency and fault tolerance, by allowing redundant data transmissions which, however, should be optimized to not spend too much energy, (d) maximizing the number of "alive" particles over time, thus prolonging the system's lifetime, and (e) balancing the energy dissipation among the sensors in the network, in order to avoid the early depletion of certain sensors and thus the breakdown of the network (see, e.g., [8]). Even energy aspects related to potentially dangerous for health radiation levels (and their auto balancing) are relevant.

We note that it is very difficult to achieve all the above goals at the same time. There even exist trade-offs between some of the goals above. Furthermore, the importance and priority of each of these goals may depend on the particular application domain. Thus, it is important to have a variety of protocols, each of which may possibly focus at some of the energy efficiency goals above (while still performing well with respect to the rest goals as well).

In this chapter, we focus on online aspects of energy efficiency related to data propagation. In particular, for routing we present two methodologically characteristic energy efficient protocols:

- *The Local Target Protocol (LTP)* that performs a local optimization trying to minimize the number of data transmissions.
- *The Probabilistic Forwarding Protocol (PFR)* that creates redundant data transmissions that are probabilistically optimized, to trade off energy efficiency with fault tolerance.

Because of the complex nature of a sensor network (that integrates various aspects of communication and computing), protocols, algorithmic solutions, and design schemes for all layers of the networking infrastructure are needed. Far from being exhaustive, we mention the need for frequency management solutions at the physical layer and Medium Access Control (MAC) protocols to cope with multi-hop transmissions at the data link layer. Our approach focuses on routing and does not directly consider other important issues like link fading and hidden station interference, i.e., we assume that other network layers are resolving these issues. Also, regarding reliability, an endemic issue to wireless sensor networks, we consider permanent physical failures, i.e., we do not deal with bad values from physical sensing or transient failures. The interested reader may use the excellent survey by Akyildiz et al. [2] for a detailed discussion of design aspects and state of the art of all layers of the networking infrastructure.

12.2 LTP: A Single-Path Data Propagation Protocol

The LTP Protocol was introduced in [3]. The authors assume the basic model M_0 defined in Sect. 12.1.3.

12.2.1 The Protocol

Let $d(p_i,\ p_j)$ be the distance (along the corresponding vertical lines toward $\mathcal{W}$) of particles $p_i,\ p_j$ and $d(p_i, \mathcal{W})$ the (vertical) distance of p_i from $\mathcal{W}$. Let *info*$(\mathcal{E})$ the information about the realization of the crucial event $\mathcal{E}$ to be propagated. Let p the particle sensing the event and starting the execution of the protocol. In this protocol, each particle p' that has received *info*$(\mathcal{E})$ does the following:

- *Search phase*: It uses a periodic low-energy directional broadcast in order to discover a particle nearer to $\mathcal{W}$ than itself (i.e., a particle p'' where $d(p'', \mathcal{W}) < d(p', \mathcal{W})$).
- *Direct transmission phase*: Then, p' sends *info($\mathcal{E}$)* to p''.
- *Backtrack phase*: If consecutive repetitions of the *search phase* fail to discover a particle nearer to $\mathcal{W}$, then p' sends *info($\mathcal{E}$)* to the particle that it originally received the information from.

Note that one can estimate an a priori upper bound on the number of repetitions of the search phase needed, by calculating the probability of success of each search phase, as a function of various parameters (such as density, search angle, transmission range). This bound can be used to decide when to backtrack.

For a graphical representation see Figs. 12.2 and 12.3.

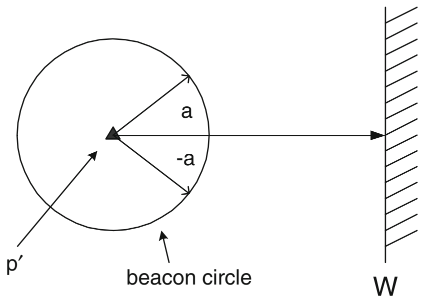

Fig. 12.2 Example of the search phase

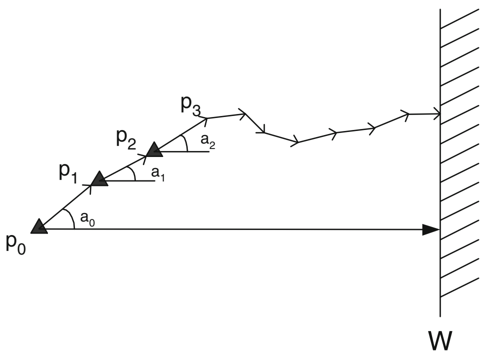

Fig. 12.3 Example of a transmission

12.2.2 Analysis of the Expected Hops Efficiency

We first provide some basic definitions.

Definition 1 Let $h_{\text{opt}}(p, \mathcal{W})$ be the (optimal) number of "hops" (direct, vertical to $\mathcal{W}$ transmissions) needed to reach the wall, in the *ideal* case when particles always exist in pairwise distances $\mathcal{R}$ on the vertical line from p to $\mathcal{W}$. Let Π be *a smart-dust propagation protocol*, using *a transmission path* of length $L(\Pi, p, \mathcal{W})$ to send info about event $\mathcal{E}$ to wall $\mathcal{W}$. Let $h(\Pi, p, \mathcal{W})$ be the actual number of hops (transmissions) taken to reach $\mathcal{W}$. The *"hops" efficiency* of protocol Π is the ratio

$$C_h = \frac{h(\Pi, p, \mathcal{W})}{h_{\text{opt}}(p, \mathcal{W})}$$

Clearly, the number of hops (transmissions) needed characterizes the (order of the) energy consumption and the time needed to propagate the information $\mathcal{E}$ to the wall. Remark that $h_{\text{opt}} = \left\lceil \frac{d(p,\mathcal{W})}{\mathcal{R}} \right\rceil$, where $d(p, \mathcal{W})$ is the (vertical) distance of p from the wall $\mathcal{W}$.

In the case where the protocol Π is randomized, or in the case where the distribution of the particles in the cloud is a random distribution, the number of hops h and the efficiency ratio C_h are random variables and one wishes to study their expected values.

The reason behind these definitions is that when p (or any intermediate particle in the information propagation to $\mathcal{W}$) "looks around" for a particle as near to $\mathcal{W}$ as possible to pass its information about $\mathcal{E}$, it may not get any particle in the perfect direction of the line vertical to $\mathcal{W}$. This difficulty comes mainly from three causes: (a) Due to the initial spreading of particles of the cloud in the area and because particles do not move, there might not be any particle in that direction. (b) Particles of sufficient remaining battery power may not be currently available in the right direction. (c) Particles may be there but temporarily "sleep" (i.e., not listen to transmissions) in order to save battery power.

Note that any given distribution of particles in the sensor cloud may not allow the ideal optimal number of hops to be achieved at all. In fact, the least possible number of hops depends on the input (the positions of the grain particles). We, however, compare the efficiency of protocols to the ideal case. A comparison with the best achievable number of hops in each input case will of course give better efficiency ratios for protocols.

To enable a first step toward a rigorous analysis of routing protocols, we make the following simplifying assumption: *The search phase always finds a* p'' (of sufficiently high battery) in the semicircle of center the particle p' currently possessing the information about the event and radius R, in the direction toward $\mathcal{W}$. Note that this assumption of always finding a particle can be relaxed in the following ways: (a) by repetitions of the search phase until a particle is found. This makes sense if at least one particle exists but was sleeping during the failed searches, (b) by considering, instead of just the semicircle, a cyclic sector defined by circles of radiuses

$\mathcal{R} - \Delta\mathcal{R}$, $\mathcal{R}$ and also take into account the density of the sensor cloud, and (c) if the protocol during a search phase ultimately fails to find a particle toward the wall, it may *backtrack*.

We also assume that the position of p'' is uniform in the arc of angle $2a$ around the direct line from p' vertical to $\mathcal{W}$. Each data transmission (one hop) takes constant time t (so the "hops" and time efficiency of our protocols coincide in this case). It is also assumed that each target selection is stochastically *independent* of the others, in the sense that it is always drawn uniformly randomly in the arc $(-a, a)$.

The above assumptions may not be very realistic in practice, however, they can be relaxed as explained above and in any case allow to perform a first effort toward providing some concrete analytical results.

Lemma 1 *The expected "hops efficiency" of the local target protocol in the a-uniform case is*

$$E(C_h) \simeq \frac{\alpha}{\sin\alpha}$$

for large h_{opt}. *Also*

$$1 \le E(C_h) \le \frac{\pi}{2} \simeq 1.57$$

for $0 \le \alpha \le \frac{\pi}{2}$.

Proof Due to the protocol, a sequence of points is generated, $p_0 = p, p_1, p_2, \ldots, p_{h-1}, p_h$ where p_{h-1} is a particle within $\mathcal{W}$'s range and p_h is beyond the wall. Let α_i be the (positive or negative) angle of p_i with respect to p_{i-1}'s vertical line to $\mathcal{W}$. It is

$$\sum_{i=1}^{h-1} d(p_{i-1}, p_i) \le d(p, \mathcal{W}) \le \sum_{i=1}^{h} d(p_{i-1}, p_i)$$

Since the (vertical) progress toward $\mathcal{W}$ is then $\Delta_i = d(p_{i-1}, p_i) = \mathcal{R}\cos\alpha_i$, we get

$$\sum_{i=1}^{h-1} \cos\alpha_i \le h_{\text{opt}} \le \sum_{i=1}^{h} \cos\alpha_i$$

From Wald's equation for the expectation of a sum of a random number of independent random variables (see [18]), then

$$E(h-1) \cdot E(\cos\alpha_i) \le E(h_{\text{opt}}) = h_{\text{opt}} \le E(h) \cdot E(\cos\alpha_i)$$

Now, $\forall i$, $E(\cos\alpha_i) = \int_{-\alpha}^{\alpha} \cos x \, \frac{1}{2\alpha} dx = \frac{\sin\alpha}{\alpha}$. Thus

$$\frac{\alpha}{\sin\alpha} \leq \frac{E(h)}{h_{\text{opt}}} = E(C_h) \leq \frac{\alpha}{\sin\alpha} + \frac{1}{h_{\text{opt}}}$$

Assuming large values for h_{opt} (i.e., events happening far away from the wall, which is the most non-trivial and interesting case in practice since the detection and propagation difficulty increases with distance) we get the result (since for $0 \leq \alpha \leq \frac{\pi}{2}$ it is $1 \leq \frac{\alpha}{\sin\alpha} \leq \frac{\pi}{2}$) . □

12.2.3 Local Optimization: The Min-Two Uniform Targets Protocol (M2TP)

We can further assume that the search phase always returns *two points* p'', p''' each uniform in $(-\alpha, \alpha)$ and that a modified protocol M2TP selects the best of the two points, with respect to the local (vertical) progress. This is in fact an optimized version of the Local Target Protocol.

In a similar way as in the proof of the previous lemma, we can prove the following result:

Lemma 2 *The expected "hops efficiency" of the "min-two uniform targets" protocol in the a-uniform case is*

$$E(C_h) \simeq \frac{\alpha^2}{2(1-\cos\alpha)}$$

for $0 \leq \alpha \leq \frac{\pi}{2}$ *and for large h.*

Now remark that

$$\lim_{\alpha\to 0} E(C_h) = \lim_{\alpha\to 0} \frac{2\alpha}{2\sin a} = 1$$

and

$$\lim_{\alpha\to\frac{\pi}{2}} E(C_h) = \frac{(\pi/2)^2}{2(1-0)} = \frac{\pi^2}{8} \simeq 1.24$$

Thus, we have the following:

Lemma 3 *The expected "hops" efficiency of the min-two uniform targets protocol is*

$$1 \leq E(C_h) \leq \frac{\pi^2}{8} \simeq 1.24$$

for large h and for $0 \leq \alpha \leq \frac{\pi}{2}$.

Remark that, with respect to the expected hops efficiency of the local target protocol, the min-two uniform targets protocol achieves, because of the one additional search, a relative gain which is $(\pi/2 - \pi^2/8)/(\pi/2) \simeq 21.5\%$. Chatzigiannakis et al. [3] also experimentally investigate the further gain of additional (i.e., $m > 2$) searches.

12.2.4 Tight Upper Bounds to the Hops Distribution of the General Target Protocol

Consider the particle p (which senses the crucial event) at distance x from the wall. Let us assume that when p searches in the sector S defined by angles $(-\alpha, \alpha)$ and radius $\mathcal{R}$, another particle p' is returned *in the sector* with some probability density $f(\vec{p'})d\mathcal{A}$, where $\vec{p'} = (x_{p'}, y_{p'})$ is the position of p' in S and $d\mathcal{A}$ is an infinitesimal area around p'.

Definition 2 (Horizontal progress) Let Δx be the projection of the line segment (p, p') on the line from p vertical to $\mathcal{W}$.

We assume that each search phase returns such a particle, with independent and identical distribution $f()$.

Definition 3 (Probability of significant progress) Let $m > 0$ be the least integer such that $\mathbf{P}\left\{\Delta x > \frac{\mathcal{R}}{m}\right\} \geq p$, where $0 < p < 1$ is a given constant.

Lemma 4 For each continuous density $f()$ on the sector S and for any constant p, there is always an $m > 0$ as above.

Proof Remark that $f()$ defines a density function $\tilde{f}()$ on $(0, \mathcal{R}]$ which is also continuous. Let $\tilde{F}()$ its distribution function. Then we want $1 - \tilde{F}(\mathcal{R}/m) \geq p$, i.e., to find the first m such that $1 - p \geq \tilde{F}(\mathcal{R}/m)$. Such an m always exists since $\tilde{F}$ is continuous in $[0, 1]$. □

Definition 4 Consider the (discrete) stochastic process P in which with probability p the horizontal progress is $\mathcal{R}/m$ and with probability q it is zero, where $q = 1 - p$. Let Q the actual stochastic process of the horizontal progress implied by $f()$.

Lemma 5 $\mathbf{P}_P\{h \leq h_0\} \leq \mathbf{P}_Q\{h \leq h_0\}$

Proof The actual process Q makes always more progress than P. □

Now let $t = \left\lceil \frac{x}{\mathcal{R}/m} \right\rceil = \left\lceil \frac{mx}{\mathcal{R}} \right\rceil$. Consider the integer random variable H such that $\mathbf{P}\{H = i\} = q^i(1 - q)$ for any $i \geq 0$. Then H is geometrically distributed. Let $H_1, \ldots, H_t$ be t random variables, independent and identically distributed according to H. Clearly then

Lemma 6 $\mathbf{P}_{\mathrm{P}}\{\textit{number of hops is } \mathrm{h}\} = \mathbf{P}\{H_1 + \cdots H_{\mathrm{t}} = h\}$

The probability generating function of H is

$$H(s) = \mathbf{P}\{H = 0\} + \mathbf{P}\{H = 1\}s + \cdots + \mathbf{P}\{H = i\}s^i + \cdots$$

i.e.,

$$H(s) = p(1 + qs + q^2s^2 + \cdots + q^i s^i + \cdots) = \frac{p}{1 - qs}$$

But the probability generating function of $\sum_t = H_1 + \cdots + H_t$ is then just $\left(\frac{p}{1-qs}\right)^t$ by the convolution property of generating functions. This is just the generating function of the t-fold convolution of geometric random variables, and it is *exactly* the distribution of the *negative binomial distribution* (see [9], vol. I. p. 253). Thus,

Theorem 1 $\mathbf{P}_P\{\text{the number of hops is } h\} = \binom{-t}{h} p^t (-q)^h = \binom{t+h-1}{h} p^t q^h$

Corollary 1 *For the process P, the mean and variance of the number of hops are*

$$E(h) = \frac{tq}{p} \qquad \mathrm{Var}(h) = \frac{tq}{p^2}$$

Note that the method sketched above finds a distribution that upper bounds the number of hops till the crucial event is reported to the wall. Since for all $f()$ it is $h \geq \frac{x}{\mathcal{R}} = h_{\text{opt}}$ we get that

$$\frac{E_P(h)}{h_{\text{opt}}} \leq \frac{\left\lceil \frac{mx}{\mathcal{R}} \right\rceil \frac{q}{p}}{x/\mathcal{R}} \leq \frac{(m+1)q}{p}$$

Theorem 2 *The above upper bound process P estimates the expected number of hops to the wall with a guaranteed efficiency ratio* $\frac{(m+1)(1-p)}{p}$ *at most.*

Example 1 When for $p = 0.5$, we have $m = 2$ and the efficiency ratio is 3, i.e., the overestimate is three times the optimal number of hops.

12.3 PFR–A Probabilistic Multi-path Forwarding Protocol

The LTP protocol, as shown in the previous section, manages to be very efficient by always selecting exactly one next-hop particle, with respect to some optimization criterion. Thus, it tries to minimize the number of data transmissions. LTP is indeed very successful in the case of dense and robust networks, since in such networks a next-hop particle is very likely to be discovered. In sparse or faulty networks, however, the LTP protocol may behave poorly, because of many backtracks due to frequent failure to find a next-hop particle. To combine energy efficiency and fault tolerance, the Probabilistic Forwarding Protocol (PFR) has been introduced. The trade-offs in the performance of the two protocols implied above are shown and discussed in great detail in [5].

The PFR protocol was introduced in [4]. The authors assume the model M_1 (defined in Sect. 12.1.3).

12.3.1 The Protocol

The PFR protocol is inspired by the probabilistic multi-path design choice for the directed diffusion paradigm mentioned in [11]. The basic idea of the protocol (introduced in [4]) lies in probabilistically favoring transmissions toward the sink within a *thin zone* of particles around the line connecting the particle sensing the event $\mathcal{E}$ and the sink (see Fig. 12.4). Note that transmission along this line is energy optimal. However, it is not always possible to achieve this optimality, basically because certain sensors on this direct line might be inactive, either permanently (because their energy has been exhausted) or temporarily (because these sensors might enter a sleeping mode to save energy). Further reasons include (a) physical damage of sensors, (b) deliberate removal of some of them (possibly by an adversary in military applications), (c) changes in the position of the sensors due to a variety of reasons (weather conditions, human interaction etc.), and (d) physical obstacles blocking communication.

The protocol evolves in two phases.

Phase 1: The "Front" Creation Phase

Initially the protocol builds (by using a limited, in terms of rounds, flooding) a sufficiently large "front" of particles, in order to guarantee the survivability of the data propagation process. During this phase, each particle having received the data to be propagated, deterministically forward it toward the sink. In particular, and for a sufficiently large number of steps $s = 180\sqrt{2}$, each particle broadcasts the information to all its neighbors, toward the sink. Remark that to implement this phase, and in particular to count the number of steps, we use a counter in each message. This counter needs at most $\lceil \log 180\sqrt{2} \rceil$ bits.

Phase 2: The Probabilistic Forwarding Phase

During this phase, each particle P possessing the information under propagation calculates an angle ϕ by calling the subprotocol "ϕ-calculation" (see description below) and broadcasts $info(\mathcal{E})$ to all its neighbors with probability $\mathbf{P}_{\text{fwd}}$ (or it does not propagate any data with probability $1 - \mathbf{P}_{\text{fwd}}$) defined as follows:

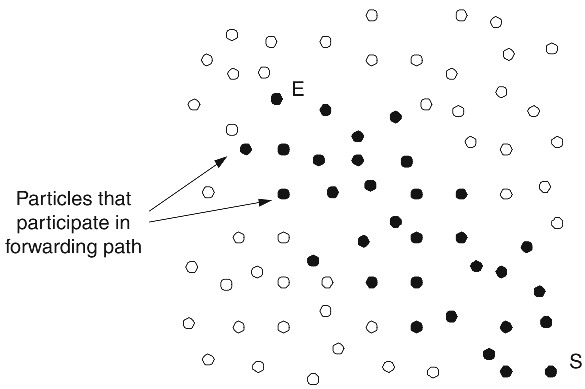

Fig. 12.4 Thin zone of particles

$$\mathbf{P}_{\text{fwd}} = \begin{cases} 1 & \text{if } \phi \geq \phi_{\text{threshold}} \\ \frac{\phi}{\pi} & \text{otherwise} \end{cases}$$

where ϕ is the angle defined by the line EP and the line PS and $\phi_{\text{threshold}} = 134°$ (the selection reasons of this $\phi_{\text{threshold}}$ can be found in [4]).

In both phases, if a particle has already broadcast $info(\mathcal{E})$ and receives it again, it ignores it. Also the PFR protocol is presented for a single event tracing. Thus no multiple paths arise and packet sizes do not increase with time.

Remark that when $\phi = \pi$ then P lies on the line ES and vice versa (and always transmits).

If the density of particles is appropriately high, then for a line ES there is (with high probability) a sequence of points "closely surrounding ES" whose angles ϕ are larger than $\phi_{\text{threshold}}$ and so that successive points are within transmission range. All such points broadcast and thus essentially they follow the line ES (see Fig. 12.4).

12.3.1.1 The ϕ-Calculation Subprotocol

Let P_{prev} the particle that transmitted $info(E)$ to P (see Fig. 12.5).

(1) When P_{prev} broadcasts $info(E)$, it also attaches the info $|EP_{\text{prev}}|$ and the direction $\overrightarrow{P_{\text{prev}}E}$.
(2) P estimates the direction and length of line segment $P_{\text{prev}}P$, as described in the model.
(3) P now computes angle $(\widehat{EP_{\text{prev}}P})$ and computes $|EP|$ and the direction of $\overrightarrow{PE}$ (this will be used in further transmissions from P).
(4) P also computes angle $(\widehat{P_{\text{prev}}PE})$ and by subtracting it from $(\widehat{P_{\text{prev}}PS})$ it finds ϕ.

Notice the following:

(i) The direction and distance from activated sensors to E is inductively propagated (i.e., P becomes P_{prev} in the next phase).
(ii) The protocol needs only messages of length bounded by $\log A$, where A is some measure of the size of the network area, since (because of (i) above) there is no cumulative effect on message lengths.

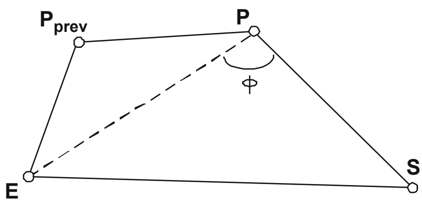

Fig. 12.5 Angle ϕ calculation example

Essentially, the protocol captures the intuitive, deterministic idea "if my distance from ES is small, then send, else do not send." Chatzigiannakis et al. [4] have chosen to enhance this idea by random decisions (above a threshold) to allow some local flooding to happen with small probability and thus to cope with local sensor failures.

12.3.2 Properties of PFR

Any protocol Π solving the data propagation problem must satisfy the following three properties:

- *Correctness*. Π must guarantee that data arrive to the position S, given that the whole network exists and is operational.
- *Robustness*. Π must guarantee that data arrive at enough points in a small interval around S, in cases where part of the network has become inoperative.
- *Efficiency*. If Π activates k particles during its operation then Π should have a small ratio of the number of activated over the total number of particles $r = \frac{k}{N}$. Thus r is an energy efficiency measure of Π.

We show that this is indeed the case for PFR.

Consider a partition of the network area into small squares of a fictitious grid G (see Fig. 12.6). Let the length of the side of each square be l. Let the number of squares be q. The area covered is bounded by ql^2. Assuming that we randomly throw in the area at least $\alpha q \log q = N$ particles (where $\alpha > 0$ a suitable constant), then the probability that a particular square is avoided tends to 0. So with very high probability (tending to 1) all squares get particles.

We condition all the analysis on this event, call it F, of at least one particle in each square.

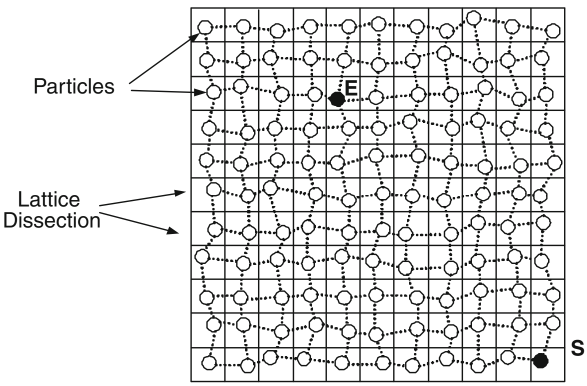

Fig. 12.6 A lattice dissection G

12.3.3 *The Correctness of* PFR

Without loss of generality, we assume each square of the fictitious lattice G to have side length 1.

Lemma 7 PFR *succeeds with probability 1* in sending the information from E to S given the event F.

Proof In the (trivial) case where $|ES| \leq 180\sqrt{2}$, the protocol is clearly correct due to front creation phase.

Let Σ a unit square of G intersecting ES in some way (see Fig. 12.7). Since a particle always exists somewhere in Σ, we will only need to examine the worst case of it being in one of the corners of Σ. Consider vertex A. The line EA is always to the left of AB (since E is at end of ES). The same is true for AS (S is to the right of B').

Let AD the segment from A perpendicular to ES (and D its intersection point) and let yy' be the line from A parallel to ES. Then,

$$\begin{aligned}\phi = (\widehat{EAS}) &= 180^\circ - (\widehat{yAE}) - (\widehat{y'AS})\\ &= 90^\circ - (\widehat{yAE}) + 90^\circ - (\widehat{y'AS})\\ &= (\widehat{yAD}) - (\widehat{yAE}) + (\widehat{DAy'}) - (\widehat{y'AS})\end{aligned}$$

Let $\widehat{\omega_1} = (\widehat{yAD}) - (\widehat{yAE})$ and $\widehat{\omega_2} = (\widehat{DAy'}) - (\widehat{y'AS})$ and, without loss of generality, let $ED < DS$. Then always $\widehat{\omega_1} > 45^\circ$, since it includes half of the 90°-angle of A in the unit square. Also

$$\sin(\widehat{y'AS}) = \sin(\widehat{ASD}) = \frac{AD}{AS} < \frac{AD}{DS}$$

$$\text{but} \quad AD \leq \sqrt{2} \quad (= AB) \quad \text{and} \quad DS \geq \frac{ES}{2} \geq 90\sqrt{2}$$

$$\Longrightarrow \sin(\widehat{y'AS}) \leq \frac{1}{90} \iff (\widehat{y'AS}) < 1^\circ$$

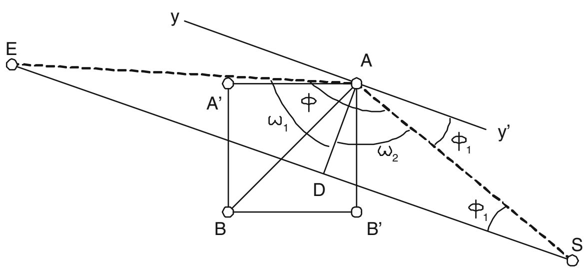

Fig. 12.7 The square Σ

(Note that here we use the elementary fact that when $x < 90°$ then $sinx < \frac{1}{90} \Longleftrightarrow x < 1°$). Then $\widehat{\omega_2} > 89°$, since $(\widehat{DAy'}) = 90°$ by construction. Thus

$$\phi = \widehat{\omega_1} + \widehat{\omega_2} > 45° + 89° = 134°$$

There are two other ways to place an intersecting to ES unit square Σ, whose analysis is similar.

But (a) the initial square (from E) always broadcasts due to the first protocol phase and (b) any intermediate intersecting square will be notified (by induction) and thus will broadcast. Hence, S is always notified, when the whole grid is operational. □

12.3.4 The Energy Efficiency of PFR

Consider the fictitious lattice G of the network area and let the event F hold. We have (at least) one particle inside each square. Now join all "nearby" particles (see Fig. 12.6) of each particle to it, thus by forming a new graph G' which is "lattice shaped" but its elementary "boxes" may not be orthogonal and may have varied length. When G's squares become smaller and smaller, then G' will look like G. Thus, for reasons of analytic tractability, we will assume in the sequel that our particles form a lattice (see Fig. 12.8). We also assume length $l = 1$ in each square, for normalization purposes. Notice, however, that when $l \to 0$ then "$G' \to G$" and thus all our results in this section hold for any random deployment "in the limit." We now prove the following:

Theorem 3 *The energy efficiency of the* PFR *protocol is* $\Theta\left(\left(\frac{n_0}{n}\right)^2\right)$ *where* $n_0 = |ES|$ *and* $n = \sqrt{N}$. *For* $n_0 = |ES| = o(n)$, *this is* $o(1)$.

Proof The analysis of the energy efficiency considers particles that are active but are as far as possible from ES. Thus the approximations we do are suitable for remote particles.

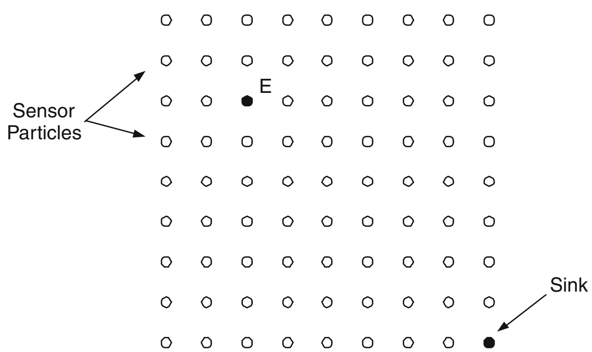

Fig. 12.8 A lattice sensor network

Here we estimate an upper bound on the number of particles in an $n \times n$ (i.e., $N = n \times n$) lattice. If k is this number then $r = \frac{k}{n^2}$ $(0 < r \leq 1)$ is the "energy efficiency ratio" of PFR.

We want r to be less than 1 and as small as possible (clearly, $r = 1$ leads to flooding). Since, by Lemma 7, ES is always "surrounded" by active particles, we will now assume without loss of generality that ES is part of a horizontal grid line (see Fig. 12.9) somewhere in the middle of the lattice.

Recall that $|ES| = n_0$ particles of all the active, via PFR, particles that continue to transmit, and let Q be a set of points whose shortest distance from particles in ES is maximum. Then L_Q is the locus (curve) of such points Q. The number of particles included in L_Q (i.e., the *area* inside L_Q) is the number k.

Now, if the distance of such points Q of L_Q from ES is ω then $k \leq 2(n_0 + 2\omega)\omega$ (see Fig. 12.10) and thus $r \leq \frac{2\omega(n_0+2\omega)}{n^2}$. Notice, however, that ω is a random variable hence we will estimate its expected value $\mathbf{E}(\omega)$ and the moment $\mathbf{E}(\omega^2)$ to get

$$\mathbf{E}(r) \leq \frac{4\mathbf{E}(\omega^2)}{n^2} + \frac{2n_0}{n^2}\mathbf{E}(\omega) \tag{12.1}$$

Now, look at points $Q : (\widehat{EQS}) = \phi < 30°$, i.e., $\phi < \frac{\pi}{6}$. We want to use the approximation $|ED| \simeq x$, where x is a random variable depending on the particle distribution.

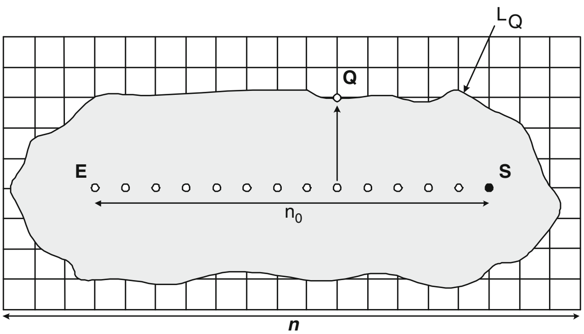

Fig. 12.9 The L_Q area

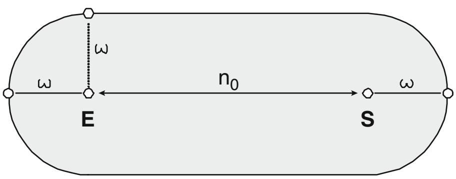

Fig. 12.10 The particles inside the L_Q area

Let $\frac{|ED|}{x} = 1 + \varepsilon$. Then a bound on the approximation factor ε determines a bound on the "cut-off" angle ϕ_0 since $|ED| \cos\left(\frac{\phi_o}{2}\right) = \varepsilon x$.

$$(1+\varepsilon)\, x \,\cos\left(\frac{\phi_o}{2}\right) = x \implies \cos\left(\frac{\phi_o}{2}\right) = \frac{1}{1+\varepsilon}$$

ϕ_0 above is constant that can be appropriately chosen by the protocol implementer. We chose $\phi_0 = 30°$ so $\varepsilon = 0.035$. Also, we remark here that the energy spent increases with $x_0 = n_0\left(1 - \frac{\xi}{2}\right)$ in the area below x_0, where $\xi = \tan 75°$, but decreases with x for $x > x_0$. This is a trade-off and one can carefully estimate the angle ϕ_0 (i.e., ϵ) to minimize the energy spent.

Let $\phi_1 = (\widehat{EQD})$, $\phi_2 = (\widehat{DQS})$ where $QD \perp ES$ and D in ES. In Fig. 12.11, $\phi = \phi_2 - \phi_1$.

We approximate ϕ by $\sin\phi_2 - \sin\phi_1$ (note: $\phi \simeq \sin\phi \geq -\sin\phi_1 + \sin\phi_2$).

$$\text{Note} \quad \sin\phi_1 \simeq \frac{|ED|}{|QE|} \simeq \frac{|ED|}{x} \quad \text{and} \quad \sin\phi_2 \simeq \frac{|DS|}{x}$$

$$\implies \sin\phi \simeq \sin\phi_2 - \sin\phi_1 \simeq \frac{|ES|}{x} = \frac{n_0}{x}$$

Thus $\frac{\phi}{\pi} \simeq \frac{n_0}{\pi x}$ for such points Q. We note that the analysis is very similar when $(\widehat{QES}) < \frac{\pi}{2}$ and thus $\phi = \phi_1 + \phi_2$.

Let $\mathcal{M}$ be the stochastic process that represents the vertical (to ES) distance of active points from ES, and $\mathcal{W}$ be the random walk on the vertical line yy' (see Fig. 12.12) such that, when the walk is at distance $x \geq x_0$ from ES then it (a) goes to $x+1$ with probability $\frac{n_0}{\pi x}$ or (b) goes to $x-1$ with probability $1 - \frac{n_0}{\pi x}$ and never goes below x_0, where x_0 is the "30°-distance" i.e. $x_0 = \frac{n_0 \xi}{2}$.

Clearly $\mathcal{W}$ dominates M, i.e., $\mathbf{P}_{\mathcal{M}}\{x \geq x_1\} \leq \mathbf{P}_{\mathcal{W}}\{x \geq x_1\}, \forall x_1 > x_0$.

Furthermore, $\mathcal{W}$ is dominated by the continuous time "discouraged arrivals" birth-death process $\mathcal{W}'$ (for $x \geq x_0$) where the rate of going from x to $x+1$ in $\mathcal{W}'$ is $\frac{\alpha}{x} = \frac{n_0/\pi}{x}$ and the rate of returning to $x-1$ is $1 - \frac{n_0}{\pi x_0} = 1 - \frac{2}{\pi\xi} = \mathbb{B}$.

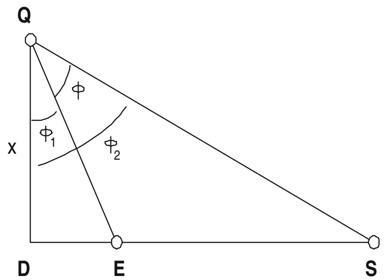

Fig. 12.11 The QES triangle

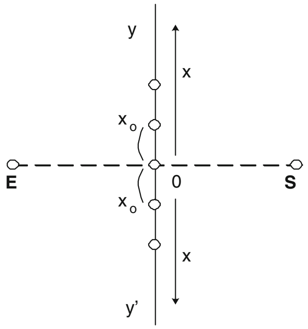

Fig. 12.12 The random walk $\mathcal{W}$

We know from [15] that for $\Delta x = x - x_0$ $(x > x_0)$

$$\mathbf{E}_{\mathcal{W}'}(\Delta x) = \frac{\alpha}{\mathbb{B}} = \frac{n_0}{\pi\left(1 - \frac{2}{\pi\xi}\right)}$$

Thus, $\mathbf{E}_{\mathcal{W}'}(x) = x_0 + \frac{\alpha}{\mathbb{B}}$, hence by domination

$$\mathbf{E}(\omega) \leq \mathbf{E}_{\mathcal{M}}(x) \leq x_0 + \frac{\alpha}{\mathbb{B}} \tag{12.2}$$

Also from [15] the process $\mathcal{W}'$ is a Poisson one and $\mathbf{P}\{\Delta x = k\} = \frac{(\alpha/\mathbb{B})^k}{k!} e^{-(\alpha/\mathbb{B})}$. From this, the variance of Δx is $\sigma^2 = \alpha/\mathbb{B}$ (again), i.e., for $\omega = x_0 + \Delta x$

$$\begin{aligned}
\mathbf{E}_{\mathcal{W}'}(\omega^2) &= \mathbf{E}_{\mathcal{W}'}\left((x_0 + \Delta x)^2\right) \\
&= x_0^2 + 2x_0\mathbf{E}_{\mathcal{W}'}(\Delta x) + \mathbf{E}_{\mathcal{W}'}(\Delta x^2) \\
&= x_0^2 + 2\frac{\alpha}{\mathbb{B}}x_0 + \left(\sigma^2 + \mathbf{E}^2(\Delta x)\right) \\
&= x_0^2 + 2\frac{\alpha}{\mathbb{B}}x_0 + \frac{\alpha}{\mathbb{B}} + \left(\frac{\alpha}{\mathbb{B}}\right)^2
\end{aligned}$$

$$\text{So} \quad \mathbf{E}_{\mathcal{W}'}(\omega^2) = \frac{3n_0^2}{4} + 2\frac{\alpha}{\mathbb{B}}\frac{n_0\,\xi}{2} + \frac{\alpha}{\mathbb{B}} + \left(\frac{\alpha}{\mathbb{B}}\right)^2$$

$$\text{where} \quad \frac{\alpha}{\mathbb{B}} = \frac{n_0}{\pi\left(1 - \frac{2}{\pi\,\xi}\right)} = \frac{n_0}{\tau}$$

$$\text{and} \quad \tau = \pi\left(1 - \frac{2}{\pi\,\xi}\right) = \pi - \frac{2}{\xi}$$

$$\text{thus} \quad \mathbf{E}_{\mathcal{W}'}(\omega^2) = n_0^2\left(\frac{3}{4} + \frac{\xi}{\tau} + \frac{1}{\tau^2}\right) + \frac{n_0}{\tau} \tag{12.3}$$

and by domination $\mathbf{E}(\omega^2) \leq \mathbf{E}_{\mathcal{M}}(\omega^2) \leq \mathbf{E}_{\mathcal{W}'}(\omega^2)$. So, finally

$$\mathbf{E}(r) \leq \frac{4\left[n_0^2\left(\frac{3}{4}+\frac{\xi}{\tau}+\frac{1}{\tau^2}\right)+\frac{n_0}{\tau}\right]}{n^2}+\frac{n_0^2}{n^2}\left(\xi+\frac{2}{\tau}\right) \tag{12.4}$$

which proves the theorem. □

12.3.5 The Robustness of PFR

We now consider the robustness of our protocol, in the sense that PFR must guarantee that data arrive at enough points in a small interval around S, in cases where part of the network has become inoperative, without, however, isolating the sink.

Lemma 8 PFR *manages to propagate the crucial data across lines parallel to ES, and of constant distance, with fixed nonzero probability (not depending on n, $|ES|$).*

Proof Here we consider particles very near the line ES. Thus the approximations that we do are suitable for nearby particles. Let Q' an active particle at vertical distance x from ES and $D \in ES$ such that $Q'D \perp ES$.

Let $yy' \parallel ES$, drawn from Q' (see Fig. 12.13) and $\phi_1 = (\widehat{yQ'E})$, $\phi_2 = (\widehat{SQ'y'})$. Then $\phi = 180° - (\phi_1 + \phi_2)$, i.e., $\frac{\phi}{\pi} = 1 - \frac{\phi_1+\phi_2}{\pi}$.

Since ϕ_1, ϕ_2 are small (for small x) we use the approximation $\phi_1 \simeq \sin\phi_1$ and $\phi_2 \simeq \sin\phi_2$

$$\phi_1 + \phi_2 \simeq \sin\phi_1 + \sin\phi_2 = \frac{x}{EQ'} + \frac{x}{Q'S}$$

Since $\phi > 90°$, ES is the biggest edge of triangle $(\widehat{EQ'S})$, thus $EQ' \leq n_0$ and $Q'S \leq n_0$. Hence,

$$\sin\phi_1 + \sin\phi_2 \leq \frac{2x}{n_0}$$

$$\text{and} \quad 1 - \frac{\sin\phi_1 + \sin\phi_2}{\pi} \geq 1 - \frac{2x}{\pi n_0} \quad \text{(for small } x\text{)}$$

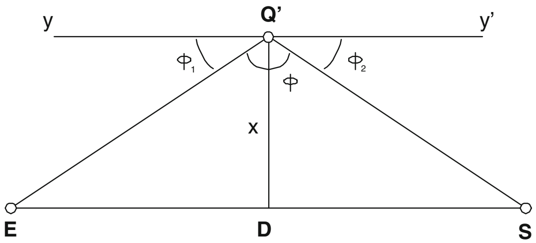

Fig. 12.13 The $Q'ES$ triangle

Without loss of generality, assume ES is part of a horizontal grid line. Here we study the case in which some of the particles on ES or very near ES (i.e. at angles $> 134°$) are not operating.

Consider now a horizontal line ℓ in the grid, at distance x from ES. Clearly, some particles of ℓ near E will be activated during the initial broadcasting phase (see Phase 1).

Let $\mathcal{A}$ be the event that all the particles of ℓ (starting from p) will be activated, until a vertical line from S is reached. Then

$$\mathbf{P}(\mathcal{A}) \geq \left(1 - \frac{2x}{\pi n_0}\right)^{n_0} \simeq e^{-\frac{2x}{\pi}} \qquad \square$$

12.4 An Experimental Comparison of LTP, PFR

We evaluate the performance of four protocols (PFR, LTP, and two variations of LTP) by a comparative experimental study. The protocols have been implemented as C++ classes using the data types for two-dimensional geometry of LEDA [17] based on the environment developed in [3, 6]. Each class is installed in an environment that generates sensor fields given some parameters (such as the area size of the field, the distribution function used to drop the particles) and performs a network simulation for a given number of repetitions, a fixed number of particles, and certain protocol parameters.

In the experiments, we generate a variety of sensor fields in a 100 m × 100 m square. In these fields, we drop $n \in [100, 3000]$ particles randomly uniformly distributed on the sensor cloud plane, i.e., for density $0.01 \leq d \leq 0.3$. Each sensor particle has a fixed radio range of $\mathcal{R} = 5m$ and $\alpha = 90°$. The particle p that initially senses the crucial event is always explicitly positioned at $(x, y) = (0, 0)$ and the sink is located at $(x, y) = (100, 100)$. Note that this experimental setup is based on and extends that used in [10, 13, 16]. We repeated each experiment for more than 5,000 times in order to achieve good average results.
We now define the efficiency measures investigated.

Definition 5 Let h_{A} (for "active") be the *number of "active" sensor particles* participating in the data propagation and let E_{TR} be the *total number of data transmissions* during propagation. Let T be the *total time* for the propagation process to reach its final position and H the *total number of "hops"* required.

Clearly, by minimizing h_{A} we succeed in avoiding flooding and thus we minimize energy consumption. Remark that in LTP we count as active those particles that transmit *info($\mathcal{E}$)* at least once.

Note that h_{A}, E_{TR}, T, and H are random variables. Furthermore, we define the success probability of the algorithm where we call success the eventual data propagation to the sink.

Definition 6 Let $\mathbf{P}_s$ be the *success probability* of the protocol.

We also focus on the study of the following parameter. Suppose that the data propagation fails to reach the sink. In this case, it is very important to know *"how close" to the sink* it managed to get. Propagation reaching close to the sink might be very useful, since the sink (which can be assumed to be mobile) could itself move (possibly by performing a random walk) to the final point of propagation and get the desired data from there. Even assuming a fixed sink, proximity to it is important, since the sink might in this case begin some "limited" flooding to get to where data propagation stopped. Clearly, the closer to the sink we get, the cheaper the flooding becomes.

Definition 7 Let F be the final position of the data propagation process. Let $\mathcal{D}$ be $\mathcal{F}$'s (Euclidean) distance from the sink $\mathcal{S}$.

Clearly in the case of total success $\mathcal{F}$ coincides with the sink and $\mathcal{D} = 0$.

We start by examining the success rate of the four protocols (see Fig. 12.14), for different particle densities. Initially, when the density is low (i.e., $d \leq 0.06$), the protocols fail to propagate the data to the sink. However, as the density increases, the success rate increases quite fast and for high densities all four protocols almost always succeed in propagating the data to the sink. Thus, all protocols are very successful. We remark a similar shape of the success rate function in terms of density. This is due to the fact that all protocols use local information to decide how to proceed by basically selecting (all protocols) the next particle with respect to a similar criterion (best progress toward the sink).

In the case when the protocols fail to propagate the data to the sink, we examine *"how close" to the sink* they managed to get. Fig. 12.15 depicts the distance of the final point of propagation to the position of the sink. Note that this figure should

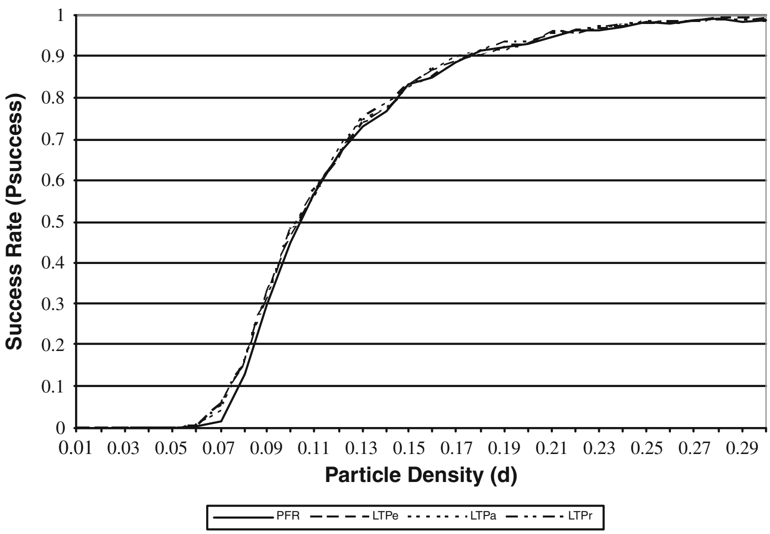

Fig. 12.14 Success probability ($\mathbf{P}_s$) over particle density $d = [0.01, 0.3]$

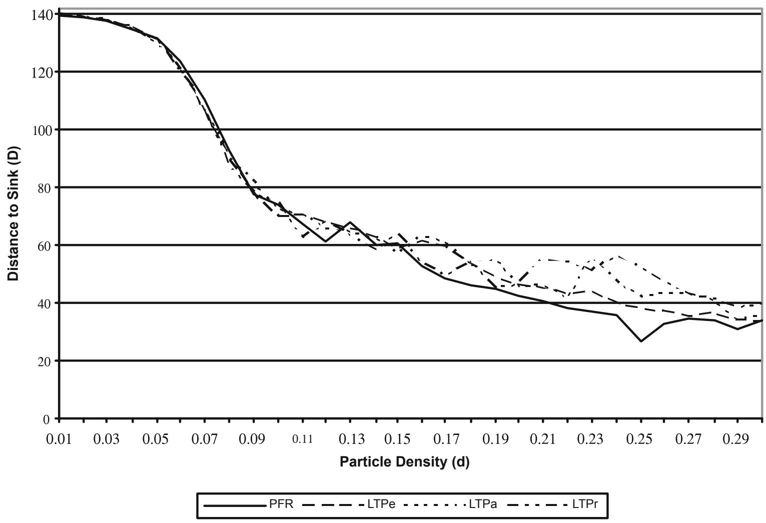

Fig. 12.15 Average distance from the sink ($\mathcal{D}$) over particle density $d = [0.01, 0.3]$

be considered in conjunction with Fig. 12.14 on the success rate. Indeed, failures to reach the sink are very rare and seem to appear in very extreme network topologies due to bad particle distribution in the network area.

Figure 12.16 depicts the ratio of active particles over the total number of particles ($r = \frac{h_A}{n}$) that make up the sensor network. In this figure we clearly see that PFR, for low densities (i.e., $d \leq 0.07$), indeed activates a small number of particles (i.e.,

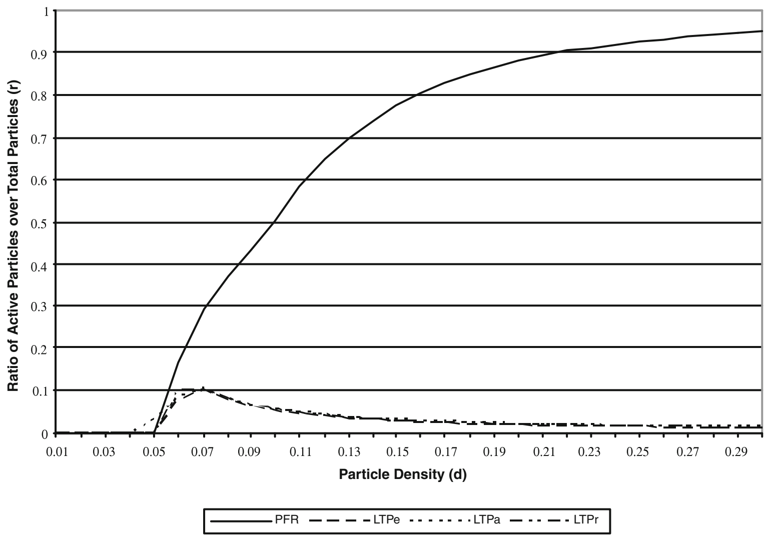

Fig. 12.16 Ratio of active particles over total particles (r) over particle density $d = [0.01, 0.3]$

$r \leq 0.3$) while the ratio r increases fast as the density of the particles increases. The LTP-based protocols seem more efficient and the ratio r seems to be independent of the total number of particles (because only one particle in each hop becomes active).

Remark 1 Because of the way PFR attempts to avoid flooding (by using angle ϕ to capture "distance" from optimality) its merits are not sufficiently shown in the setting considered here. We expect PFR to behave significantly better with respect to energy in much larger networks and in cases where the event is sensed in an average place of the network. Also, stronger probabilistic choices (i.e., $\mathbf{P}_{\text{fwd}} = \left(\frac{\phi}{\pi}\right)^{\alpha}$, where $a > 1$ a constant) may further limit the number of activated particles.

Furthermore, examining the total number of transmissions performed by the particles (see Fig. 12.17), it is evident that because the LTP-based protocols activate a small number of particles, the overall transmissions are kept low. This is a surprising finding, since the PFR protocol was originally designed to work without the need of any control messages so that the energy consumption is low. However, the comparative study clearly shows that avoiding the use of control messages does not achieve the expected results. So, even though all four protocols succeed in propagating the data, it is evident that the LTP-based protocols are more energy efficient in the sense that less particles are involved in the process.

We continue with the following two parameters: (a) the "hops" efficiency and (b) the time efficiency, measured in terms of rounds needed to reach the sink. As can be seen in Fig. 12.18, all protocols are very efficient, in the sense that the number of hops required to get to the sink tends below 40 even for densities $d \geq 0.17$. The value 40 in our setting is close to optimality since in an ideal placement, the diagonal line is of length $100\sqrt{2}$ and since for the transmission range $\mathcal{R} = 5$ the optimal number of hops (in an ideal case) is roughly 29. In particular PFR achieves

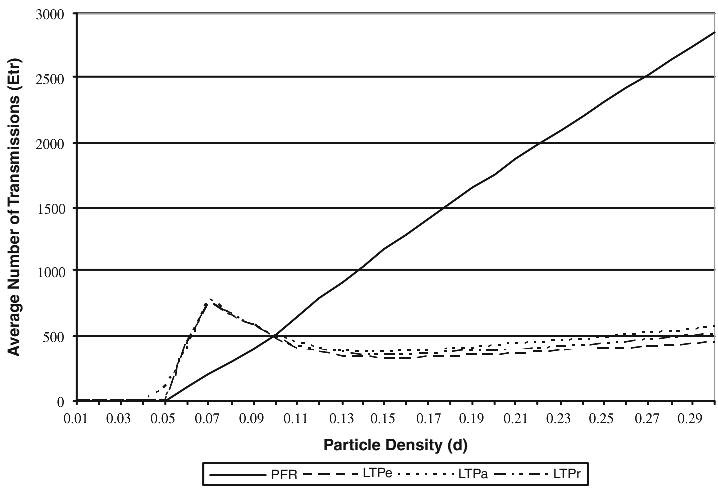

Fig. 12.17 Average number of transmissions (E_{TR}) over particle density $d = [0.01, 0.3]$

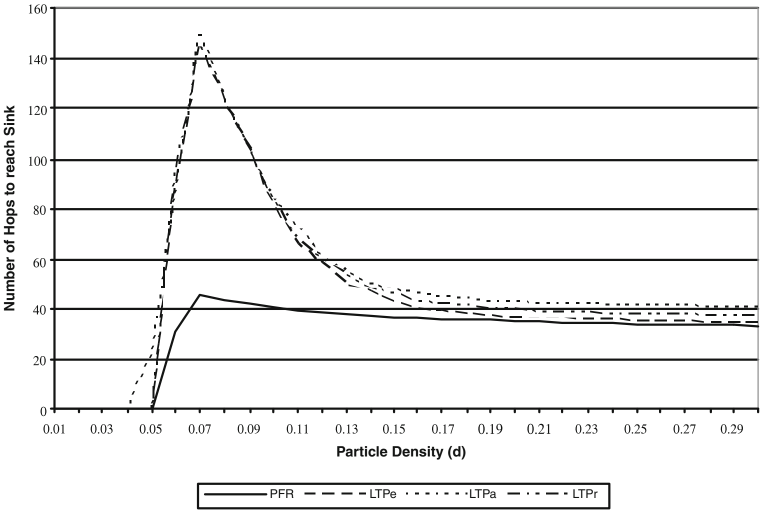

Fig. 12.18 Average number of hops to reach the sink (H) over particle density $d = [0.01, 0.3]$

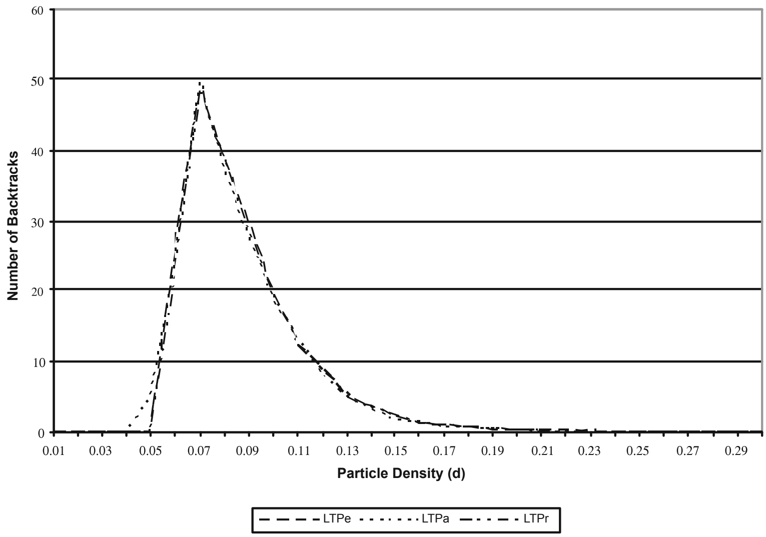

Fig. 12.19 Average number of backtracks over particle density $d = [0.01, 0.3]$

this for very low densities ($d \geq 0.07$). On the other hand, the LTP-based protocols exhibit a certain pathological behavior for low densities (i.e., $d \leq 0.12$) due to a high number of executions of the backtrack mechanism in the attempt to find a particle closer to the sink (see also Fig. 12.19).

Finally, in Fig. 12.19 we compare the three LTP-based protocols and the number of backtracks invoked in the data propagation. It is evident that for very low particle densities (i.e., $d \leq 0.12$), all three protocols perform a large number of backtracks in order to find a valid path toward the sink. As the particle density increases, the number of backtrack reduces fast enough and almost reaches zero.

12.5 Conclusions

We investigated some important aspects of online energy optimization in sensor networks, like minimizing the total energy spent in the network, minimizing the number (or the range) of data transmissions, combining energy efficiency and fault tolerance (by allowing redundant data transmissions which, however, should be optimized to not spend too much energy). Since it is very difficult (if possible at all) to achieve all the above goals at the same time we presented two characteristic protocols, each of which focuses on some of the energy efficiency goals above (while still performing well with respect to the rest goals as well). In particular, we presented: (a) The Local Target Protocol (LTP) that performs a local optimization trying to minimize the number of data transmissions and (b) the Probabilistic Forwarding Protocol (PFR) that creates redundant data transmissions that are probabilistically optimized, to trade off energy efficiency with fault tolerance.

Open issues for further research include considering the impact on protocols' performance of other important types of faults (malicious, byzantine). Also, it is worth investigating other performance aspects like congestion (possibly using game-theoretic methods).

Acknowledgments This work has been partially supported by the IST Programme of the EU under contract number IST-2005-15964 (**AEOLUS**).

References

1. I.F. Akyildiz, W. Su, Y. Sankarasubramaniam, and E. Cayirci. Wireless sensor networks: A survey. *In the Journal of Computer Networks*, 38: 393–422, 2002.
2. I.F. Akyildiz, W. Su, Y. Sankarasubramaniam, and E. Cayirci. A survey on sensor networks. *IEEE Communications Magazine*, 102–114, August 2002.
3. I. Chatzigiannakis, S. Nikoletseas, and P. Spirakis. Smart dust protocols for local detection and propagation. In: *Proceedings 2nd ACM Workshop on Principles of Mobile Computing* – POMC'2002, pages 9–16. Also, in the *ACM Mobile Networks (MONET) Journal, Special Issue on Algorithmic Solutions for Wireless, Mobile, Adhoc and Sensor Networks*, MONET 10 (1): 133–149, 2005.
4. I. Chatzigiannakis, T. Dimitriou, S. Nikoletseas, and P. Spirakis. A Probabilistic algorithm for efficient and robust data propagation in smart dust networks. *In the Proceedings of the 5th European Wireless Conference on Mobile and Wireless Systems beyond 3G (EW 2004)*, pages 344–350, 2004. Also, in the *Ad-Hoc Networks Journal, Elsevier*, 4(5): 621–635, 2006.
5. I. Chatzigiannakis, T. Dimitriou, M. Mavronicolas, S. Nikoletseas, and P. Spirakis. A comparative study of protocols for efficient data propagation in smart dust networks. In *Proceedings International Conference on Parallel and Distributed Computing* – EUPOPAR 2003 (Distinguished Paper), pages 1003–1016. Also in the *Parallel Processing Letters (PPL) Journal*, 13(4): 615–627, 2003.

6. I. Chatzigiannakis and S. Nikoletseas. A Sleep-awake protocol for information propagation in smart dust networks. In: *Proceedings 3rd Workshop on Mobile and Ad-Hoc Networks (WMAN)*–IPDPS Workshops, IEEE Press, page 225, 2003. Also, in the *ACM/Baltzer Journal of Mobile Networks and Applications* (MONET), 9(4): 319–332, 2004.
7. D. Estrin, R. Govindan, J. Heidemann, and S. Kumar. Next century challenges: Scalable coordination in sensor networks. In: *Proceedings 5th ACM/IEEE International Conference on Mobile Computing* – MOBICOM'1999, IEEE Computer Society, Seattle.
8. H. Euthimiou, S. Nikoletseas, and J. Rolim. Energy balanced data propagation in wireless sensor networks. In: *Proceedings 4th International Workshop on Algorithms for Wireless, Mobile, Ad-Hoc and Sensor Networks (WMAN '04), IPDPS* 2004, 2004. Also, in the *Wireless Networks (WINET) Journal*, 12(6): 691–707, 2006.
9. W. Feller. An introduction to probability theory and its applications, volume I, Wiley, 1968.
10. W. R. Heinzelman, A. Chandrakasan, H. Balakrishnan: Energy-Efficient Communication Protocol for Wireless Microsensor Networks. In: *Proceedings 33rd Hawaii International Conference on System Sciences* – HICSS'2000, pages 876–882, IEEE Computer Society, Maui, HI, USA.
11. C. Intanagonwiwat, R. Govindan, and D. Estrin. Directed Diffusion: A scalable and robust communication paradigm for sensor networks. In: *Proceedings 6th ACM/IEEE International Conference on Mobile Computing* – MOBICOM'2000, pages 56–67, IEEE Computer Society, Boston.
12. C. Intanagonwiwat, R. Govindan, D. Estrin, J. Heidemann, and F. Silva: Directed diffusion for wireless sensor networking. Extended version of [11].
13. C. Intanagonwiwat, D. Estrin, R. Govindan, and J. Heidemann: Impact of network density on data aggregation in wireless sensor networks. Technical Report 01–750, University of Southern California Computer Science Department, November 2001.
14. J.M. Kahn, R.H. Katz, and K.S.J. Pister: Next Century Challenges. Mobile networking for smart dust. In: *Proceedings 5th ACM/IEEE International Conference on Mobile Computing*, pages 271–278, September 1999.
15. L. Kleinrock. Queueing systems, Theory, Wiley Volume. I, page 100, 1975.
16. A. Manjeshwar and D.P. Agrawal. TEEN: A routing protocol for enhanced efficiency in wireless sensor networks. In: *Proceedings 2nd International Workshop on Parallel and Distributed Computing Issues in Wireless Networks and Mobile Computing*, satellite workshop of *16th Annual International Parallel & Distributed Processing Symposium* – IPDPS'02, IEEE Computer Society, Fort Lauderdale, Florida.
17. K. Mehlhorn and S. Näher. *LEDA: A platform for combinatorial and geometric computing*. Cambridge University Press, Cambridge, 1999.
18. S. M. Ross. *Stochastic processes, 2nd Edition*. Wiley, 1995.
19. P. Triantafilloy, N. Ntarmos, S. Nikoletseas, and P. Spirakis. Nanopeer networks and P2P Worlds. In: *Proceedings 3rd IEEE International Conference on Peer-to-Peer Computing*, pages 40–46, Linkoeping, Sweden, 2003.