Surveillance coverage of sensor networks under a random mobility strategy*

George Kesidis
EE & CSE Depts, Penn State University
University Park, PA, USA
kesidis@engr.psu.edu

Takis Konstantopoulos
ECE Dept, University of Texas
Austin, TX, USA
takis@ece.utexas.edu

Shashi Phoha
ARL, Penn State University
University Park, PA, USA
sxp26@psu.edu

Abstract

We consider the problem of surveillance of a region undertaken by a group of mobile sensors. Random mobility strategies are discussed in terms of coverage efficiency, communication and reliability in hostile environments. Under a Brownian motion random mobility strategy for the sensor grid, the distribution of the time-until-detection of slowly moving (point) targets is studied. Both two and three dimensional environments are considered. We obtain explicit formulas in three dimensions and bounds in two.

INTRODUCTION

Sensor network applications range from passive surveillance (search) to simultaneous tracking of multiple targets and distributed identification of their co-ordinated movement. The targets will have specific space-time neighborhoods (from stationary point-targets to large, rapidly moving ones). Example targets range from individual humans roaming through a crowd to platoons of tanks in a desert war theatre to point explosions to activation sites of chemical or biological weapons. A wide variety of sensors could also be used including image (video, 3-dimensional, infrared, omnidirectional), acoustic microphones, seismic or radiation meters, RADAR and LADAR. The sensor nodes themselves could be fixed (to, e.g., monitor an airport or power plant) or could form a mobile expeditionary sensor grid.

Sensor nodes may have limited amounts of available energy stored in batteries which may require energy efficient routing and communication medium access control. The environment in which the sensors and their targets operate could range from land (desert, jungle, mountains), sea and air. Furthermore, environmental conditions could include high levels of noise, low levels of ambient light, cloud cover and other occlusions, and enemy activity that can compromise target search and tracking and network communication.

Individual sensors may need to co-ordinate in a distributed fashion to maintain communication connectivity [8, 13, 5, 1], target contact, and surveillance coverage to prespecified degrees of confidence. These are often conflicting goals requiring basic trade-offs.

The focus of this paper is to study target detection performance for a sensor network under random mobility. This paper is organized as follows. We first define notation for our subsequent problem formulation and give a brief overview of mobility strategies in target-search mode. Deterministic and random search strategies are compared. Next, we described the problem of detection of slowly moving (point) targets in two and three dimensional environments, again under a random mobility mechanism for the sensor grid. For a two-dimensional planar environment, we give a bound on the tail of the distribution of time-until-detection of a point target. A more precise calculation is given for a three-dimensional environment. Our problem formulation allows us to consider a large number of sensors operating in a large region. Finally, design issues pertaining to the single parameter of mobility, the variance σ^2, are discussed.

TARGET SEARCH STRATEGIES

Consider a contiguous region A in the plane. Extensions to three-dimensions (undersea, air or space environments) are straight-forward. Let N be the number of mobile sensors in A that form the nodes of our network. The position of the node i at time t is denoted by $X_i(t)$. Each node is assumed to have a maximum radius of surveillance r_s.

At time t, the total surveillance coverage of the N nodes is the intersection of the set

$$\bigcup_{i=1}^{N} B(X_i(t), r_s)$$

with the region A. Here,

$$B(x,r) = \{y : |y-x| \le r\}$$

denotes a closed disk of radius r and center x. The distance between two points $x, y \in A$ is the Euclidean distance $|x-y| = \sqrt{\sum_j (x_j - y_j)^2}$.

We assume in the following that the maximum coverage area $N\pi r_s^2$ (maximum assuming no overlapping disks) is significantly less than the area $|A|$ of A. Thus, mobility of the nodes is necessary to maintain surveillance coverage of the whole region A to some degree of confidence as described below.

*This research was funded by the Defense Advanced Research Projects Agency (DARPA) and administered by the Army Research Office (ARO) under ESP MURI Award No. DAAD19-01-1-0504. Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of the DARPA or the ARO. The first author's research was supported in part by NSF grant ITR-0312558 and second author's research was supported in part by NSF grant ANI-9903495.

Let λ be the mean density of nodes in the given region, i.e., $\lambda = N/|A|$ or $\lambda = EN/|A|$ if N is random.

The objective of surveillance is to gather information about certain target objects that move through the region A. These objects exist in a so-called time-space neighborhood characterized by a time-varying function $\Omega(t)$ for times t in some finite interval of (continuous) time $[t_0, t_1]$ where Ω takes values in the subsets of A. For convenience of notation, we can define $\Omega(t) = \emptyset$ for $t \notin [t_0, t_1]$. In the following, we simply assume that if

$$B(X_i(t), r_s) \cap \Omega_j(t) \neq \emptyset$$

then sensor node i $\mathrm{detects}$ object j at time t.

Certain targets may be $\mathrm{stationary}$ in the sense that, for a fixed x and r, $\Omega(t) = B(x, r) \subset A$ for all times $t \in [t_0, t_1]$ where x is the center (ground zero) of, for example, an explosion and $t_1 - t_0$ is its duration. An expanding object's detection region grows as a function of time, e.g., a biological weapon whose agents diffuse through the air; i.e., $\Omega(t) = B(x, r(t))$ for some increasing function $r(\cdot)$. Finally, a single moving object could be characterized by a (one dimensional) path in A, $x(t)$ for $t \in [t_0, t_1]$ as well as a radius r, i.e., $\Omega(t) = B(x(t), r)$.

We obviously assume that no sensor node is $\mathrm{a\ priori}$ aware of any such attributes of a target object. In [7], a data fusion framework for sensor data is given and, in particular, a mechanism is described for estimation of object velocity involving a triangulation between two nodes sensing the object.

Problems in surveillance coverage

Mobility strategies have two modes of operation depending on the presence (data acquisition, target tracking) or absence (search) of a $\mathrm{detected}$ target object(s). In search mode, mobility strategies can be roughly classified into two groups: random and deterministic. In deterministic strategies, region A could be partitioned a priori into N pieces and a sensor node assigned to each piece that becomes the home area of the node. The node sweeps its home area in, e.g., a deterministic zigzag pattern of a lawn-mower, searching for targets. Under random search strategies, each node moves at random or "diffuses" through A. Random mobility can be made somewhat more efficient by adopting strategies wherein nodes locally "repel" each other and are then less likely to visit areas very recently visited. Alternatively, proximal nodes can dynamically elect a "cluster head" which remains active while the other proximal nodes "sleep" to conserve energy and reduce communication congestion in the cluster neighborhood [15, 4, 12]. Clearly, hybrid and hierarchical mobility patterns could also be devised, e.g., a coarser partition of A in which each element of the partition has more than one node and each node employs a random sweep pattern within its (now larger and shared) home area.

In data-acquisition mode, the first node to sense an object could be deemed the "co-ordinator" for that object and the alarm to the rest of the network. Certain nodes would then deterministically drift toward the target. In the following, we focus on sensor networks only in search mode.

Robustness of mobility strategies

A sensor node may fail because of electro-mechanical problems, unforeseen obstructions or hazards (including traps placed by the enemy), exhaustion of its power supply, or by overt enemy activity leading to destruction or capturing/hijacking of the node. When a node fails in the deterministic mobility setting, the remaining nodes need to quickly become aware of the failure (note that this assumption would require a nontrivial intrusion detection system). An advantage of random mobility strategies is that they require minimal coordination (communication overhead) especially when a node fails and experience graceful degradation in surveillance coverage in that event. On the other hand, random mobility strategies will, in general, have poorer target detection performance.

Sensor networks with uncontrolled random mobility

Recently, networks of extremely small sensors have been proposed, in particular the "smart dust" project [9, 10]. In such networks, control of sensor mobility may highly limited by energy supply and environmental conditions, an example of the latter being turbulent air. For such networks, mobility may be implicitly random, i.e., mobility may be modeled with diffusion dynamics together with a drift imparted at the time of deployment.

POINT TARGET DETECTION UNDER BROWNIAN MOVEMENT

Consider an arbitrary point in $\mathbb{R}^2$ or $\mathbb{R}^3$, representing a stationary point target, taken to be the origin and taken to come into existence at time zero (both without loss of generality since our sensor mobility models are spatial and temporal translation invariant). In this section, we study the distribution of time required by a sensor grid, under a Brownian mobility strategy, to detect the object. Our problem formulation allows us to consider a large number of sensors operating in a large region.

We assume the initial positions of the sensors $\mathbf{N}(0) = \{X_n(0)\}_{n=1}^{\infty}$ are distributed as a Poisson point process with intensity λ nodes/m^2 or nodes/m^3 (depending on whether the network is two- or three-dimensional). Recall, at this point, the notion of a $\mathrm{Boolean\ Model}$ [14]. It is constructed by drawing i.i.d. copies of some random set (e.g., a disk or a ball) that are known as grains and by translating each copy at the point of a Poisson process; these points are known as germs and are taken to be independent of the grains.

Below we model the network as a $\mathrm{Brownian\ Boolean\ Model}$ (BBM) whereby the nodes execute independent Brownian motions, $X_n(t)$, $n = 1, 2, \ldots$. The assumption of i.i.d. node displacements means that the sensor positions $\mathbf{N}(t)$ form a Poisson process for all $t \geq 0$ [6].

The three-dimensional environment

In this section, each stochastic process

$$X_n(t) = (X_n^1(t), X_n^2(t), X_n^3(t))$$

is a 3-dimensional Brownian motion with independent Cartesian coordinates and variance coefficient σ^2 (measured in m^2/s). Without loss of generality, suppose that a target is located at the origin. Let $B(0,r)$ be the ball of radius $r \equiv r_s$ (the surveillance radius) centered at the origin. At each point of time, the area that can be monitored by the sensor network is represented by the set

$$Z(t) = \bigcup_n (X_n(t) \oplus B(0,r)) = \left(\bigcup_n X_n(t)\right) \oplus B(0,r),$$

where $A \oplus B$ denotes Minkowski addition of two sets $A, B \subset \mathbb{R}^3$:

$$A \oplus B = \{a + b : a \in A, b \in B\}.$$

Thus, at each t, the random set $Z(t)$ is a Boolean model. We are interested in the time required for the target detection:

$$T \equiv \min\{t \geq 0 : 0 \in Z(t)\}.$$

Define

$$\mathcal{Z}(t) = \bigcup_{s \in [0,t]} Z(s).$$

We have

$$\mathsf{P}(T > t) = \mathsf{P}(0 \notin Z(s) \text{ for all } s \in [0,t]) = \mathsf{P}(0 \notin \mathcal{Z}(t)).$$

Now note that, for each t, $\mathcal{Z}(t)$ is itself a Boolean model. Indeed, we may take as centers (or germs) the initial Poisson points, and as sets around them (or grains) i.i.d. copies of the set

$$\mathcal{W}(t) = \bigcup_{s \in [0,t]} X(s) \oplus B(0,r),$$

i.e., a ball carried around by the trajectory of a single Brownian motion X between 0 and t that is a so-called Wiener sausage. Hence, $\mathsf{P}(T > t)$ is actually the volume fraction of the Boolean model with Wiener sausage grains [14]:

$$\mathsf{P}(T > t) = \exp(-\lambda \mathsf{E}\,\mathrm{vol}(\mathcal{W}(t))) \tag{1}$$

where T is the target hitting time of a single grain. Clearly, by Fubini's theorem,

$$\begin{aligned} \mathsf{E}\,\mathrm{vol}(\mathcal{W}(t)) &= \mathsf{E}\int_{\mathbb{R}^3} \mathbb{1}\{x \in \mathcal{W}(t)\}\mathrm{d}x \\ &= \int_{\mathbb{R}^3} \mathsf{P}(x \in \mathcal{W}(t))\mathrm{d}x. \end{aligned}$$

Note that

$$\begin{aligned} x \in \mathcal{W}(t) &\iff \exists s \in [0,t],\ \exists y \in B(0,r) :\ x = X(s) + y \\ &\iff \exists s \in [0,t] :\ x - X(s) \in B(0,r) \\ &\iff \exists s \in [0,t] :\ |x - \sigma W(s)| \leq r \\ &\iff \exists s \in [0,t] :\ |(x/\sigma) - W(s)| \leq r/\sigma, \end{aligned}$$

where W is a standard Brownian motion (variance 1) with $W(0) = 0$.

Letting H_r^x denote the first hitting time of the ball $B(0,r)$ by a standard Brownian motion W started at $x \in \mathbb{R}^3$, we therefore have

$$x \in \mathcal{W}(t) \iff H_{r/\sigma}^{x/\sigma} \leq t.$$

Hence,

$$\mathsf{E}\,\mathrm{vol}(\mathcal{W}(t)) = \int_{\mathbb{R}^3} \mathsf{P}(H_{r/\sigma}^{x/\sigma} \leq t)\mathrm{d}x.$$

Now consider again the Bessel process $R(t) = |x + W(t)|$ for a standard Brownian motion started at x. It is known that R is a one-dimensional diffusion with $R(0) = |x| = \rho$. Define $\hat{H}_r^\rho = \min\{t \geq 0 : R_t \leq r\}$ and note that since there is isotropy (rotational invariance), the hitting time depends on the initial position only through its magnitude, i.e., $\hat{H}_r^\rho = H_r^x$ for all x such that $\rho = |x|$. Thus,

$$\mathsf{E}\,\mathrm{vol}(\mathcal{W}(t)) = \int_0^\infty 4\pi\rho^2 \mathsf{P}(\hat{H}_{r/\sigma}^{\rho/\sigma} \leq t)\mathrm{d}\rho. \tag{2}$$

The random variable $\hat{H}_r^\rho$ has a well-known distribution [2]:

$$P(\hat{H}_r^\rho \leq t) = \begin{cases} 1 & \rho \leq r \\ \int_0^t \frac{r(\rho - r)}{\rho\sqrt{2\pi}s^{3/2}} \exp\left(-\frac{(\rho-r)^2}{2s}\right) \mathrm{d}s & \rho \geq r. \end{cases} \tag{3}$$

We see that

$$P(\hat{H}_{r/\sigma}^{\rho/\sigma} \leq t) = P(\hat{H}_r^\rho \leq \sigma^2 t).$$

Hence we may take $\sigma = 1$, to alleviate notation, and in the end replace t by $\sigma^2 t$. With this in mind, and using the expression above, we have

$$\begin{aligned} \mathsf{E}\,\mathrm{vol}(\mathcal{W}(t)) = \int_0^r 4\pi\rho^2\mathrm{d}\rho + \\ \int_r^\infty 4\pi\rho^2 \left[\int_0^t \frac{r(\rho-r)}{\rho\sqrt{\frac{2}{\pi}}s^{3/2}} \exp\left(-\frac{(\rho-r)^2}{2s}\right)\mathrm{d}s\right]\mathrm{d}\rho. \end{aligned}$$

The first integral is clearly $(4/3)\pi r^3$, the volume of the ball $B(0,r)$. For the second integral we use Fubini and a change of variables (setting $\rho - r = x$) to get that the second integral equals

$$\begin{aligned} &\sqrt{8\pi}r \int_0^t s^{-3/2}\left[\int_r^\infty \rho(\rho - r)\exp\left(-\frac{(\rho-r)^2}{2s}\right)\mathrm{d}\rho\right]\mathrm{d}s \\ &= \sqrt{8\pi}r\int_0^t s^{-3/2}\left[\int_0^\infty x(x+r)\exp(-x^2/2s)\mathrm{d}x\right]\mathrm{d}s \\ &= \sqrt{8\pi}r\int_0^t s^{-3/2} I(s)\mathrm{d}s. \end{aligned}$$

Let

$$\varphi_s(x) = \frac{1}{\sqrt{2\pi s}}\exp(-x^2/2s)$$

be the density of $\sqrt{s}Y$, where Y is a standard normal distributed random variable. Thus,

$$\begin{aligned}
I(s) &= \int_0^\infty x(x+r)\exp(-x^2/2s)\mathrm{d}x \\
&= \sqrt{2\pi s}\int_{-\infty}^\infty x(x+r)\mathbf{1}\{x\geq 0\}\varphi_s(x)\,\mathrm{d}x \\
&= \sqrt{2\pi s}\mathsf{E}[\sqrt{s}Y(\sqrt{s}Y+r)\mathbf{1}\{Y\geq 0\}] \\
&= \sqrt{2\pi s}[s\mathsf{E}(Y^2\mathbf{1}\{Y\geq 0\})+r\sqrt{s}E(Y\mathbf{1}\{Y\geq 0\})] \\
&= \sqrt{2\pi s}\left(\frac{s}{2}+r\sqrt{\frac{s}{2\pi}}\right) \\
&= s^{3/2}\sqrt{\pi/2}+rs.
\end{aligned}$$

Collecting the calculations together, we get

$$\begin{aligned}
\mathsf{E}\,\mathrm{vol}(\mathcal{W}(t)) &= \frac{4\pi}{3}r^3+\sqrt{8\pi}r\int_0^t s^{-3/2}(\sqrt{\frac{\pi}{2}}s^{3/2}+rs)\mathrm{d}s \\
&= \frac{4\pi}{3}r^3+4\pi rt+4r^2\sqrt{2\pi t}.
\end{aligned}$$

Now recall that the actual variance is σ^2 (hence replace t by $\sigma^2 t$ in the last expression) and finally substitute in the expression (1) to get:

$$\mathsf{P}(T>t) = \exp\left(-\left(\frac{4\pi}{3}\lambda r^3+4\pi\lambda r\sigma^2 t+4\sqrt{2\pi}\lambda r^2\sigma\sqrt{t}\right)\right)$$

This is a distribution with a Weibull-type tail.

The two-dimensional environment

In a two-dimensional (planar) environment A, the expression for the distribution of the hitting time $\hat{H}$ given in (3) above is not available. We can, however, use the Chernoff upper bound:

$$\begin{aligned}
\mathsf{P}(\hat{H}_r^\rho\leq t) &= 1-\mathsf{P}(\hat{H}_r^\rho>t) \\
&\geq 1-\exp\left(-\sup_\alpha\{\alpha t-\mathsf{E}\exp(\alpha\hat{H}_r^\rho)\}\right)
\end{aligned}$$

where the expression for $\mathsf{E}\exp(\alpha\hat{H}_r^\rho)$ is given as a ratio of modified Bessel functions in 2.0.1, p. 297, of [2]. This lower bound could then be substituted into (2) to obtain a lower bound on $\mathsf{E}\,\mathrm{vol}(\mathcal{W}(t))$ and thus obtain an upper bound on the tail of the target detection time distribution (1).

Random mobility design

A design objective here could be as follows: Suppose that a point target is required to be detected within a prespecified time t with an error probability less than $e^{-\beta}$, where β is a positive number.

Our design variable is the variance σ^2 of the node mobility. When r is very small, the target detection goal can be achieved if we choose σ^2 roughly larger than $\beta/(4\pi\lambda rt)$:

$$\sigma^2 \gtrsim \frac{\beta}{4\pi\lambda rt}.$$

This is found by solving a quadratic inequality and taking asymptotics when r becomes small. We omit the details of the computation.

We next compute the expectation of the target hitting time T as follows. First rewrite

$$\mathsf{P}(T>t)=Ce^{-at-b\sqrt{t}},$$

with the obvious choice of the constants. The integration $\mathsf{E}T=\int_0^\infty \mathsf{P}(T>t)\mathrm{d}t$ can be performed to give

$$\mathsf{E}T=\frac{C}{a}\left(1-\sqrt{\pi}\frac{b}{2\sqrt{a}}e^{b^2/4a}\bar{\Phi}(b/2\sqrt{a})\right)$$

where $\bar{\Phi}(x)$ is the tail of the standard normal distribution function. Using the inequality

$$\bar{\Phi}(x)\leq\frac{e^{-x^2/2}}{x\sqrt{2\pi}}$$

bounds can be obtained. Finally, by substituting the values of the constants, we obtain an exact expression:

$$\mathsf{E}T=\frac{1}{\sigma^2}e^{-\frac{4\pi}{3}\lambda r^3}\left(\frac{1}{4\pi\lambda r}-\sqrt{\frac{r}{8\pi\lambda}}e^{2\lambda r^3}\bar{\Phi}\left(\sqrt{2\lambda r^3}\right)\right).$$

This is a rapidly decreasing function of λ (exponential decay).

SUMMARY

We considered the problem of surveillance of a potentially large region undertaken by a potentially large group of mobile sensors. Under a random mobility strategy for the sensor grid, the distribution of the contact time between two nodes and the distribution of the time-until-detection of slowly moving (point) targets were studied. Both two and three dimensional environments were considered. Finally, design issues pertaining to the single parameter of mobility, the variance σ^2, were discussed.

REFERENCES

[1] Bettstetter, C. On the minimum node degree and connectivity of a wireless multihop network. Proc. ACM MobiHOC, Lausanne, Switzerland, June 2002.

[2] Borodin, A.N. and Salminen, P. Handbook of Brownian Motion–Facts and Formulae. Birkhäuser, Boston, 1996.

[3] J. Canny. Complexity of Robot Motion Planning. MIT Press, Cambridge, MA, June 1988.

[4] C. Chevally, R.E. Van Dyck and T.A. Hall, Self-organization for Wireless Sensor Networks. In Proc. CISS, Princeton, March 2002.

[5] Dowell, L.J. and Bruno, M.L. Connectivity of random graphs in mobile sensor networks: validation of Monte Carlo simulation results. Proc. ACM RBAC, Berlin, 2000.

[6] Frey, A. and Schmidt, V. Marked point processes in the plane I. Adv. Performance Anal. 1, 65-110, 1998.

[7] Friedlander, D. and Phoha, S. Semantic information fusion for coordinated signal processing in mobile

sensor networks. Intern. J. High Performance Computing Appl. Special issue on sensor networks, 16, August 2002.
[8] Gupta, P. and Kumar, P.R. Critical power connectivity in wireless networks. In: Stochastic Analysis, Control, Optimization, and Applications: A Volume in Honor of W.H. Fleming, W.M. McEneany, G. Yin, and Q. Zhang (eds.) Birkhäuser, Boston, 547-556, 1998.
[9] Kahn, J.M., Katz, R.H. and Pister, K.S.J. Mobile networking for smart dust. In Proc. ACM/IEEE Int. Conf. on Mobile Computing and Networking, Seattle, 1999.
[10] Kahn, J.M., Katz, R.H. and Pister, K.S.J. Emerging Challenges: mobile networking for smart dust. Journal of Communications and Networks, Vol. 2, No. 3, 2000.
[11] G. Kesidis, R. Rao and S. Phoha, "Mobility management of ad-hoc sensor networks using distributed annealing", CSE Dept, Penn State University, Technical Report, CSE 03-017, Sept. 29, 2003.
[12] Krishnamachari, B., Wicker, S.B., and Bejar, R. Phase transition phenomena in wireless ad hoc networks. Proc. IEEE Symposium on Ad Hoc Wireless Networks, in conjunction with IEEE GLOBECOM, San Antonio, TX, November 2001.
[13] Q. Li, R. Peterson, M. DeRosa, and D. Rus. Reactive Behavior in Self-reconfiguring Sensor Networks. in Proc. ACM MOBICOM'02, Atlanta, 2002.
[14] Santi, P., Blough, D.M., and Vainstein, F. A probabilistic analysis for the range assignment problem in ad hoc networks. Proc. ACM MobiHOC, Long Beach, CA, 2001.
[15] Stoyan, D., Kendall, W.S., and Mecke, J. Stochastic Geometry and Its Applications. Wiley, New York, 1996.
[16] L. Subramanian and R.H. Katz. An architecture for building self-configurable systems. In Proc. IEEE MobiHoc, 2000.