Smart Dust Protocols for Local Detection and Propagation *

[Extended Abstract]

Ioannis Chatzigiannakis
Computer Technology Institute
and Patras University
261 10 Patras, Greece
ichatz@cti.gr

Sotiris Nikoletseas
Computer Technology Institute
and Patras University
261 10 Patras, Greece
nikole@cti.gr

Paul Spirakis
Computer Technology Institute
and Patras University
261 10 Patras, Greece
spirakis@cti.gr

ABSTRACT

Smart Dust is a set of a vast number of ultra-small fully autonomous computing and communication devices, with very restricted energy and computing capabilities, that co-operate to quickly and efficiently accomplish a large sensing task. Smart Dust can be very useful in practice i.e. in the local detection of a remote crucial event and the propagation of data reporting its realization. In this work we make an effort towards the research on smart dust from a basic algorithmic point of view. We first provide a simple but realistic model for smart dust and present an interesting problem, which is how to propagate efficiently information on an event detected locally. Then we present smart dust protocols for local detection and propagation that are simple enough to be implemented on real smart dust systems, and perform, under some simplifying assumptions, a rigorous average case analysis of their efficiency and energy consumption (and their interplay). This analysis leads to concrete results showing that our protocols are very efficient.

Categories and Subject Descriptors

C.2.2 [**Computer - Communication Networks**]: Network Protocols—*Routing Protocols*; C.2.4 [**Computer - Communication Networks**]: Distributed Systems—*Distributed applications*; C.3 [**Special - Purpose and Application - Based Systems**]: Real-time and embedded systems; C.4 [**Performance of Systems**]: Fault Tolerance; G.3.4 [**Probability and Statistics**]: Stochastic processes; F.2.0 [**Analysis of Algorithms and Problem Complexity**]: General

General Terms

Algorithms, Design, Performance, Theory, Experimentation

*This work has been partially supported by the IST Program of the European Union under contract numbers IST-1999-14186 (ALCOM-FT) and IST-2001-33116 (FLAGS).

POMC'02, October 30-31, 2002, Toulouse, France.

Keywords

Sensor Networks, Smart Dust, Local Detection, Propagation, Average Case Analysis

1. INTRODUCTION

Networked sensors (or Smart Dust) are very large systems, comprised of a vast number of homogenous ultra-small fully autonomous computing and communication devices that co-operate to achieve a large sensing task. Each device has one or more sensors, embedded processors and low-power radios, and is normally battery operated. Examining each such single device individually, might appear to have small utility. The realization of Smart Dust however, lies in using and co-coordinating a vast number of such devices.

Smart Dust is a useful case of dynamic environments of networked sensors that are spread over a global system and try to communicate and compute efficiently and quickly, having only partial knowledge of the global conditions and having poor energy and computing resources. Typically, these networked sensors coordinate to perform a common task. Designing protocols to coordinate such systems (i.e. create a dynamic and efficient network of these sensors) and monitoring their behavior as they operate in complex and dynamic global environments is of great importance for information gathering and processing in many practical situations.

As an example, [4] points that integrated low-power sensing devices will permit remote object monitoring and tracking in inhospitable physical environments such as remote geographic regions or toxic urban locations. They will also enable low maintenance sensing in the field (vehicles, equipment, personnel), the office buildings (projectors, furniture, books, people), the factory floor (motors, small robotic devices).

There are many possible models for such networked sensors. In this work, we consider networked sensors where (a) all nodes in the network are homogenous and constrained by low availability of resources (energy, communication) and (b) the data being sensed by the nodes must be transmitted to *a fixed control center located far away from the sensors.* Thus direct communication between the sensor nodes and the control center is impossible and/or expensive, since there are no "high-energy" nodes through which communication can proceed. This is the general framework for MIT's μAMPS project [12], which focuses on innovative energy-optimized solutions at all levels of the system hierarchy, from the physical layer and communication protocols up to the application layer.

To motivate the challenges in designing such sensor networks, we can consider the following scenario where local detection and fast propagation to the authorities of the realization of a crucial event can be achieved using smart dust. Think about thousand of disposable sensors scattered (e.g. thrown from an aircraft) over a forest. Each of these sensors can monitor the temperature at a single, very small geographical area. The sensors coordinate to establish an efficient, dynamic and short-lived communication network, dividing the task of monitoring the terrain and offering continuous monitoring of the environment in order to alert the authorities as soon as possible after a forest fire is detected by some sensor.

Several aspects of such systems of autonomous networked entities emerge, which are quite different from those posed by standard computer networks. Such aspects include the very poor and highly restricted resources (e.g. very low battery power, low computing capabilities, total absence of synchrony and anonymity). Network protocols must be designed to achieve fault tolerance in the presence of individual node failure while minimizing energy consumption. Another important aspect is the scalability to the change in network size, node density and topology. The network topology changes over time as some nodes may die, or possibly because new nodes join later.

This work, continuing our line of research on communication in ad-hoc mobile networks [2, 3], is an attempt towards capturing the underlying foundational and algorithmic issues in the design of such systems, abstracting accurately enough the real technological specifications involved and providing some first concrete results for the efficiency of a variety of smart dust protocols using an average case analysis. We focus in this paper on the efficient use of smart dust in local detection and propagation protocols. We first provide an abstract model for smart dust systems which is simple enough to allow an analysis to develop, being however at the same time quite realistic, in terms of the technological specifications of real smart dust systems it captures. Then we define the problem of local detection and propagation using smart dust and also propose some concrete measures for the average case analysis of protocols for this problem.

Our Results: For the local detection and propagation problem using smart dust, we provide two protocols. Both protocols are simple enough to be implemented in real smart dust systems despite the severe energy and computing power limitations of such systems. Furthermore, we give a rigorous average case analysis for the efficiency of these protocols. We consider a variety of performance and robustness criteria, such as propagation time, number of particle to particle transmissions (which also characterizes energy consumption and time efficiency, assuming an efficient MAC protocol) and fault-tolerance:

1. Our first protocol, which we call the "local target protocol", uses a fast and cheap search phase which is assumed to always return a nearby particle towards the authorities, uniformly in some range. It achieves a propagation time and an energy consumption whose expected ratio over the optimal solutions is at most $\pi/2 \simeq 1.57$.

2. Our second protocol (the "min-two uniform targets" protocol) applies the simple idea of getting at least two particles towards the authorities and selecting the best in terms of propagation progress. It is in fact an optimized version of the local target protocol, and has an expected time and energy ratio over the optimal solutions which is at most $\pi^2/8 \simeq 1.24$.

3. Next we provide tight upper bounds to the distribution of the number of particle to particle data transmissions of a generalized target protocol.

Discussion of Selected Related Work: In the last few years, Sensor Networks have attracted a lot of attention from researchers at all levels of the system hierarchy, from the physical layer and communication protocols up to the application layer.

At the MAC level, many researchers have done research work in an effort to minimize the power consumption. [17] presents a contention-based protocol that tries to minimize energy consumption due to node idle listening, by avoiding the overhearing among neighboring nodes. A recent work [19] exploits a similar method for energy savings, and further reduce idle listening by avoiding any use of out-of-channel signaling. Additionally, their protocol trades off per-node fairness for further energy savings.

For establishing communication and routing information to the control center, mobile ad-hoc routing protocols ([15]) may be used in sensor networks. However, although protocols for mobile ad-hoc networks take into consideration energy conservation issues, most of them are not really suitable for sensor networks. [11] presents a routing protocol suitable for sensor networks that makes *greedy* forwarding decisions using only information about a node's immediate neighbors in the network topology. This approach achieves high scalability as the density of the network increases. [7] presents a clustering-based protocol that utilizes randomized rotation of local cluster heads to evenly distribute the energy load among the sensors in the network. [13] introduces a new energy efficient routing protocol that does not provide periodic data monitoring (as in [7]), but instead nodes transmit data only when sudden and drastic changes are sensed by the nodes. As such, this protocol is well suited for time critical applications and compared to [7] achieves less energy consumption and response time.

A family of negotiation-based information dissemination protocols suitable for wireless sensor networks is presented in [8]. Sensor Protocols for Information via Negotiation (SPIN) focus on the efficient dissemination of individual sensor observations to all the sensors in a network. However, in contrast to classic flooding, in SPIN sensors negotiate with each other about the data they possess using meta-data names. These negotiations ensure that nodes only transmit data when necessary, reducing the energy consumption for useless transmissions.

A data dissemination paradigm called *directed diffusion* for sensor networks is presented in [9], where data-generated by sensor nodes is named by attribute-value pairs. An observer requests data by sending *interests* for named data; data matching the interest is then "drawn" down towards that node by selecting a single path or through multiple paths by using a low-latency tree. [10] presents an alternative approach that constructs a greedy incremental tree that is more energy-efficient and improves path sharing.

We note that, as opposed to the work presented in this paper, the above research focuses on energy consumption without examining the time efficiency of their protocols. Furthermore, these works contain basically protocol design and technical specifications, while quantitative aspects are only experimentally evaluated and no theoretical analysis is given. Note also that our protocols are quite general in the sense that (a) do not assume global network topology information, (b) do not assume geolocation information (such as GPS information) and (c) use very limited control message exchanges, thus having low communication overhead.

2. THE MODEL

Smart dust is comprised of a vast number of ultra-small homogenous sensors, which we call "grain" particles. Each smart-dust *grain particle* is a fully-autonomous computing and communication device, characterized mainly by its available power supply (battery) and the energy cost of computation and transmission of data. Such particles cannot move.

Each particle is equipped with a set of monitors (sensors) for light, pressure, humidity, temperature etc. Each particle has two communication modes: a *broadcast* (digital radio) *beacon mode* (for low energy - short signals) and a *directed to a point* actual data transmission mode (usually via a laser beam). Also, in a variation of our model capturing energy saving specifications, each particle may alternate (independently of other particles) between a *sleeping* and an *awake* mode. During sleeping periods grain particles cease any communication with the environment, thus they are unable to listen, receive and propagate data transmitted by other particles.

We adopt here (as a starting point) a two-dimensional (plane) framework: A *smart dust cloud* (a set of particles) is spread in an area (for a graphical presentation, see fig. 1). Note that a two-dimensional setting is also used in [7, 8, 9, 10, 13].

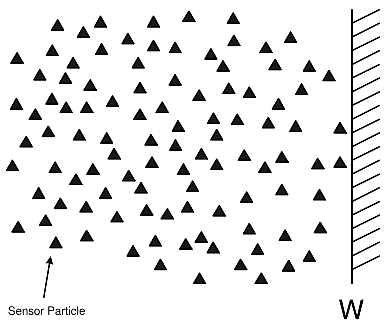

Figure 1: A Smart Dust Cloud

DEFINITION 2.1. *Let d (usually measured in numbers of particles/m^2) be the* density *of particles in the area. Let $\mathcal{R}$ be the maximum (beacon/laser) transmission range of each grain particle. A* receiving wall $\mathcal{W}$ *is defined to be an infinite line in the smart-dust plane. Any particle transmission within range $\mathcal{R}$ from the wall $\mathcal{W}$ is received by $\mathcal{W}$.*

We assume that $\mathcal{W}$ has very strong computing power, able to collect and analyze received data and has a constant power supply and so has no energy constraints. The wall represents in fact the authorities (the fixed control center) who the realization of a crucial event should be reported to. Note that a wall of appropriately big (finite) length suffices. We plan to conduct an analysis of the (expected and/or with high probability) deviation of the transmitted data from the vertical to the wall position in order to provide upper bounds on the wall's length needed.

Furthermore, we assume that there is a set-up phase of the smart dust network, during which the smart cloud is dropped in the terrain of interest, when using special control messages (which are very short, cheap and transmitted only once) each smart dust particle is provided with the direction of $\mathcal{W}$. By assuming that each smart-dust particle has individually *a sense of direction* (e.g. through its magnetometer sensor), and using these control messages, each particle is aware of the general location of $\mathcal{W}$.

We feel that our model, although simple, depicts accurately enough the technological specifications of real smart dust systems. Similar models are being used by other researchers in order to study sensor networks (see [7, 13]). In contrast to [9, 11], our model is weaker in the sense that *no geolocation abilities* are assumed (e.g. a GPS device) for the smart dust particles leading to more generic and thus stronger results. In [6] a thorough comparative study and description of smart dust systems is given, from the technological point of view. In the Appendix we report some basic technical characteristics which we took into account when defining the model of smart dust we use here.

3. THE PROBLEM

An *adversary* $\mathcal{A}$ selects a single particle, p, in the plane-cloud and allows it to monitor a *local crucial event* $\mathcal{E}$. The general *propagation problem* $\mathcal{P}$ is the following:

> "How can particle p, via cooperation with the rest of the cloud, propagate information about event $\mathcal{E}$ to the receiving wall $\mathcal{W}$?"

DEFINITION 3.1. *Let $h_{opt}(p, \mathcal{W})$ be the (optimal) number of "hops" (direct, vertical to $\mathcal{W}$ transmissions) needed to reach the wall, in the* ideal *case in which particles always exist in pair-wise distances $\mathcal{R}$ in 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 number of hops (transmissions) taken to reach $\mathcal{W}$. The* "hops" efficiency *of protocol π is the ratio*

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

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

In the case where 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 we study here 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 $\mathcal{W}$ passing from p. 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 available in the right direction. c) Particles may temporarily "sleep" (i.e. not listen to transmissions) in order to save battery power.

Remark: Note that any given distribution of particles in the smart dust 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 have chosen, however, to compare the efficiency of our 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 our protocols.

4. THE LOCAL TARGET PROTOCOL

Let $d(p_i, p_j)$ the distance (along the corresponding vertical lines towards $\mathcal{W}$) 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. In this protocol, each particle p' that has received *info($\mathcal{E}$)* from p (via, possibly, other particles) does the following:

- *Search Phase*: It uses a periodic low energy broadcast of a beacon 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'' via a direct line (laser) time consuming transmission.
- *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 p (i.e. to the particle that it originally received the information).

Note that one can estimate an a-priori upper bound on the number of repetition of the search phase needed, by using the probability of success of each search phase. This bound can be used to decide when to backtrack.

Also note that the maximum distance $d(p', p'')$ is $\mathcal{R}$, i.e. the beacon transmission range (for a graphical representation see figures 2,3).

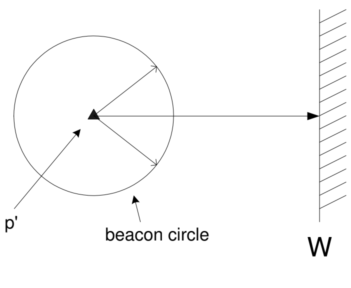

Figure 2: Example of the Search Phase

To enable a first step towards a rigorous analysis of smart dust protocols, we make the following simplifying assumption. *The search phase takes zero time* and always finds a

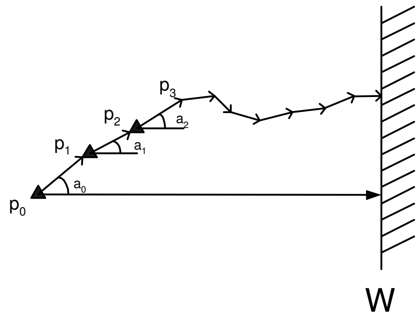

Figure 3: Example of a Transmission

p'' (of sufficiently high battery) in the semicircle of center p, in the direction towards $\mathcal{W}$. Note that this assumption on 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) we may consider, 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 smart cloud, (c) if the protocol during a search phase ultimately fails to find a particle towards the wall, it may *backtrack*.

In this analysis we do not consider the energy spent in the search phase. Note however that even the case where this is comparable to the energy spent in actual data transmission, the number of hops accounts for both (total energy spent is upper bounded by a multiple of actual data transmission energy).

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). We also assume that each target selection is random *independent* of the others, in the sense that it is always drawn uniformly in the arc $(-\alpha, \alpha)$.

We are aware of the fact that the above assumptions may not be very realistic in practice, however, they allows us to perform a first effort towards providing some concrete analytical results.

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 part of 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 towards $\mathcal{W}$ is then $\Delta_i = d(p_{i-1}, p_i) = \mathcal{R}\cos\alpha_i$, we get:

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

From Wald's equation for the expectation of a sum of a random number of independent random variables (see [16]), then

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

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

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

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

LEMMA 4.1. *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* $0 \le \alpha \le \frac{\pi}{2}$.

5. LOCAL OPTIMIZATION - THE "MIN TWO UNIFORM TARGETS" PROTOCOL

Note that the same basic framework holds for any situation in which the local (vertical) progress in the direction towards $\mathcal{W}$ (i.e. Δ_i) is of the same, independent, distribution. i.e. it always holds (via the Wald's equation) that

$$\frac{\mathcal{R}}{E(\Delta_i)} \;\le\; \frac{E(h)}{h_{opt}} \;\le\; \frac{\mathcal{R}}{E(\Delta_i)} + \frac{1}{h_{opt}} \qquad \Rightarrow$$

$$\Rightarrow \quad E(C_h) \;\simeq\; \frac{\mathcal{R}}{E(\Delta_i)} \;=\; \frac{1}{E(\cos\alpha_i)} \tag{1}$$

for large h. To understand the power of this, let us assume that the search phase always returns *two points* p'', p''' each uniform in $(-\alpha,\,\alpha)$ and that the protocol selects the best of the two points, with respect to the local (vertical) progress.

Let α_{i1}, α_{i2} the angles of the particles found and let $\alpha_i = \min\{|\alpha_{i1}|,\,|\alpha_{i2}|\}$. Then, for any $0 \le \phi \le \alpha$, it is:

$$\begin{aligned} \mathbb{P}\{\alpha_i > \phi\} \;&=\; \mathbb{P}\{|\alpha_{i1}| > \phi \,\cap\, |\alpha_{i2}| > \phi\} \\ &=\; \left(\frac{2\alpha - 2\phi}{2\alpha}\right)^2 \;=\; \left(\frac{\alpha-\phi}{\alpha}\right)^2 \end{aligned}$$

Thus, the distribution function of α_i, for any $0 \le \phi \le \alpha$, is

$$F_{\alpha_i}(\phi) \;=\; \mathbb{P}\{\alpha_i \le \phi\} \;=\; 1 - \frac{(\alpha-\phi)^2}{\alpha^2} \;=\; \frac{2\alpha\phi - \phi^2}{\alpha^2}$$

and the probability density function is, for any $0 \le \phi \le \alpha$:

$$f_{\alpha_i}(\phi) \;=\; \frac{d}{d\phi}\mathbb{P}\{\alpha_i \le \phi\} \;=\; \frac{2}{\alpha}\left(1 - \frac{\phi}{\alpha}\right)$$

The expected local progress is:

$$E(\cos\alpha_i) \;=\; \int_0^{\alpha}\cos\phi\, f_{\alpha_i}(\phi)\,d\phi \;=\; \frac{2(1-\cos\alpha)}{\alpha^2} \tag{2}$$

LEMMA 5.1. *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* $0 \le \alpha \le \frac{\pi}{2}$ *and for large* h.

We remark that

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

and

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

LEMMA 5.2. *The expected "hops efficiency of the min-two uniform targets protocol is* $1 \le E(C_h) \le \frac{\pi^2}{8} \simeq 1.24$ *for large* h *and for* $0 \le \alpha \le \frac{\pi}{2}$.

We remark that, w.r.t. 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 $(\pi/2 - \pi^2/8)/(\pi/2) \simeq 21.5\%$. We experimentally investigate the further gain of additional (i.e. $m > 2$) searches in section 7.

6. 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 6.1. (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 6.2. (Probability of significant progress) *Let* $m > 0$ *be the least integer such that* $\mathbb{P}\{\Delta x > \frac{\mathcal{R}}{m}\} \ge p$, *where* $0 < p < 1$ *is a given constant.*

LEMMA 6.1. *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 $1 - \tilde{F}(\mathcal{R}/m) \ge p$ i.e. to find the first m such that $1 - p \ge \tilde{F}(\mathcal{R}/m)$. Such an m always exists since $\tilde{F}$ is continuous in $[0, 1]$. $\square$

DEFINITION 6.3. *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 6.2. $\mathbb{P}_P\{h \le h_0\} \le \mathbb{P}_Q\{h \le h_0\}$

PROOF. The actual process Q makes always more progress than P. $\square$

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 $\mathbb{P}\{H = i\} = q^i(1-q)$ for any $i \ge 0$. 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 6.3. $\mathbb{P}_P\{$*number of hops is* $\mathrm{h}\} = \mathbb{P}\{\mathrm{H}_1 + \cdots \mathrm{H}_t = \mathrm{h}\}$

The probability generating function of H is

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

i.e.

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

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 theorem 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 [5], Vol. I. p. 253). Thus,

Theorem 6.4. $\mathbb{P}_P\{$*the number of hops is* $h\} = \binom{-t}{h}p^t(-q)^h = \binom{t+h-1}{h}p^tq^h$

Corollary 6.5. *For the process P, the mean and variance of the number of hops are:*

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

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 $h \geq \frac{x}{\mathcal{R}} = h_{opt}$ we get that

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

Theorem 6.6. *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: When for $p = 0.5$ we have $m = 2$ and the efficiency ratio is 3 i.e. the overestimate is 3 times the optimal number of hops.

7. EXPERIMENTAL EVALUATION

In this section we report on two sets of experiments that aim to validate the theoretical analysis of the previous section. We have implemented the two protocols using C++ and the data types for two-dimensional geometry of LEDA [14]. Each experiment was conducted for more than 10,000 times in order to achieve good average results.

In the first set of experiments, we investigate (a) the impact of the angle α and (b) the number of targets found during the *search phase*, on the hops efficiency of the protocol when considering the ideal case where the search phase always finds a particle (of sufficiently high battery) in $(-\alpha, \alpha)$ (we call the measured efficiency, the *ideal hops efficiency*). We observe that for both protocols, as $\alpha \to 0$, the ideal hops efficiency $C_h \to 1$. Actually, the ideal C_h initially decreases very fast with increasing α, while having a limiting behavior of no further significant improvement when $\alpha \leq 40$. Figure 5 shows the effect of finding more than one target during the search phase; as the number of targets increases, the ideal hops efficiency $C_h \to 1$. We note a similar threshold behavior, for a total number of 4 targets.

In order to study the performance of the two protocols in more realistic cases, we generate a variety of sensor fields

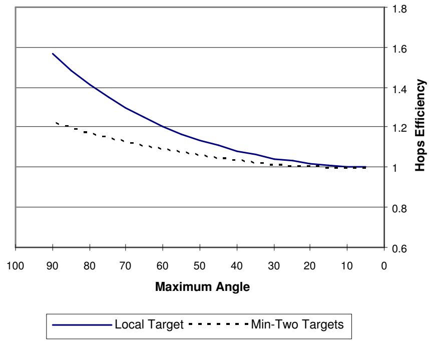

Figure 4: Ideal Hops Efficiency for angles $\alpha \in [5, 90]$

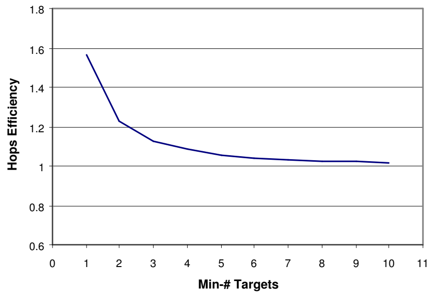

Figure 5: Ideal Hops Efficiency for different number of targets

in a 100m by 100m square. In these fields, we drop $n \in [100, 5000]$ particles uniformly distributed on the smart-dust plane, i.e. $0.01 \leq d \leq 0.5$. Each smart dust particle has a radio range of $\mathcal{R} = 5m$. Finally, the particle p selected by the adversary $\mathcal{A}$ is always located at position $(x, y) = (0, 50)$ and the wall is located at $x = 100$. In this set of experiments, the particle p' discovered in the search phase can be located anywhere within the cyclic sector defined by circles of radiuses $0, \mathcal{R}$ and angles $(-\alpha, \alpha)$. Note that this experimental setup is based on that used in [7, 10, 13]. Also, remark that the efficiency is *measured over the successful tries*, i.e. we do not take into account those runs that *backtracked*, however we keep track of the total number of times that the protocol was required to backtrack.

In figure 6 we observe that opposed to the ideal case (i.e. when the search phase always returns a particle on the semicircle), we do not get significant improvement in the hops efficiency as the angle α is reduced. This is basically because the discovered particle p' might be close to p and thus the local improvement made is of limited significance. Note that the min-two uniform targets protocol achieves better efficiency compared to the local target protocol.

Figure 7 depicts the effect of density d on the hops efficiency of the two protocols. Interestingly, we observe that

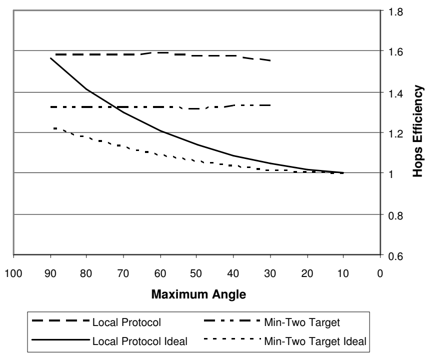

Figure 6: Hops Efficiency for angles $\alpha \in [5, 90]$ for $d = 0.3$

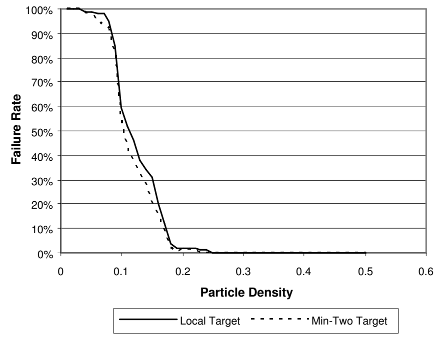

Figure 8: Failure rate for density $d \in [0.01, 0.5]$ and $\alpha = 90$

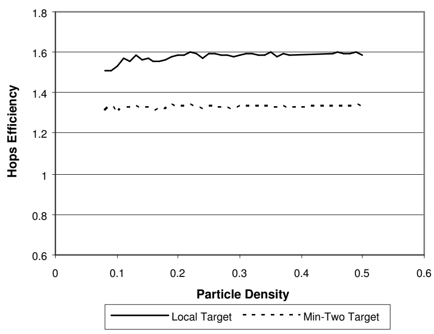

Figure 7: Hops Efficiency for density $d \in [0.01, 0.5]$ and $\alpha = 90$

even for quite low density of particles (i.e. $d \leq 0.2$) the hops efficiency remains unaffected. This is a result of our choice not to include the failed searches in our measurements, that is, the measurements include only the search phases that resulted in finding a particle p closer to $\mathcal{W}$. To get a more complete view on the effect of density, figure 8 shows the failure rate (i.e. the number of times that the protocols backtracked) for different values of d. We observe that for low density (i.e. $d \leq 0.1$) both protocols almost always use the backtrack mechanism, while when $d \geq 0.2$ the failure rate drops very fast to zero. This can be justified by taking into account the average degree of each particle for various density d.

8. CONCLUSIONS AND FUTURE WORK

We presented here a model for Smart Dust and two basic protocols (and their average case analysis) for local detection and propagation. We plan to investigate protocols that trade-off hops efficiency and time, as well as study the fault-tolerance of protocols as a function of smart dust parameters (such as density of the cloud, the energy saving characteristics etc.). Also, we are currently working towards the design of local protocols than can monitor the spreading of a time-sequence of events (i.e. tracking protocols). Finally, we plan to provide performance comparisons with other protocols mentioned in the related work section.

9. REFERENCES

[1] Berkeley Wireless Research Center: `http://bwrc.eecs.berkeley.edu` (2001).

[2] I. Chatzigiannakis, S. Nikoletseas, and P. Spirakis: An Efficient Communication Strategy for Ad-hoc Mobile Networks. In *Proc. 15th International Symposium on Distributed Computing* – DISC'2001. Also Brief Announcement in *20th Annual Symposium on Principles of Distributed Computing* – PODC'2001.

[3] I. Chatzigiannakis, S. Nikoletseas and P. Spirakis: On the Average and Worst-case Efficiency of Some New Distributed Communication and Control Algorithms for Ad-hoc Mobile Networks. Invited paper, in *Proc. 1st ACM Workshop on Principles of Mobile Computing* – POMC'2001.

[4] D. Estrin, R. Govindan, J. Heidemann and S. Kumar: Next Century Challenges: Scalable Coordination in Sensor Networks. In *Proc. 5th ACM/IEEE International Conference on Mobile Computing* – MOBICOM'1999.

[5] W. Feller. An Introduction to Probability Theory and its Applications. *John Wiley*, New York, 1957.

[6] S.E.A. Hollar: COTS Dust. Msc. Thesis in Engineering-Mechanical Engineering, University of California, Berkeley, USA, 2000.

[7] W. R. Heinzelman, A. Chandrakasan and H. Balakrishnan: Energy-Efficient Communication Protocol for Wireless Microsensor Networks. In *Proc. 33rd Hawaii International Conference on System Sciences* – HICSS'2000.

[8] W. R. Heinzelman, J. Kulik and H. Balakrishnan: Adaptive Protocols for Information Dissemination in Wireless Sensor Networks. In *Proc. 5th ACM/IEEE International Conference on Mobile Computing* – MOBICOM'1999.

[9] C. Intanagonwiwat, R. Govindan and D. Estrin: Directed Diffusion: A Scalable and Robust Communication Paradigm for Sensor Networks. In *Proc. 6th ACM/IEEE International Conference on Mobile Computing* – MOBICOM'2000.

[10] 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.

[11] B. Karp: Geographic Routing for Wireless Networks. Ph.D. Dissertation, Harvard University, Cambridge, USA, 2000.
[12] μ-Adaptive Multi-domain Power aware Sensors: `http://www-mtl.mit.edu/research/icsystems/uamps`, April, 2001.
[13] A. Manjeshwar and D.P. Agrawal: TEEN: A Routing Protocol for Enhanced Efficiency in Wireless Sensor Networks. In *Proc. 2nd International Workshop on Parallel and Distributed Computing Issues in Wireless Networks and Mobile Computing (PDCIWNMC'02), 16th Annual International Parallel & Distributed Processing Symposium* –IPDPS'2002.
[14] K. Mehlhorn and S. Näher: LEDA: A Platform for Combinatorial and Geometric Computing. *Cambridge University Press*, 1999.
[15] C.E. Perkins: Ad Hoc Networking. *Addison-Wesley*, Boston, USA, January, 2001.
[16] S. M. Ross: Stochastic Processes, 2nd Edition. *John Wiley and Sons, Inc.*, 1995.
[17] S. Singh and C.S. Raghavendra: PAMAS: Power aware multi-access protocol with signalling for ad hoc networks. In *ACM Computer Communication Review*, vol. 28, no. 3, pp. 5-25, July 1998.
[18] Wireless Integrated Sensor Networks: `http:/www.janet.ucla.edu/WINS/`, April, 2001.
[19] W. Ye, J. Heidemann and D. Estrin: An Energy-Efficient MAC Protocol for Wireless Sensor Networks. In *Proc. 12th IEEE International Conference on Computer Networks* – INFOCOM'2002.

APPENDIX

A. TECHNOLOGICAL SPECIFICATIONS OF SMART DUST DEVICES

Each smart dust device has a size on the order of a cubic inch which contains the four basic subsystems: power, computation, sensors and communication.

One of the tightest constraints on the circuit design is the operating voltage. Since the device operates using batteries, and because the RF chip sets in the design require a very narrow operating range centered at `3 Volts`, [6] focuses on `3V` portable power sources (in particular Lithium batteries that operate at the desired voltage, with `285 mWatt-hrs/gram` energy density and `10mA` max current at `10.5 grams`).

To operate for long periods of time, the grain particles are designed to enter sleep mode in order to save power. During this period, virtually all microprocessors on the market can go into low power modes that consume about 1μW (on average).

Taking into consideration the computational needs of the system, [6] uses an `8 bit` microcontroller (Amter AVR AT90 LS8538) that provides a minimum of `1.25 MIPS/mA`, at an operating voltage of `2.7-5.5V` with `4K` flash memory. An alternative to this is the `16/32 bit` Strong Thumb processor (Amter AT91M404000) that provides `0.6 MIPS/mA`, but operates at `2.7-3.6V` and offers external memory capabilities.

Sensors provide the raw data the grain particles need to gain a high level understanding of the environment. The sensors themselves can be classified into two groups: (1) weather monitors, and (2) motion detectors. [6] uses a magnetometer, accelerometer, light sensor, temperature sensor, pressure sensor and humidity sensor. All sensors consume 200-650μA and operate at a voltage range of `2.7-9V`.

Finally, regarding the communication subsystem, we will report three different technologies. The TR1000 RF Monolithis transceiver is designed for short-range wireless data applications, consumes an average of `18 mWatts` of power, and provides bit rates up to `19.2kbps` at distances of `5-30m`. The Bluetooth compliant Philstar PH2401 transceiver chip consumes less than `20mA` of power with throughput of `1 Mbit/s` at distances of up to `10m`. Finally, an alternative to RF communication is the one-way asynchronous laser communication which consumes `50mA` of power and provides bit rate up to `5kbps` at distances of `21.4km`.

It is worth noting that currently, a number of research institutions in the U.S. are working on centimeter-scale distributed sensor networks [1, 18].