

Asymptotics of Query Strategies over a Sensor Network

Sanjay Shakkottai*

February 11, 2004

Abstract

We consider the problem of a user querying for information over a sensor network, where the user does not have prior knowledge of the location of the information. We consider three information query strategies: (i) a Source-only search, where the source (user) tries to locate the destination by initiating query which propagates as a continuous time random walk (Brownian motion); (ii) a Source and Receiver Driven “Sticky” Search, where both the source and the destination send a query or an advertisement, and these leave a “sticky” trail to aid in locating the destination; and (iii) where the destination information is spatially cached (i.e., repeated over space), and the source tries to locate any one of the caches. After a random interval of time with average t , if the information is not located, the query times-out, and the search is unsuccessful.

For a source-only search, we show that the probability that a query is unsuccessful decays as $(\log(t))^{-1}$. When both the source and the destination send queries or advertisements, we show that the probability that a query is unsuccessful decays as $t^{-5/8}$. Further, faster polynomial decay rates can be achieved by using a finite number of queries or advertisements. Finally, when a spatially periodic cache is employed, we show that the probability that a query is unsuccessful decays no faster than t^{-1} . Thus, we can match the decay rates of the source and the destination driven search with that of a spatial caching strategy by using an appropriate number of queries.

The sticky search as well as caching utilize memory that is *spatially distributed* over the network. We show that *spreading the memory* over space leads to a decrease in memory requirement, while maintaining a polynomial decay in the query failure probability. In particular, we show that the memory requirement for spatial caching is larger (in an order sense) than that for sticky searches. This indicates that the appropriate strategy for querying over large sensor networks with little infrastructure support would be to use multiple queries and advertisements using the sticky search strategy.

*This research was partially supported by NSF Grants ACI-0305644 and CNS-0325788. The author is with the Wireless Networking and Communications Group, Department of Electrical and Computer Engineering, The University of Texas at Austin, shakkott@ece.utexas.edu. A shorter version of this paper appears in the Proceedings of IEEE Infocom, Hong Kong, March 2004.

1 Introduction

With the availability of cheap wireless technology and the emergence of micro-sensors based on MEMS technology [10, 23], sensor networks are anticipated to be widely deployed in the near future. Such networks have many potential applications, both in the military domain (eg. robust communication infrastructure or sensing and physical intrusion detection), as well as commercial applications such as air or water quality sensing and control. Such networks are usually characterized by the absence of any large-scale established infrastructure, and nodes cooperate by relaying packets to ensure that the packets reach their respective destinations.

An important problem in sensor networks is that of querying for information. The query type could range from trying to determine the location of a particular node to querying for particular information. This problem has recently received a lot of attention [8, 3, 2, 21], and is also related to routing over such networks [6, 20, 9, 7].

We consider a problem where a querying node (the source/user) transmits a query for some information (the destination), which is located at a (normalized) distance of '1' from the source. However, we assume that the source has no prior knowledge of the location of the information (the destination), nor that they are separated by a distance of '1' (i.e., *no prior destination location information is available*). Further, we assume that the nodes *do not have any "direction" information*. In other words, nodes only know who their local neighbors are, but do not have their geographical position or direction information.

We consider three search/query strategies: **(i)** A Source-only search, where the source (i.e., the user) tries to locate the destination by initiating a query which propagates as a continuous time random walk (Brownian motion); **(ii)** a Source and Receiver Driven "Sticky" Search, where both the source and the destination send a query or an advertisement, and these leave "sticky" trails to aid in locating the destination; and **(iii)** where the destination information is spatially cached (i.e., repeated periodically over space), and the source tries to locate any one of the caches.

As an aside, we note that if partial destination location information was available, one could design strategies which explicitly use this knowledge. For instance, if it was known a-priori that the destination was at a distance '1' from the source, constrained flooding would be a candidate strategy for locating the position of the destination. Further, the availability of direction information (for

instance, with GPS equipped nodes) would enable strategies similar to the sticky search, but route along “straight-lines” in the network [2] instead of random walks.

In this paper, we study query strategies in the absence of location/direction information. We consider the search strategies based on random walks, and associate an (exponentially distributed) random time-out interval, after which a query ceases to propagate. Thus, if the destination is not found before the time-out, the query is unsuccessful. We derive the asymptotic behavior of the querying strategies discussed above, and discuss their relative performance.

1.1 Related Work

As discussed earlier, there has been much work on querying and routing with information constraints in sensor networks. There have been various protocols and algorithms that have been recently developed for these networks [8, 3, 2, 21, 6, 20, 9, 19, 12, 13]. These protocols span various types – from using geographical information [9, 12, 8, 6] to (limited) flooding [19, 21] to gossip [13, 19] – and use various combinations of the three schemes we have considered. In Section 2, we shall discuss [2] in detail.

We use Brownian motion based models for analysis. Related work includes [7], where the author studies optimal placement of limited routing information over a sensor network. The author shows that uniform placement over space will not be optimal. Instead, the routing information needs to be “concentrated” according to a specific pattern (a bicycle wheel spokes like pattern). Other related work includes [14], where the authors consider an object at origin and an infinitely large number of mobile nodes (initially placed according to a spatial Poisson process) looking for the object. They show that the probability that the object is not located decays exponentially in time. For the schemes we have described in Section 1, there does not seem to be an analytical comparison in literature. In this paper, we derive the asymptotic performance of these schemes, and discuss their trade-offs.

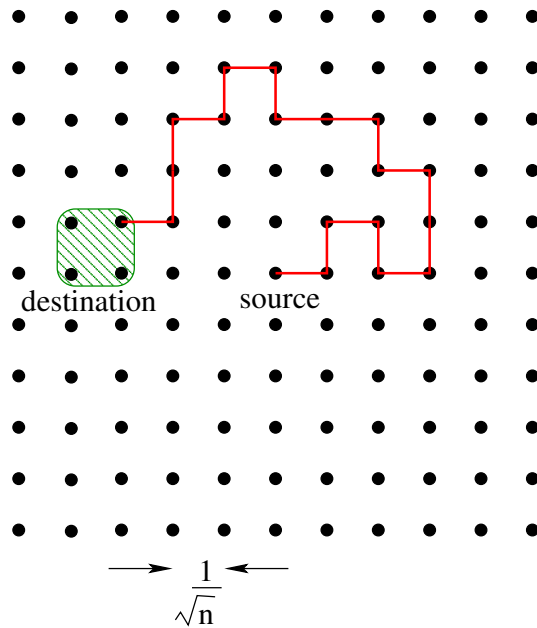


Figure 1: A random walk based search on a sensor grid

2 Querying Models for a Sensor Network

Let us consider a regular sensor grid network [22], with n sensor nodes per unit area as shown in Figure 1. For such a model, nodes are spaced regularly with the inter-node distance being $1/\sqrt{n}$. Let us assume that each node can communicate a distance $1/\sqrt{n}$. In other words, the nodes can communicate with each of its adjacent neighbors.

We are interested in the situation where a node located at origin (also referred to as the *source node*) is interested in locating some particular information, which is located at some region or node on the grid, henceforth referred to as the *destination*. We assume that the destination is located at a (normalized) distance of '1' from the source node. However, we assume that the source node has *no prior knowledge* of the location of the destination. Thus, the distance normalization of '1' is used merely for comparison and computation of the performance of various strategies that we will consider, and this information is not available to the routing/query forwarding algorithms in the network.

We consider the following three query strategies for locating the destination.

(i) Source-only Search: In this strategy, the source transmits a query requesting the destination in-

formation. As the location of the destination is unknown, the query propagates through the network as a random walk. In other words, each intermediate node along the query path picks a neighbor at random (i.e., with equal probability of 0.25) and forwards the query to it (see Figure 1). The destination is located at a normalized distance of '1' from the source. We assume that all nodes that are a distance less than ϵ (for some fixed $0 < \epsilon < 1$) from the destination know the location of the destination, or equivalently, they form a spatial memory structure (for example, all the nodes in the shaded region in Figure 1 have the information, or possibly store the information in a distributed manner). Thus, if the query enters a region which is at a distance less than ϵ , the query is said to be successful. The destination can then pass the information back to the source node by various means depending on the infrastructure available (for example, using an addressing mechanism where the routing tables in intermediate nodes are updated by the source query).

Further, we assume that there is a time-out associated with the query. At each intermediate node along the path of the query, the node could do one of two things: it could decide to not forward the query with some positive probability, or it could decide to forward the query (to a random neighbor). Thus, for such a forwarding model, the query will terminate after a geometrically distributed number of hops (if the destination has not been found).

(ii) Source-Receiver “Sticky” Search: The second strategy we consider is a sticky search. Unlike the previous case where the source alone sends a query, we now consider the case where *both* the source and destination send “probes” into the network. As before, the source sends (one or more) queries. These queries propagate as a random walk, with a geometrically distributed time-out. However, in addition, the destination also sends probes into the network to advertise itself. These advertisements propagate over the network in a manner similar to that of a query, i.e., as a random walk with a geometrically distributed time-out. Henceforth, we will refer to queries and/or advertisements as *probes* that are transmitted by the source and/or destination.

In addition, the source and destination probes leave a “sticky” trail as they traverse the networks (see Figure 2). That is, each node in the network through which the source query passes through *remembers* that a query passed through it and is searching for some particular information. Similarly, each node in the network through which the destination advertisement passes through *remembers* that a probe passed through it advertising the presence of some information. Thus, if the source (destination) probe passes through a node through which the destination (source) previously passed

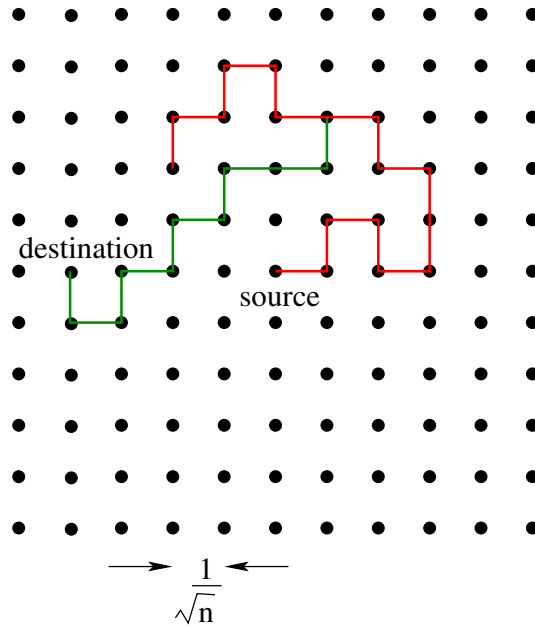


Figure 2: A source and destination based “sticky” random walk on a sensor grid

through, it can simply trace backward along the destination (source) probe’s path to reach the destination (source), see Figure 2. This scheme has been proposed by Braginsky and Estrin [2], and the authors develop the rumor routing protocol for sensor networks based on this idea.

(iii) Spatially-periodic Caching: Finally, the third strategy we consider is spatial caching. As in the previous cases, we assume source initiates a query which propagates as a random walk with a geometrically distributed time-out. However, in this case, we assume that the destination information is assumed to be periodically cached along a regular grid as shown in Figure 3. Thus, we assume that there is some infrastructure which allows the dispersion of the information to each of the caches. Further, we assume that each of the caches is of radius $0 < \epsilon < 1$. Thus, as in the source-only search strategy, *each* cache of radius ϵ could either be a distributed memory structure, or else one of the nodes in each cache could have the data, and let all the other nodes within the cache point toward it.

2.1 A Brownian Motion Based Model

In the previous section, we considered a random walk on a grid, with step size $1/\sqrt{n}$ and a geometrically distributed time-out. In this section, we will informally show that for n large enough, this model can be approximated by a two dimensional planar model, where the source and destination

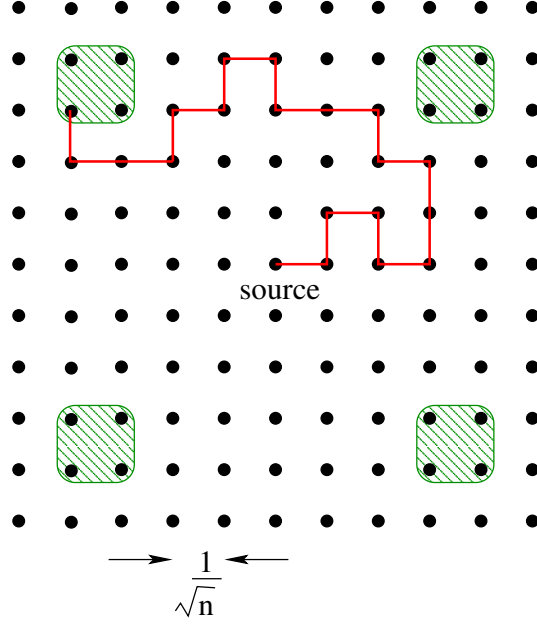


Figure 3: Periodic caching over a sensor grid

probes propagate as a Brownian motion with an exponentially distributed time-out. Thus, in the rest of the paper, we will consider Brownian motion based models for querying over a sensor network.

Let us rotate the coordinate axes by 45 degrees, and let $(X(i), Y(i))$ be the random variables which correspond to the discrete step taken at time i , with respect to the new coordinate system. Thus, $(X(i), Y(i))$ can take one of four values: $(1/\sqrt{2}, 1/\sqrt{2})$, $(-1/\sqrt{2}, 1/\sqrt{2})$, $(1/\sqrt{2}, -1/\sqrt{2})$, or $(-1/\sqrt{2}, -1/\sqrt{2})$, each with equal probability, depending on which direction the random walk propagates.

Thus the location of the random walk $(L_x(j), L_y(j))$ at time j is given by

$$L_x(j) = \frac{1}{\sqrt{n}} \sum_{i=0}^j X(i)$$

$$L_y(j) = \frac{1}{\sqrt{n}} \sum_{i=0}^j Y(i)$$

Now, let us define the continuous time processes

$$B_x^n(s) = \frac{1}{\sqrt{n}} \sum_{i=0}^{\lfloor 2ns \rfloor} X(i) \tag{1}$$

$$B_y^n(s) = \frac{1}{\sqrt{n}} \sum_{i=0}^{\lfloor 2ns \rfloor} Y(i). \tag{2}$$

Thus, at any continuous time s , $B_x^n(s)$ is the x-coordinate of the random walk after $\lfloor 2ns \rfloor$ time-steps (and correspondingly $B_y^n(s)$), and $\lfloor \cdot \rfloor$ is the integer floor function (i.e., the integer part of the argument). From standard theory for Brownian motion [11], we can show that the processes described in (1) and (2) converges (in a suitable sense) to a two dimensional Brownian motion $B(s) = (B_x(s), B_y(s))$ where $B_x(s)$ and $B_y(s)$ are independent one-dimensional Brownian motions. This follows from a central limit theorem type of an argument for functions¹.

In addition, we can show that under the time scaling that has been considered here (the mean time-out of the geometric random variable is scaled as $2nt$), the geometrically distributed time-out converges to an exponentially distributed time-out. We will denote the random variable corresponding to the time-out by τ , where $\tau \sim \exp(1/t)$. In other words, τ is exponentially distributed, with mean duration $E(\tau) = t$.

Thus, in the rest of this paper, we will study unit variance Brownian motion based models for querying, and the details of the models will be described in the appropriate sections of this paper.

3 Main Results and Discussion

The main results in this paper are the following:

- (i) For a source-only search, in Section 4, we show that the probability that a query is unsuccessful decays as $\frac{1}{\log(t)}$, where t is the mean time-out interval.
- (ii) Next, in Section 5, we consider the case where both the source and the destination send sticky probes (thus, memory in the network is utilized). We show that the probability that a query is unsuccessful decays as $\frac{1}{t^{5/8}}$. Further, with k probes (i.e., the source sends multiple queries, and the destination sends multiple advertisements), the probability that a query is unsuccessful decays *at least* as fast as $t^{-5k/8}$.
- (iii) In Section 6, we consider a spatially periodic caching strategy. We show that the probability that a query is unsuccessful decays *no faster than* $\frac{1}{t}$. We also provide an order computation which argues that spatial caching in fact leads to an order $\frac{1}{t}$ decay.

¹We scale time by $2n$ in (1) and (2) to ensure that the limiting Brownian motion has unit variance.

(iv) The sticky search as well as caching utilize memory that is *spatially distributed* over the network. In Section 7, we consider these strategies, as well as a fixed routing strategy that “directs” the query toward the destination. We show that *spreading the memory* leads to a decrease in memory requirement while achieving the same query failure probability.

The results in Section 4 shows that the probability that a query is unsuccessful decays as $(\log t)^{-1}$ for a single source-only search. By using multiple source probes, it is easy to see that the probability decays as $(\log t)^{-k}$, where k is the number of probes. On the other hand, the source and receiver driven search, as well as spatial caching utilize memory in the network. *This enables the decay probability to change from logarithmic decay to polynomial decay.* Thus, the implication is that *no matter how many finite number of source queries are used, we cannot match query techniques which utilize memory that is spatially distributed over the network.* In Section 8, we will see that there is a marked difference in performance between the source-only search and the performance of query/search strategies which use memory.

Next, for the source and receiver driven sticky search, the probability that a query is unsuccessful decay *at least* as fast as $t^{-k5/8}$, where k is the number of probes. However, from the analysis in Section 6 for spatially periodic caching, the probability that a query is unsuccessful decays *at most* as fast as t^{-k} , where k is the number of probes.

Thus, as both strategies have polynomial decay laws, by choosing the *number of probes* appropriately, it is possible to make the decay asymptotics for both the source and receiver driven search, as well as that due to spatial caching to be the same.

Finally, in Section 7, we present an argument to show that the memory requirement for caching is larger (in an order-wise sense) than that for sticky searches. While the spatial caching strategy requires memory that is proportional to the node density n (i.e., $\Theta(n)$), the sticky search strategy requires memory that scales ² as $\Theta(n/\log(n))$.

Further, a caching strategy requires a higher degree of cooperation and infrastructure in the network than to simply send advertisements. These arguments indicate that the appropriate strategy for querying over large sensor networks with little infrastructure support would be to use multiple probes with a sticky strategy.

²A function $f(n)$ is said to be $\Theta(n)$ if there exist positive constants $c_2 > c_1 > 0$ such that for all n large enough, $c_1 \leq f(n)/n \leq c_2$.




Figure 4: A Source driven search

The heuristic reason for the decrease in memory requirement with the sticky strategy is the following. While caching “concentrates” the spatial memory into small regions of space, the sticky strategy distributes it over the path of the query and advertisement. In other words, the sticky strategy *spatially spreads* the memory.

In Section 7, to illustrate the effect of spreading the memory, we consider a *routing strategy* where *every node in the network* points the query toward the destination. In other words, instead of randomly routing a query, there is a small bias at each node that tends to direct the query toward the destination. While this strategy requires *far more infrastructure support* to setup the routing tables at every node, studying this strategy provides insight into spatially spreading the memory. We show that such a strategy also *leads to a polynomial decay in the probability that a query fails*. However, the memory required is far smaller than caching or the sticky strategy, and grows only as $\Theta(\sqrt{n})$.

Thus, the results in this paper illustrate the trade-offs involved in spatial memory, infrastructure and the probability of success in routing a query.

4 Source-only Search

As described in Section 2.1, we consider a query which propagates as a two-dimensional planar Brownian motion, denoted by $B(s)$, and ceases after an exponentially distributed time-out τ , with $E(\tau) = t$.

We assume that the destination is at a (normalized) distance ‘1’ from the source. Without loss of generality, we assume that it is located at $(-1,0)$ on the x-y plane (see Figure 4). The

destination advertises that it has the required information to a small neighborhood (of radius ε) around itself. Thus, the query successfully finds the destination if the Brownian motion trajectory enters the circular region of radius ε about $(-1, 0)$.

Let us define

$$\hat{B}(s) = \begin{cases} B(s) & s \leq \tau \\ B(\tau) & s > \tau \end{cases}$$

Thus, $\hat{B}(s)$ is a Brownian motion that has been “stopped” at a random time τ . Next, we fix some small value of $\varepsilon > 0$, ($\varepsilon < 1$) and let $y = (-1, 0)$. Let $\mathcal{A} = \{x \in \mathfrak{R}^2 \mid d(x, y) \leq \varepsilon\}$ be a ball of radius ε centered at $(-1, 0)$. Define

$$p_{sd}(t) = \Pr(\hat{B}(s) \text{ does not intersect } \mathcal{A})$$

Thus, $p_{sd}(t)$ is the probability that the query does not locate the information before time-out.

Proposition 4.1 *Fix any $\varepsilon \in (0, 1)$. We have*

$$\lim_{t \rightarrow \infty} (\log t) p_{sd}(t) = -2 \log \varepsilon \quad (3)$$

Proof: Let us translate the axes such that the receiver is centered at origin. To do so, let us define

$$\begin{aligned} W(s) &= B(s) - (-1, 0), \\ \hat{\mathcal{A}} &= \{x \in \mathfrak{R}^2 \mid |x| \leq \varepsilon\} \end{aligned}$$

Thus, $W(s)$ is the shifted Brownian motion which starts at $(1, 0)$, and $\hat{\mathcal{A}}$ is a small circular region about origin. Then, from the translation invariance property of Brownian motion [11], we have

$$p_{sd}(t) = \Pr(\hat{W}(s) \text{ does not intersect } \hat{\mathcal{A}})$$

Let $R(t) = |W(t)|$ be the Euclidean distance of $W(t)$ from the origin. $R(t)$ is known as the Bessel process of order zero. Then, we have

$$\begin{aligned} p_{sd}(t) &= \Pr(\hat{W}(s) \text{ does not intersect } \hat{\mathcal{A}}) \\ &= \Pr\left(\inf_{0 \leq s \leq \tau} R(s) > \varepsilon\right), \end{aligned}$$

with $R(0) = 1$. From standard formulas for a Bessel process [1, page 373, 1.2.2], we have that

$$\Pr \left(\inf_{0 \leq s \leq \tau} R(s) > \varepsilon \right) = 1 - \frac{K_0(\sqrt{2/t})}{K_0(\varepsilon\sqrt{2/t})} \quad (4)$$

where $K_0(x)$ is the modified Bessel function of order zero, and is given by

$$\begin{aligned} K_0(x) &= I_0(x) \log \left(\frac{2}{x} \right) - C \\ &\quad + \sum_{k=1}^{\infty} \frac{x^{2k}}{2^{2k}(k!)^2} \psi(k+1), \end{aligned} \quad (5)$$

$$\psi(k+1) = -C + \sum_{i=1}^k \frac{1}{i}, \quad (6)$$

$$I_0(x) = \sum_{k=0}^{\infty} \frac{1}{(k!)^2} \left(\frac{x}{2} \right)^{2k}, \quad (7)$$

where $C = 0.5772$ is the Euler's constant. Let us denote

$$g(x) = \sum_{k=1}^{\infty} \frac{x^{2k}}{2^{2k}(k!)^2} \psi(k+1)$$

To show (3), we first prove the following.

$$\lim_{x \rightarrow 0} |g(x)| = 0 \quad (8)$$

$$\lim_{x \rightarrow 0} |g(x) - g(\varepsilon x)| = 0 \quad (9)$$

These follow directly from the definition of $g(\cdot)$. To see that the limits are well defined, we observe that

$$\psi(k+1) < -C + k < k!.$$

Thus, we have

$$\begin{aligned} |g(x)| &< \sum_{k=1}^{\infty} \frac{1}{2^{2k}(k!)^2} x^{2k} (k!) \\ &= \sum_{k=1}^{\infty} \frac{1}{k!} \left(\frac{x}{2} \right)^{2k} \\ &= e^{x^2/4} - 1 \end{aligned}$$

Thus, (8) follows. The proof of (9) follows from the triangle inequality. Similarly, we have

$$\lim_{x \rightarrow 0} I_0(x) = 1 \quad (10)$$

$$\lim_{x \rightarrow 0} |I_0(x) - I_0(\varepsilon x)| \log(1/x) = 0 \quad (11)$$

Equation (10) follows directly from (7) by substitution. To show (11), for any fixed $0 < \varepsilon < 1$, and for any $x \geq 0$, we have from (7), we have

$$\begin{aligned}
|I_0(x) - I_0(\varepsilon x)| &= \sum_{k=1}^{\infty} \frac{(1 - \varepsilon^{2k})}{2^{2k}(k!)^2} x^{2k} \\
&< \sum_{k=1}^{\infty} \frac{1}{2^{2k}(k!)^2} x^{2k} \\
&< \sum_{k=1}^{\infty} \frac{1}{2^{2k}k!} x^{2k} \\
&= e^{x^2/4} - 1 \\
&\leq \frac{x^2}{4},
\end{aligned}$$

Thus, we have

$$\begin{aligned}
0 &\leq \lim_{x \rightarrow 0} |I_0(x) - I_0(\varepsilon x)| \log(1/x) \\
&\leq \lim_{x \rightarrow 0} \frac{x^2}{4} \log(1/x) \\
&= \lim_{y \rightarrow \infty} \frac{\log y}{4y^2} \\
&= 0
\end{aligned}$$

Next, from (4), we have

$$\begin{aligned}
\log(t)p_{sd}(t) &= \log(t) \left(1 - \frac{K_0(\sqrt{2/t})}{K_0(\varepsilon\sqrt{2/t})} \right) \\
&= \log(t) \left[\frac{\begin{pmatrix} I_0(\varepsilon\sqrt{2/t}) \log(2/(\varepsilon\sqrt{2/t})) \\ -I_0(\sqrt{2/t}) \log(2/\sqrt{2/t}) \\ -g(\sqrt{2/t}) + g(\varepsilon\sqrt{2/t}) \end{pmatrix}}{\begin{pmatrix} I_0(\varepsilon\sqrt{2/t}) \log(2/(\varepsilon\sqrt{2/t})) \\ -C + g(\varepsilon\sqrt{2/t}) \end{pmatrix}} \right] \\
&= \log(t) \left[\frac{\begin{pmatrix} \frac{1}{2} \log(2t)(I_0(\varepsilon\sqrt{2/t}) - I_0(\sqrt{2/t})) \\ -I_0(\varepsilon\sqrt{2/t}) \log(\varepsilon) \\ +(g(\varepsilon\sqrt{2/t}) - g(\sqrt{2/t})) \end{pmatrix}}{\begin{pmatrix} \frac{1}{2} I_0(\varepsilon\sqrt{2/t}) \log(2t) \\ -C + g(\varepsilon\sqrt{2/t}) \\ -I_0(\varepsilon\sqrt{2/t}) \log(\varepsilon) \end{pmatrix}} \right]
\end{aligned}$$

Now, from (10) and (8), it follows that

$$\lim_{t \rightarrow \infty} \frac{\log(t)}{\begin{pmatrix} \frac{1}{2} I_0(\varepsilon \sqrt{2/t}) \log(2t) \\ -C + g(\varepsilon \sqrt{2/t}) \\ -I_0(\varepsilon \sqrt{2/t}) \log(\varepsilon) \end{pmatrix}} = 2 \quad (12)$$

Finally, combining the various limits, from (12), (10), (11) and (9), we have

$$\lim_{t \rightarrow \infty} \log(t) p_{sd}(t) = -2 \log(\varepsilon).$$

■

We finally comment that by using multiple independent queries, it is clear that Proposition 4.1 implies

$$\lim_{t \rightarrow \infty} (\log t)^k p_{sd} = (-2 \log(\varepsilon))^k,$$

where k is the number of queries.

5 Source and Receiver Driven “Sticky” Search

In the previous section, we assumed that the destination advertises in a small neighborhood about itself. Instead, in this section, we consider the case where *both* the source and destination send probes into the network. As before, the source sends (one or more) queries, which are described by means of Brownian motions with exponential time-outs. The destination also sends probes into the network to advertise itself. These advertisements propagate over the network also as Brownian motions with exponential time-outs. In addition, as discussed in Section 2, the probes each leave “sticky” trails as they traverse the networks (see Figure 5), where intermediate nodes in the network remember the query or advertisement that passed through it.

Let the source initiate m queries, each of which is described by a Brownian motion $B_i^{src}(t), i = 1, 2, \dots, m$. Over any interval of time $[0, T]$, we denote $B_i^{src}[0, T], i = 1, 2, \dots, m$ to be the corresponding trajectories. Similarly, let the destination initiate n advertisements, each of which is described by a Brownian motion $B_j^{dst}(t), j = 1, 2, \dots, n$, with the corresponding trajectories over $[0, T]$ denoted by $B_j^{dst}[0, T], j = 1, 2, \dots, n$.




Figure 5: Source-destination “sticky” search, where the query and the advertisement both leave trails in the network

Then, in recent work by Lawler et. al. [16, 17, 18] on intersection exponents for Brownian motion, the authors show the following result. Let us define

$$q_{mn}(T) = \Pr \left[\left(\bigcup_{i=1}^m B_i^{\text{src}}[0, T] \right) \cap \left(\bigcup_{i=1}^n B_i^{\text{dst}}[0, T] \right) = \emptyset \right]$$

Thus, $q_{mn}(t)$ is the probability that the two “bundles” of Brownian motions will never intersect over the time interval $[0, T]$.

Theorem 5.1 [Lawler, Schramm and Werner 2000] *There exist $c > 0$ such that for any $T \geq 1$,*

$$\frac{1}{c} T^{-\eta(m,n)} \leq q_{mn}(T) \leq c T^{-\eta(m,n)}$$

where

$$\eta(m,n) = \frac{1}{96} \left[\left(\sqrt{24m+1} + \sqrt{24n+1} - 2 \right)^2 - 4 \right]$$

In the problem we consider on source and receiver driven search, each Brownian motion has an independent time-out. Let $\tau_i^{\text{src}}, i = 1, 2, \dots, m$, and $\tau_j^{\text{dst}}, j = 1, 2, \dots, n$, be independent, identically distributed exponential random variables, with $E(\tau_i^{\text{src}}) = E(\tau_j^{\text{dst}}) = t$. Let us denote

$$p_{mn}(t) = \Pr \left[\left(\bigcup_{i=1}^m B_i^{\text{src}}[0, \tau_i^{\text{src}}] \right) \cap \left(\bigcup_{j=1}^n B_j^{\text{dst}}[0, \tau_j^{\text{dst}}] \right) = \emptyset \right]$$

Thus, $p_{mn}(t)$ is the probability that the two “bundles” of Brownian motions³, each with an independent time-out will never intersect. In our case, we have randomness due to two sources: the path of the Brownian motions, and the exponential time-outs. Interestingly, we observe that for $n + m \leq 3$, the decay function of $p_{mn}(t)$ depends only on the randomness due to the path of the Brownian motions. Using Theorem 5.1, we will derive the exact asymptote. However, for larger numbers of Brownian motions, the randomness due to the exponential time-outs could also become a factor, and we will derive an upper bound on the non-intersection probability.

Proposition 5.1 *There exist finite, strictly positive constants c_1, c_2 such that*

$$\begin{aligned} c_1 &\leq \liminf_{t \rightarrow \infty} t^{5/8} p_{11}(t) \\ &\leq \limsup_{t \rightarrow \infty} t^{5/8} p_{11}(t) \leq c_2 \end{aligned} \quad (13)$$

$$\begin{aligned} c_1 &\leq \liminf_{t \rightarrow \infty} t p_{21}(t) \\ &\leq \limsup_{t \rightarrow \infty} t p_{21}(t) \leq c_2 \end{aligned} \quad (14)$$

Further, for any $k \geq 1$, there exists a finite, positive constant, $c_3(k) > 0$, such that

$$\limsup_{t \rightarrow \infty} t^{5k/8} p_{kk}(t) \leq c_3(k) \quad (15)$$

Proof: We first consider (13). Let us define $\bar{\tau} = \max\{\tau_1^{src}, \tau_1^{dst}\}$, and $\underline{\tau} = \min\{\tau_1^{src}, \tau_1^{dst}\}$. Then, we have

$$\begin{aligned} p_{mn}(t) &= \Pr \left[\mathbf{B}_1^{src}[0, \tau_1^{src}] \cap \mathbf{B}_1^{dst}[0, \tau_1^{dst}] = \emptyset \right] \\ &\leq \Pr \left[\mathbf{B}_1^{src}[0, \underline{\tau}] \cap \mathbf{B}_1^{dst}[0, \underline{\tau}] = \emptyset \right] \\ &= \int_0^\infty q_{11}(s) \frac{2}{t} e^{-2s/t} ds, \end{aligned}$$

where the last step follows from the definition of $q_{mn}(T)$, and the fact that min of two independent

³From symmetry, we have $p_{mn}(t) = p_{nm}(t)$.

exponential random variables is also an exponential random variable. Thus, we have

$$\begin{aligned}
p_{nm}(t) &= \int_0^1 q_{11}(s) \left(\frac{2}{t}\right) e^{-2s/t} ds \\
&\quad + \int_1^\infty q_{11}(s) \left(\frac{2}{t}\right) e^{-2s/t} ds \\
&\leq \int_0^1 \frac{2}{t} e^{-2s/t} ds \\
&\quad + c \int_1^\infty t^{-5/8} \left(\frac{2}{t}\right) e^{-2s/t} ds
\end{aligned} \tag{16}$$

$$= (1 - e^{-2/t}) + \frac{c\Gamma(3/8, 2/t)}{t^{5/8}}, \tag{17}$$

where $\Gamma(\cdot, \cdot)$ is the incomplete Gamma function. Equation (16) follows from Theorem 5.1, and the fact that $q_{11}(s) \leq 1$. Equation (17) follows from expressions for exponential integrals [5, Section 3.381].

Now, observe that $(1 - e^{-2/t}) = 2/t + o(1/t)$. Further, it follows from the definition of the incomplete Gamma function that $\lim_{t \rightarrow \infty} \Gamma(3/8, 2/t) = \Gamma(3/8) \approx 2.27$. Thus, the upper bound in (13) follows. Similarly, the lower bound in (13) can be derived using the definition of $\bar{\tau}$. The proof of (14) is identical. We skip the details for brevity.

To show (15), observe that

$$p_{kk}(t) \leq [p_{11}(t)]^k \tag{18}$$

To see this, recall that we have k independent Brownian motions originating from the source, which are indexed by $i = 1, 2, \dots, k$; and k independent Brownian motions originating from the destination, which are indexed by $j = 1, 2, \dots, k$. By considering intersections *only between equal indices* (i.e., $i = j$), and neglecting all intersections with indices $i \neq j$, the bound in (18) follows from the independence of the paths and the time-outs. Thus, (15) follows. ■

Remark 5.1 *In Proposition 5.1, we have the exact asymptotics when $n + m \leq 3$. However, when $n + m \geq 4$, we provide only an upper bound.*

This is explained by observing that when $n + m \leq 3$, the upper bound based on the minimum of the time-outs is dominated by the randomness in the Brownian paths, and not the randomness




Figure 6: A spatially periodic cache

in the time-out. This can be seen from (17). The first term in the summation corresponds to the randomness in the time-out, and the second term corresponds to the randomness in the paths. The first term is of order $1/t$, while the second term decays slower. Thus, the asymptote is dominated by the randomness in the Brownian path. However, for $n + m \geq 4$, the asymptote due to the Brownian path decays as $1/t^a$ for some $a > 1$. Thus, a bound using the minimum of time-outs will not provide a tight bound. In general, it is possible that the exact asymptote will depend on both the randomness due to the paths as well as the randomness due to the time-outs.

6 Spatially Periodic Caching

Unlike in the previous sections where the target information was localized, in this section, we consider asymptotics of caching (see Section 2). Assume that the source initiates a query from origin. The destination information is assumed to be periodically cached along a grid (see Figure 6), with the distance from the origin to the nearest cache being normalized to '1'. A small ϵ is chosen, and each cache is assumed to be of radius ϵ .

Let us define $p_{spc}(t)$ to be the probability that the query times-out before it reaches any cache (such as for the trajectory in Figure 6).

Proposition 6.1 Fix any $\varepsilon \in (0, 1)$. We have

$$\lim_{t \rightarrow \infty} t p_{spc}(t) \geq \frac{(1 - \varepsilon)^2}{2} \quad (19)$$

Proof: The result follows from asymptotics of a killed Bessel process. As in Proposition 4.1, let $R(t) = |B(t)|$ be the Euclidean distance of $B(t)$ from the origin.

Suppose that the Bessel process $R(t)$, (the “radius” process) does not exceed $1 - \varepsilon$ before time-out (see Figure 6). Then, this clearly implies that the Brownian motion does not hit a cache before time-out. Thus, we have

$$p_{spc}(t) \geq \Pr \left(\sup_{0 \leq s \leq \tau} R(s) < 1 - \varepsilon \right),$$

with $R(0) = 0$. From standard formulas for a Bessel process [1, page 373, 1.1.2], we have that

$$\Pr \left(\sup_{0 \leq s \leq \tau} R(s) < 1 - \varepsilon \right) = 1 - \frac{1}{I_0((1 - \varepsilon)\sqrt{2/t})}$$

where $I_0(x)$ is defined in (7). Thus, we have

$$t p_{spc}(t) \geq \frac{t(I_0((1 - \varepsilon)\sqrt{2/t}) - 1)}{I_0((1 - \varepsilon)\sqrt{2/t})} \quad (20)$$

Consider the numerator in (20). We have

$$\begin{aligned} t(I_0((1 - \varepsilon)\sqrt{2/t}) - 1) &= t \sum_{k=1}^{\infty} \frac{1}{(k!)^2} \left(\frac{1 - \varepsilon}{2} \sqrt{\frac{2}{t}} \right)^{2k} \\ &= t \sum_{k=1}^{\infty} \frac{1}{(k!)^2} (1 - \varepsilon)^{2k} \left(\frac{1}{2t} \right)^k \\ &= \frac{(1 - \varepsilon)^2}{2} \sum_{k=1}^{\infty} \frac{1}{k^2} \frac{1}{((k-1)!)^2} (1 - \varepsilon)^{2(k-1)} \left(\frac{1}{2t} \right)^{k-1} \\ &= \frac{(1 - \varepsilon)^2}{2} \sum_{l=0}^{\infty} \frac{1}{(l+1)^2} \frac{1}{((l)!)^2} (1 - \varepsilon)^{2l} \left(\frac{1}{2t} \right)^l \\ &= \frac{(1 - \varepsilon)^2}{2} \\ &\quad + \sum_{l=1}^{\infty} \frac{(1 - \varepsilon)^2}{2(l+1)^2} \frac{1}{((l)!)^2} (1 - \varepsilon)^{2l} \left(\frac{1}{2t} \right)^l \end{aligned} \quad (21)$$

Substituting (21) in (20), and taking the limit as $t \rightarrow \infty$, the result follows. ■

The bound derived in this section is optimistic, in the sense that it under-estimates the probability that the query does not locate a cache (note that for a “perfect” search strategy, this probability should be zero). Thus, the result says that the “miss” probability is no better than order $1/t$. In the next section, we present an approximate computation which suggests that the bound captures the correct order of decay with respect to t .

Remark 6.1 *We comment that the computation used in Proposition 6.1 could also be used to model a sensor network with multiple exit points to a wired Internet. Another example could consist of a large number of sensor nodes which measure some physical quantity, and report this data to one of many “fusion centers.” If the location of the fusion centers are not known, the analysis in this section and that in Section 6.1 could be used to determine the time-out that is needed for the message to be transferred to any one fusion center with a low probability of failure.*

6.1 A Computation for Correctness of Order

In this section, we present an approximate computation that indicates that the optimistic bound in Proposition 6.1 is of correct order (with respect to t). As in the previous section, consider a spatially periodic cache as shown in Figure 7, with distance between consecutive caches being $\sqrt{2}$. In addition, we consider a boundary (indicated by the dotted curve in Figure 7), which consists of a circle with four protuberances. For small ϵ , the curve is approximately a circle of radius ‘1’.

Now consider a Brownian motion which starts at origin. There are three possible cases to consider: **(i)** due to time-out, the Brownian motion terminates before it hits the dotted curve in Figure 7 (i.e., before it travels a distance (approximately) ‘1’ from origin) which is illustrated by the thick trajectory in Figure 7, **(ii)** the Brownian motion hits the dotted curve, but *does not* hit a cache, and **(iii)** the Brownian motion hits the dotted curve, and hits a cache.

Consider Case (ii), where the Brownian motion hits the dotted curve, but *does not* hit a cache. This is illustrated in Figure 7 by the trajectory which hits some point ‘x’ on the dotted curve. From the independent increments property of the path of the Brownian motion [11], and the memoryless

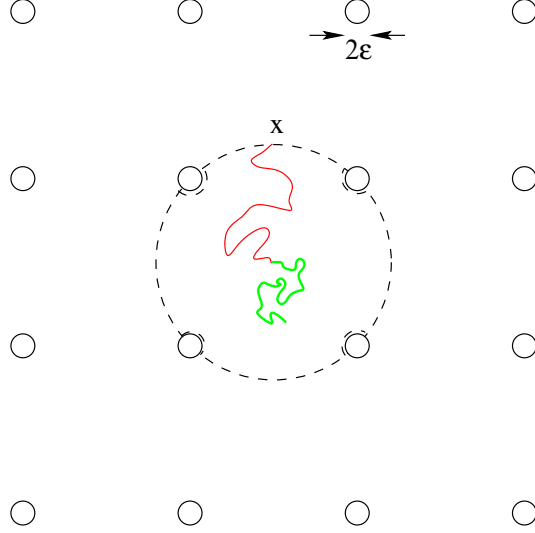


Figure 7: Approximation by periodic “restarting”

property of the exponential time-out, we can ‘restart’ the Brownian motion from location ‘x’ independent of both the path and the time it took to get there. In other words, we can reset the time to zero and consider a new exponential time-out as well as a new Brownian motion with initial value of ‘x’, independent of the past history.

As an approximation, instead of restarting at ‘x’, we will restart the Brownian motion at origin. This seems to be a conservative approximation. To see this, suppose that ‘x’ is very close to a cache. Then, restarting at origin will lead to a smaller cache hitting probability than restarting at ‘x’. However, for a small cache radius ϵ , the probability that ‘x’ is close to a cache is small. For a small value of ϵ , given that the Brownian motion trajectory hits the dotted curve, ‘x’ is (approximately) uniformly distributed over the dotted curve, and the probability that it hits a cache is approximately $(4 \times 2\epsilon)/(2\pi)$. Thus, with large probability, ‘x’ is not very close to a cache.

Now, let $q(t) = 1 - p_{spc}(t)$, be the probability that the Brownian motion trajectory hits a cache before time-out. Next, let $f(\epsilon)$ to be the probability that the trajectory hits a cache on the boundary of the dotted curve given that the Brownian motion trajectory hits the dotted curve. From the discussion in the previous paragraph, for a small value of ϵ , we have

$$f(\epsilon) \approx \frac{4\epsilon}{\pi}$$

Finally, let $g(\epsilon, t)$ be the probability that the trajectory hits the dotted curve (irrespective of whether it hits a cache or not) before time-out. For small ϵ and large t , from the discussion in Proposition 6.1,

we have

$$g(\epsilon, t) \approx 1 - \frac{(1 - \epsilon)^2}{2t}$$

With the above definitions, and our approximation (of restarting the Brownian motion at origin), we have

$$\begin{aligned} q(t) \approx & [g(\epsilon, t)f(\epsilon)] + [g(\epsilon, t)(1 - f(\epsilon))][g(\epsilon, t)f(\epsilon)] \\ & + [g(\epsilon, t)(1 - f(\epsilon))]^2[g(\epsilon, t)f(\epsilon)] + \dots \end{aligned}$$

The first term corresponds to the probability that the Brownian motion trajectory hits a cache when it first hits the dotted curve (Case (iii) we considered earlier). The second term corresponds to the probability that initially Brownian motion hits the dotted curve, but *does not* hit a cache. However, the restarted Brownian motion hits a cache. A similar reasoning holds for the third term, and so on. Thus, summing the geometric series, we have

$$q(t) \approx \frac{g(\epsilon, t)f(\epsilon)}{1 - g(\epsilon, t)(1 - f(\epsilon))}$$

From the definition of $q(t)$, we now have

$$\begin{aligned} p_{spc}(t) &= 1 - q(t) \\ &\approx \frac{1 - g(\epsilon, t)}{1 - g(\epsilon, t)(1 - f(\epsilon))} \end{aligned}$$

Simplifying, and using the approximate expression for $g(\epsilon, t)$, we have

$$p_{spc} \approx \frac{(1 - \epsilon)^2}{2} \frac{1}{f(\epsilon)t + \frac{(1 - \epsilon)^2}{2}(1 - f(\epsilon))}$$

Thus, for large t , and fixed $\epsilon > 0$, this indicates that

$$p_{spc}(t) \sim \frac{1}{t},$$

i.e., this suggests that the bound derived in Proposition 6.1 is of the correct order with respect to t .

7 Memory Requirement and Comparison

As we have discussed in Section 3, for a source only search, the results in Section 4 shows that the probability decays to zero only logarithmically fast (i.e., $(\log t)^{-k}$, where k is the number of probes).

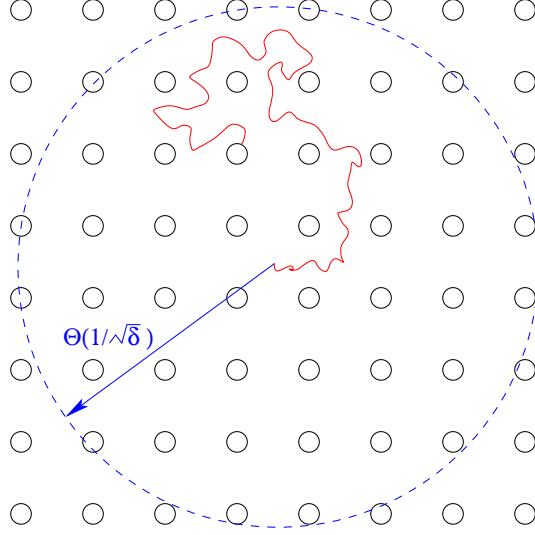


Figure 8: Order computation for memory with periodic caching

On the other hand, the sticky search as well as caching utilize memory in the network, and *this enables the decay probability to change from logarithmic decay to polynomial decay*. Thus, as both are polynomial decay laws, by choosing the number of probes appropriately, it is possible to make the decay asymptotics for the sticky search and due to spatial caching to be the same.

7.1 Memory Comparison of Sticky Search and Caching

Now, let us consider the memory requirements for the spatially periodic caching strategy. From properties of Brownian motion, over an interval of time $[0, T]$, the Brownian motion traverses a distance \sqrt{T} . Thus, given that the probability that a query fails is at most δ , the search time-out should be set to have a mean of (order) $(\delta)^{-1}$ (because $p_{spc}(t) \sim 1/t$). Over this interval of time, the Brownian motion will traverse a distance of order $\sqrt{\delta^{-1}}$ from origin (see Figure 8). Thus, when spatially periodic caching is used, we need to make sure that caching is present at least over the region of radius $\sqrt{\delta^{-1}}$ from origin.

As there is one cache per unit area (a unit square about origin contains four “quarter” caches), we need a total of $\pi/\delta - 1$ caches⁴, where we subtract ‘1’ in the expression to discount for the “original” destination.

⁴The area of a circle with radius $\sqrt{\delta^{-1}}$.

On the other hand, the Lebesgue measure of the trajectory of Brownian motion (i.e., the “area occupied” by the Brownian motion) is zero (though its Hausdorff dimension is 2 which indicates that is not far from having positive Lebesgue measure [4]).

The analysis in this paper is a continuum approximation of a large sensor network. Let us now consider the discrete grid model with n nodes per unit area. As the caching strategy requires memory corresponding to a strictly positive area in the limiting regime, it follows that for large n , the memory requirement will need to scale *linearly* with the number of nodes per unit area. In other words, memory in $\Theta(n)$ nodes will be used to aid the query.

However, a sticky search uses memory only along the path of the Brownian motion. Thus, this strategy will require memory to scale *more slowly than the number of nodes per unit area* (i.e., *memory in $o(n)$ nodes will be used to aid the query*), as the Lebesgue measure of the Brownian trajectory is zero. In fact, the average number of *distinct nodes* that a random walk traverses over n time-steps scales as $\frac{n}{\log(n)}$ (see [15]).

Thus, the above argument indicates that the memory requirement for caching is larger (in an order-wise sense) than that for sticky searches.

Further, a caching strategy requires a much higher degree of cooperation in the network than to simply send sticky probes. These arguments indicate that the appropriate strategy for querying over large sensor networks would be to use multiple sticky probes.

7.2 Spatial Spreading and Memory Requirements

In the previous section, we compared the caching strategy and the sticky search. We showed that both have polynomial decays, but the amount of memory used was order-wise different. The heuristic reason for such a difference is that a sticky search “spreads” the spatial memory whereas a caching strategy “concentrates” the spatial memory.

In this section, we consider a strategy that spreads memory over space in a radially symmetric manner as shown in Figure 9. With such a spreading scheme, the question we are interested in is the following: *How much memory is required in order to result in a failure probability that decays polynomially?*

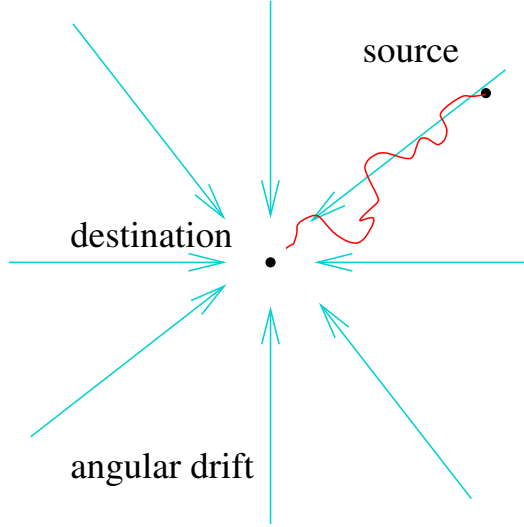


Figure 9: A radially symmetric drift pointing toward the destination.

The main result in this section is that *by spreading the memory in a radially symmetric and “uniform” manner, one needs memory that grows only as $\Theta(\sqrt{n})$, while maintaining polynomial decay in failure probability.*

We recall from the previous section that a sticky search requires spatial memory that grows as $\Theta(n/\log(n))$. At first sight, the radial-symmetric drift strategy seems much better as it requires order-wise less memory. However, note that querying with a radially symmetric drift requires *much more infrastructure* as compared to the sticky search. In particular, *all the nodes in the network* need to have an approximate sense of the location of the destination. On the other hand, *a sticky search requires no additional network support*. Thus, comparing querying with a radially symmetric drift with a sticky search is not a fair comparison. Nevertheless, the results in this section are of interest as they *quantify the effect of spatially spreading the memory*.

As before, we consider a single query generated by a source that is initially located at the coordinates $(1, 0)$ on the x-y plane, and a destination that is located at $(0, 0)$. At any point in space that is a distance r from the origin, we consider a radially symmetric drift $a(r)$ over space that is of the form

$$a(r) = -\mu - \frac{1}{2r}, \quad (22)$$

for some $\mu \geq 0$. The intuition for such a drift component is the following: The first component, μ corresponds to a constant “push” toward the origin. The second component, $\frac{1}{2r}$, corresponds to memory

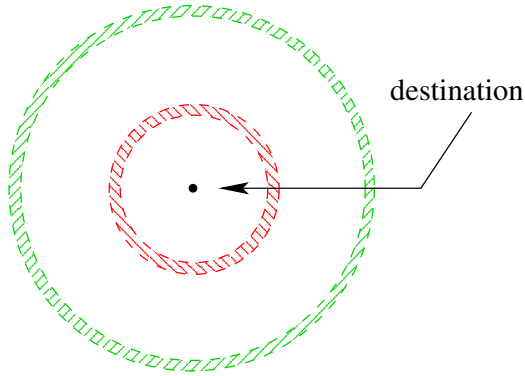


Figure 10: A constant amount of spatial memory is “spread” over each infinitesimal ring about the destination. Thus, the memory “density” decreases with distance.

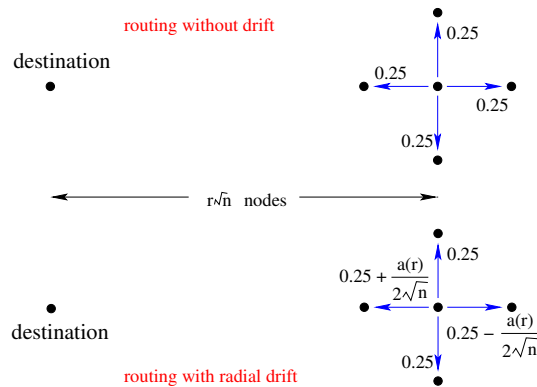


Figure 11: Radial drift over the grid network. A distance of r in the diffusion model corresponds to a $r\sqrt{n}$ hops. With drift, a distance dependent bias, $a(r)$, is introduced. From the diffusion scaling, this corresponds to a $\Theta(1/\sqrt{n})$ bias per node toward the destination.

per node that decay with distance. This decay corresponds to a *constant amount of memory for each infinitesimal ring* about origin (see Figure 10). This follows from the fact that the area of each ring of width dx is given by $2\pi r dx$. Thus, the amount of memory per ring is simply $(2\pi r dx) \times \frac{1}{2r}$. The heuristic for such a component is that *far away from the destination, only an approximate sense of direction is required*. However, *close to the destination, more precise information is required*. From a mathematical perspective, such a drift is chosen such that it compensates for the drift component of the Bessel process (radial distance from the destination) corresponding to the two-dimensional Brownian motion of the query.

Over the grid network described in Section 2.1, random routing at a node can be represented

by means of a *routing vector* with all four components⁵ (corresponding to the four neighbors) being 0.25. This means that the node will pick any one of the four neighbors with equal probability while routing a query (see Figure 11).

On the other hand, with the radial drift strategy over the grid network, instead of routing to any neighbor with equal probabilities of 0.25, there is now a small bias to route toward the destination. From the diffusion scaling, it can be shown that a drift of $a(r)$ at a spatial location that is a distance r (or equivalently, a distance of $r\sqrt{n}$ nodes) from origin corresponds to a bias in the routing probabilities that is of order $a(r)/\sqrt{n}$ (see Figure 11).

The sticky search strategy requires each node to remember the previous hop node that forwarded the query to it. This can be represented by each node having a routing vector that has a component of '1' pointing toward the node from which the query first arrived, and '0' in other directions. Nodes through which the query *did not* pass through would have a routing vector with all components being 0.25 (corresponding to random routing).

A measure of the spatial memory required at each node for a routing strategy is the deviation of the routing vector from the corresponding routing vector for the random strategy⁶. This measure captures the *certainty* with which a node knows the appropriate routing decision in order to reach the destination. In other words, suppose that a strategy requires a routing vector $[p_{00} p_{01} p_{10} p_{11}]$, a measure of the certainty (and thus, the memory required) at the node is computed as $\sum_{i,j} |p_{ij} - 0.25|$.

Thus, with a sticky search, the memory required is $\Theta(1)$ per node through which a query passed. This leads to a total memory requirement that *scales as the number of distinct nodes that a query passes through*. From standard results on a random walk [15], this is of order $\Theta(n/\log(n))$ as discussed earlier in Section 7.

On the other hand, the radial drift strategy requires memory that is of order $|a(r)|/\sqrt{n}$ for *each node*, where r is the distance of the node from the destination. Using an argument analogous to that in Section 7.1, we can “truncate” the infinite plane to a finite region about the destination (illustrated in Figure 8). We now compute the memory required with a radial drift strategy by first decomposing

⁵A routing vector $[p_{up} p_{down} p_{left} p_{right}]$ is vector of non-negative components which sum to 1.

⁶Such a measure can be also be interpreted as a continuum “spreading” of a fixed amount of routing resource (memory) as in [7]. Thus, as the author discusses in [7], a practical implementation could consist of a small number of nodes having precise direction information, and most nodes having no information, resulting in an overall *drift* in a particular direction.

$a(r)$ into two components: (i) the constant drift μ , and (ii) the radially decreasing component $\frac{1}{2r}$. The memory required for the constant component scales as $\Theta(1/\sqrt{n})$ per node. As there are n nodes per unit area, the memory required due to this component is of order $\Theta(\sqrt{n})$. The radially decreasing component $\frac{1}{2r}$ satisfies the property that the amount of memory per unit ring about the destination is constant (i.e., $\Theta(1)$ per ring). As the discrete grid can be decomposed into $\Theta(\sqrt{n})$ rings, it follows that the total memory required for the second component scales as $\Theta(\sqrt{n})$. Hence, it follows that the total memory required with the radial drift strategy grows as $\Theta(\sqrt{n})$. The above discussion is summarized in the following Proposition.

Proposition 7.1 *The spatial memory requirement to support the radial drift strategy described in (22) scales as $\Theta(\sqrt{n})$, where n is the node density.*

Next, we derive the query failure probability with this strategy. With the radial drift strategy in (22), and with a fixed μ , let us define $p_{rd,\mu}(t)$ to be the probability that the query does not locate the information before time-out.

Proposition 7.2 *With the radial drift strategy described in (22), and with $\mu = 0$, we have*

$$\lim_{t \rightarrow \infty} \sqrt{t} p_{rd,0}(t) = 1. \quad (23)$$

Further, for any $\mu > 0$, we have

$$\lim_{t \rightarrow \infty} t p_{rd,\mu}(t) = \frac{1}{2\mu}. \quad (24)$$

Proof: We consider a two dimensional Brownian motion $B(s) = (B_x(s), B_y(s))$, with initial condition $(B_x(0), B_y(0)) = (1, 0)$. As before, the query is assumed to time-out after an exponentially distributed interval of time τ , with $E(\tau) = t$. Let us define

$$\bar{\tau} = \inf\{s \geq 0 : (B_x(s), B_y(s)) = (0, 0)\}$$

to be the first time that the planar Brownian motion hits the origin. Fix any time $s \in [0, \bar{\tau})$, and a radial drift field given by (22). It follows from Ito's formula (see [11, 7]) that at any time s , the radial distance of the location of the Brownian query $R(s)$ from destination (located at $(0, 0)$) is described by the stochastic differential equation

$$dR(s) = \frac{1}{2R(s)} dt + a(R(s)) dt + d\tilde{B}(s),$$

with $R(s) \in [0, \infty)$, $a(\cdot)$ given by (22), and $\tilde{B}(s)$ is a one-dimensional Brownian motion. Thus, we have

$$\begin{aligned} p_{rd,\mu}(t) &= Pr(\tau < \bar{\tau}) \\ &= Pr(\inf_{0 \leq s \leq \tau} R(s) > 0) \end{aligned}$$

From standard results on Brownian motion [1, page 251], it follows that

$$p_{rd,\mu}(t) = 1 - e^{-\mu + \sqrt{\mu^2 + 2/t}} \quad (25)$$

By taking appropriate limits as $t \rightarrow \infty$, (23) and (24) follow. ■

Remark 7.1 *We note that the radially symmetric routing strategy has been studied in [7] in the context of determining the optimal way of distributing a given amount of memory over space to minimize the expected hitting time of the destination. The author in [7] showed that a better strategy than radially distributing memory would be to “concentrate” memory along radial bands (a bicycle wheel spoke like pattern).*

In this paper, we consider only the radially symmetric strategy, as our objective is the quantify the gain due to spreading. However, using the results in [7] and combining with computations similar to that in Section 6, we can show that using a spoke-like routing strategy is order-wise the same as radially symmetric routing (with respect to the large t asymptote). In other words, it can be shown that a spoke-like routing strategy will result in at best $\frac{1}{t}$ decay in failure probability, with a memory that grows as $\Theta(\sqrt{n})$. As our main motivation in this section is to demonstrate that spreading is beneficial, and this has been shown using a simpler radially-symmetric scheme, we skip the details.

8 Simulation Results over a Discrete Grid

In this section, we present simulation results for a grid network. We consider a regular grid with $n = 100$ nodes per unit area (see Figure 1), and the spacing between neighboring nodes is $1/\sqrt{n} = 0.1$. Thus, in this section, we study simulation results for the random walk models described in Section 2 and the scaling described in Section 2.1.

Mean Time (t)	S-O S	S-R S	S-P C
1	0.975	0.515	0.659
10	0.826	0.132	0.125
100	0.646	0.016	0.014

Table 1: Comparison of query strategies over a discrete grid. We consider a grid with density of 100 nodes per unit area. The queries and advertisements propagate as a random walk with a geometrically distributed time-out. The numbers in the table represent probabilities of unsuccessful query. Abbreviations key: Source-Only Search (S-O S), Source-Receiver “sticky” search (S-R S), Spatially-periodic caching (S-P C)

From the scaling described in Section 2.1, an interval of time T will correspond to $2nT$ discrete time-steps. Thus, exponentially distributed time-out with mean t will correspond to a geometrically distributed random variable with mean $2nt$.

In Table 1, we present simulation results for the mean time-out $t = 1, 10$ and 100 respectively, with the cache radius, $\varepsilon = 0.15$.

We see that for the source-only search, the probability that a query is unsuccessful decays extremely slowly with t . However, for the source and receiver driven search as well as spatial caching, the probability that a query is unsuccessful decays much faster. This agrees with the asymptotes which we have earlier discussed, where-in, we have shown that using spatial memory considerably increases the decay rate with respect to t , the mean time-out.

On the other-hand, the probability that the destination is not found is approximately the same for the source and receiver driven search, as well as with spatial caching. It is seen that for large t , spatial caching only slightly out-performs a sticky strategy. As both are polynomial decays, this simulation result is not very surprising, given the analytical results in the previous sections. Thus, the numerical results validate the asymptotic results in this paper.

9 Conclusion

In this paper, we have considered the the problem of a user querying for information over a sensor network, where the user does not have prior knowledge of the location of the information. We have shown that for a source-only search, the probability that a query is unsuccessful decays to zero

only logarithmically fast (i.e., $(\log t)^{-1}$). However, for schemes that utilize spatial memory in the network (the source and receiver driven search as well as spatial caching), this probability decays polynomially fast (i.e., $t^{-\alpha}$, for some $\alpha > 0$ depending on the scheme).

The sticky search as well as caching utilize memory that is *spatially distributed* over the network. We have shown that *spreading the memory* over space leads to a decrease in memory requirement, while maintaining a polynomial decay in the query failure probability. In particular, the memory requirement for the sticky search is smaller (in an order sense) than that for spatial caching, without requiring much infrastructure support. Thus, the results in this paper make a strong case for utilizing a source and receiver driven sticky search in a sensor network with little infrastructure.

References

- [1] A. N. Borodin and P. Salminen. *Handbook of Brownian Motion - Facts and Formulae*. Birkhauser, 2002.
- [2] D. Braginsky and D. Estrin. Rumor routing algorithm for sensor networks. In *First Workshop on Sensor Networks and Applications (WSNA)*, September 2002.
- [3] M. Chu, H. Haussecker, and F. Zhao. Scalable information-driven sensor querying and routing for ad hoc heterogeneous sensor networks. Technical Report P2001-10113, Xerox PARC, 2001.
- [4] J. F. Le Gall. *Some properties of planar Brownian motion*. Ecole d' Ete de Probabilites XX-1990, Lecture Notes in Mathematics 1527. Springer-Verlag, 1990.
- [5] I. S. Gradshteyn and I. M. Ryzhik. *Table of Integrals, Series and Products*. Academic Press, 2000.
- [6] M. Grossglauser and M. Vetterli. Locating nodes with EASE: Last encounter routing in Ad Hoc networks through mobility diffusion. In *Proceedings of IEEE Infocom*, San Francisco, CA, June 2003.
- [7] B. Hajek. Minimum mean hitting times of Brownian motion with constrained drift. In *Proceedings of the 27th Conference on Stochastic Processes and Their Applications*, July 2000.
- [8] C. Intanagonwiwat, R. Govindan, and D. Estrin. Directed diffusion: A scalable and robust communication paradigm for sensor networks. In *Proceedings of ACM Mobicom*, Boston, MA, August 2000.
- [9] R. Jain, A. Puri, and R. Sengupta. Geographical routing using partial information for wireless ad hoc networks. *IEEE Personal Communications*, 8(1):48–57, February 2001.
- [10] J. M. Kahn, R. H. Katz, and K. S. J. Pister. Mobile networking for smart dust. In *Proceedings of ACM Mobicom*, Seattle, WA, August 1999.
- [11] I. Karatzas and S. Shreve. *Brownian Motion and Stochastic Calculus*. Springer, New York, NY, 1996.
- [12] B. Karp and H. T. Kung. GSPR: Greedy perimeter stateless routing for wireless networks. In *Proceedings of the ACM/IEEE International Conference on Mobile Computing and Networking*, pages 243–254, Boston, MA, August 2000.

- [13] A.-M. Kermarrec, L. Massoulie, and A.J. Ganesh. Scamp: Peer-to-peer lightweight membership service for large-scale group communication. In *Proceedings of the Third International Workshop on Networked Group Communications (NGC 2001)*, London, UK, November 2001.
- [14] G. Kesidis, T. Konstantopoulos, and S. Phoha. Surveillance coverage and communication connectivity properties of ad hoc sensor networks under a random mobility strategy. preprint, 2002.
- [15] H. Larralde, P. Trunfio, S. Havlin, H. Stanley, and G. Weiss. Number of distinct sites visited by N random walkers. *Phys. Rev. A*, 45, 1992.
- [16] G. F. Lawler, O. Schramm, and W. Werner. Values of Brownian intersection exponents I: half-plane exponents. *Acta Mathematica*, 187:237–273, 2001.
- [17] G. F. Lawler, O. Schramm, and W. Werner. Values of Brownian intersection exponents II: plane exponents. *Acta Mathematica*, 187:275–309, 2001.
- [18] G. F. Lawler, O. Schramm, and W. Werner. Values of Brownian intersection exponents III: two-sided exponents. *Ann. Inst. Henri Poincaré*, 38:109–123, 2002.
- [19] M. J. Lin, K. Marzullo, and S. Masini. Gossip versus deterministic flooding: Low message overhead and high reliability for broadcasting on small networks. Technical Report CS1999-0637, University of California at San Diego, 18 1999.
- [20] M. Mauve, J. Widmer, and H. Hartenstein. A survey on position-based routing in mobile ad-hoc networks. *IEEE Network Magazine*, 15, November 2001.
- [21] N. Sadagopan, B. Krishnamachari, and A. Helmy. The ACQUIRE mechanism for efficient querying in sensor networks. In *IEEE International Workshop on Sensor Network Protocols and Applications (SNPA'03)*, May 2003.
- [22] S. Shakkottai, R. Srikant, and N. B. Shroff. Unreliable sensor grids: Coverage, connectivity and diameter. In *Proceedings of IEEE Infocom*, San Francisco, CA, June 2003.
- [23] K. Sohrabi, J. Gao, V. Ailawadhi, and G.J. Pottie. Protocols for self-organization of a wireless sensor network. *IEEE Personal Communications*, 7(5):16–27, October 2000.