

Trail: A Distance Sensitive Sensor Network Service For Distributed Object Tracking

VINODKRISHNAN KULATHUMANI, ANISH ARORA, MUKUNDAN SRIDHARAN

Dept. of Computer Science and Engineering

The Ohio State University

and

MURAT DEMIRBAS

Dept. of Computer Science and Engineering

University at Buffalo, SUNY

Distributed observation and control of mobile objects via static wireless sensors demands timely information in a *distance sensitive* manner: information about closer objects is required more often and more quickly than that of farther objects. In this paper, we present a wireless sensor network protocol, *Trail*, that supports distance sensitive tracking of mobile objects for in-network subscribers upon demand. *Trail* achieves a *find* time that is linear in the distance from a subscriber to an object, via a distributed data structure that is updated only locally when the object moves. Notably, *Trail* does not partition the network into a hierarchy of clusters and clusterheads, and as a result *Trail* has lower maintenance costs, is more locally fault-tolerant and it better utilizes the network in terms of load balancing and minimizing the size of the data structure needed for tracking. Moreover, *Trail* is reliable, and energy-efficient, despite the network dynamics that are typical of wireless sensor networks. *Trail* can be refined by tuning certain parameters, thereby yielding a family of protocols that are suited for different application settings such as rate of queries, rate of updates and network size. We evaluate the performance of *Trail* by analysis, simulations in a 90-by-90 sensor network, and experiments on 105 Mica2 nodes in the context of a pursuer-evader control application.

Categories and Subject Descriptors: C.2.1 [**Computer Communication Networks**]: Distributed Networks; C.2.1 [**Computer Communication Networks**]: Wireless Communication; C.2.2 [**Computer Communication Networks**]: Network Protocols

General Terms: Algorithms, Design, Reliability, Performance Evaluation

Additional Key Words and Phrases: Distributed tracking, Network protocol, Applications of sensor actuator networks, Data storage and query, Fault-Tolerance, Scalability, Energy efficiency

This paper is a substantially extended version of the paper that appeared in the *Proceedings of European Conference on Wireless Sensor Networks (EWSN), 2007*.

Permission to make digital/hard copy of all or part of this material without fee for personal or classroom use provided that the copies are not made or distributed for profit or commercial advantage, the ACM copyright/server notice, the title of the publication, and its date appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists requires prior specific permission and/or a fee.

© 2006 ACM 1529-3785/2006/0700-0001 \$5.00

1. INTRODUCTION

Tracking of mobile objects has received significant attention in the context of cellular telephony, mobile computing, and military applications [Abraham et al. 2004; Dolev et al. 1995; Awerbuch and Peleg 1995; Demirbas et al. 2004]. In this paper, we focus on the tracking of mobile objects using a network of static wireless sensors. Examples of such applications include those that monitor objects [Arora et al. 2004; He et al. 2006; Arora and Ramnath 2004], as well as applications that “close the loop” by performing tracking-based control; an example is a pursuer-evader tracking application, where a controller’s objective is to minimize the catch time of evaders.

We are particularly interested in large scale WSN deployments. Large networks motivate several tracking requirements. First, queries for locations of objects in a large network should not be answered from central locations as the source of the query may be close to the object itself but still have to communicate all the way to a central location. Such a centralized solution not only increases the latency but also depletes the intermediate nodes of their energy. Plus, answering queries locally may also be important for preserving the correctness of applications deployed in large WSNs. As a specific example, consider an intruder-interceptor application, where a large number of sensor nodes lie along the perimeter that surrounds a valuable asset. Intruders enter the perimeter with the intention of crossing over to the asset and the objective of the interceptors is to “catch” the intruders as far from the asset as possible. In this case, it has been shown [Cao et al. 2006] that there exist Nash equilibrium conditions which imply that, for satisfying optimality constraints, the latency with which an interceptor requires information about the intruder it is tracking depends on the relative locations of the two: the closer the distance, the smaller the latency. This requirement is formalized by the property of *distance sensitivity* for querying, i.e., the cost in terms of latency and number of messages for returning the location of a mobile object grows linearly in terms of the distance between the object and the querier.

Second, tracking services for large networks must eschew solutions with disproportionate update costs that update object locations across the network even when the object moves only by a small distance. This requirement is formalized by the property of *distance sensitivity* for updates, i.e., the cost of an update is proportional to the distance moved by the object. Querying, by itself, for events or information of interest in WSNs has received significant attention [Intanogonwivat et al. 2003; Ratnasamy et al. 2002; Liu et al. 2004] and some of them focus on distance sensitive querying [Ratnasamy et al. 2002; Sarkar et al. 2006; Funke et al. 2006]. In those solutions, information from a source is published across the network in such a way that queries can be resolved in a distance sensitive manner but the solutions do not address the update of the published information when an object moves. The focus of our work is on the tracking of mobile objects and we additionally require that the *update* cost is distance sensitive.

Contributions: In this paper, we use geometric ideas to design an energy-efficient, fault-tolerant and hierarchy-free WSN service, *Trail*, that supports tracking-based WSN applications. The specification of *Trail* is to return the location of a particular

object in response to an in-network subscriber issuing a *find* query regarding that object. *Trail* has a *find* cost that is linear ($O(d_f)$) in terms of the distance (d_f) of the subscriber from the object. To this end, *Trail* maintains a tracking data structure by propagating mobile object information only locally, and satisfying the distance sensitivity requirement for the track updates. The amortized cost of updating a track when an object moves a distance d_m is $O(d_m * \log(d_m))$.

A basic *Trail* protocol can be refined by tuning certain parameters, thus resulting in a family of *Trail* protocols. Appropriate refinements of the basic *Trail* protocol are well suited for different network sizes and *find/update* frequency settings: One refinement is to tighten its tracks by progressively increasing the rate at which the tracking structure is updated; while this results in updating a large part of the tracking structure per unit move, which is for large networks still *update* distance sensitive, it significantly lowers the *find* costs for objects at larger distances. Another refinement increases the number of points along a track, i.e., progressively loosens the tracking structure in order to decrease the *find* costs and be more *find-centric* when object *updates* are less frequent or objects are static. Moreover, *Trail* scales well to networks of all sizes. As network size decreases, *Trail* gradually eschews local explorations and updates and thus increasingly centralizes *update* and *find*.

We evaluate the performance of *Trail* by simulations in a 90×90 sensor network and experiments on 105 Mica2 nodes in the Kansei testbed [Arora et al. 2006]. This implementation has been used to support a distributed intruder interceptor tracking application where the goal of the interceptor is to catch the intruders as far away from an asset as possible.

Overview of solution: *Trail* maintains tracks from each object to only one terminating point, namely, the center of the network C ; these tracks are *almost* straight to the center, with a stretch factor of at most 1.2 times the distance to C . Note that if the track to an object P is required to be always a straight line from C to the current location of P resulting in a stretch factor equal to 1, then every time the object moves, the track has to be updated starting from C , which would not be a distance sensitive approach. Therefore, we form the track as a set of *trail segments* and update only a portion of the structure depending upon the distance moved. Thus given a terminating point *Trail* maintains a track with lengths from the terminating point that is almost close to minimum. This is important because longer tracks have a higher cost of initialization and given that network nodes may fail due to energy depletion or hardware faults, longer tracks increase the probability of a failed node along a track as well as increase the cost of detecting and correcting failures in the track.

Given the track to an object, a *find* operation explores along circles of exponentially increasing radii until the track is intersected and then follows the track to the current location of the object. The tracks maintained in *Trail* only contain pointers to the current location of an object and not the state information of the object. Publishing the current state (namely location) of the object along all points in track will violate distance sensitivity of updates because every *move* of the object will result in updating the entire track. Following the *trail* of an object from any

location leads to the current location of the object which contains the state of the object. Yet *Trail* is distance sensitive in terms of the *find* in the sense that the total cost of reaching the track for an object, and following the track to reach the object is proportional to the distance of the finder from the object. Note that a *find* explores along circles until a radii that is at most half the distance of the finder to C and then searches at C where the track is certain to be found. Thus C serves as a worst case landmark for finding objects in the network.

In our solution, we make some design decisions like choosing a single point to terminate tracks from all points in the network and avoiding hierarchy in maintaining the tracks. In Section 7, we analyze these aspects of our solution and compare them with other possible approaches.

Organization of the paper: In Section 2, we describe the system model and problem specification. In Section 3, we design the basic *Trail* protocol for a 2-d real plane. Then, in Section 4, we present an implementation of the basic *Trail* protocol for a 2-d sensor network. In Section 5, we discuss refinements of the basic *Trail* protocol. In Section 6, we present results of our performance evaluation. In Section 7, we analyze some design decisions made in our solution and compare them with other possible approaches. In Section 8, we discuss related work and, in Section 9, we make concluding remarks and discuss future work.

2. MODEL AND SPECIFICATION

The system consists of a set of *mobile objects*, and a network of static *nodes* that each consist of a sensing component and a radio component. Tracking applications execute on the mobile objects and use the sensor network to track desired mobile objects. Object detection and association services execute on the nodes, as does the desired *Trail* network tracking service.

The object detection and association service assigns a unique id, P , to every object detected by nodes in the network and stores the state of P at the node j that is closest to the object P . This node is called the *agent* for P and can be regarded as the node where P resides. The problem of detecting objects in the network and uniquely associating them with previous detections is thus orthogonal to the tracking service that we discuss in this paper. Detection and association services can be implemented in a centralized [Sinopoli et al. 2003] or distributed [Shin et al. 2003] fashion; the latter approach would suit integration with the tracking service that we discuss in this paper. As a network level service, we do not make any assumptions about the quality of these detections and leave the implications of false detections and associations to be dealt with at the application layer.

Trail Network Service: *Trail* maintains an in-network tracking structure, $trail_P$, for every object P . *Trail* supports two functions: $find(P, Q)$, that returns the state of the object P , including its location at the current location of the object Q issuing the query and $move(P, p', p)$ that updates the tracking structure when object P moves from location p' to location p .

Definition 2.1 find(P, Q) Cost. The cost of the $find(P, Q)$ function is the total communication cost of reaching the current location of P starting from the current location of Q .

Definition 2.2 $\text{move}(P, p', p)$ *Cost.* The cost of the $\text{move}(P, p', p)$ function is the total communication cost of updating trail_P to the new location p and deleting the tracking structure to the old location p' .

We note that our network service does not assume knowledge of the motion model of objects being tracked, in contrast to some query services [Lu et al. 2005], and as such the scope of every query in our case is the entire network as opposed to a certain locality. Nor does it assume a bound on the number of querying objects in the network or any synchrony between concurrent queries.

Network Model: To simplify our presentation, we first describe *Trail* in a 2-d real continuous plane. We then refine the *Trail* protocol to suitably implement in a random connected deployment of a wireless sensor network. In this model, we impose a virtual grid on the random deployment and snap each node to its nearest grid location (x, y) . Each node is aware of this location. We refer to unit distance as the one hop communication distance. $\text{dist}(i, j)$ now stands for distance between nodes i and j in these units. We describe this model in more detail and the implementation of *Trail* in this discrete model in Section 4.

Fault Model: In the wireless sensor network, we assume that nodes can fail due to energy depletion or hardware faults or there could be insufficient density at certain regions, thus leading to *holes* in the network. However, we assume that the network may not be partitioned; there exists a path between every pair of nodes in the network.

3. TRAIL

In this section, we use geometric ideas to design *Trail* for a bounded 2-d real plane. Let C denote the center of this bounded plane.

3.1 Tracking Data Structure

We maintain a tracking data structure for each object in the plane. Let P be an object being tracked, and p denote its location on the plane. We denote the tracking data structure for object P as trail_P . Before we formally define this tracking structure, we give a brief overview.

Overview: If trail_P is defined as a straight line from C to P , then every time the object moves, trail_P has to be updated starting from C . This would not be a distance sensitive approach. Hence we form trail_P as a set of *trail segments* and update only a portion of the structure depending upon the distance moved. The number of *trail segments* in trail_P increases as $\text{dist}(p, C)$ increases. Rather, the end points of the *trail segments* serve as marker points to update the tracking structure when an object moves. The point from where the update is started depends on the distance moved. Only, when P moves a sufficiently large distance, trail_P is updated all the way from C . We now formally define trail_P .

Definition 3.1 trail_P . The tracking data structure for object P , denoted by trail_P , for $\text{dist}(p, C) \geq 1$ is a path obtained by connecting any sequence of points $(C, N_{mx}, \dots, N_k, \dots, N_1, p)$ by line segments, where $mx \geq 1$, and there exist auxiliary points $c_{1..c_{mx}}$ that satisfy the properties (P1) to (P3) below. mx is defined as

$\lceil (\log_2(\text{dist}(C, p_o))) \rceil - 1$, where p_o is the location of p when trail_P was initialized or updated starting from C .

For brevity, let N_k be the level k vertex in trail_P ; let the level k trail segment in trail_P be the segment between N_k and N_{k-1} ; let $\text{Seg}(x, y)$ be any line segment between points x and y in the plane.

- (P1): $\text{dist}(c_k, N_k) = 2^k$, ($m \geq k \geq 1$).
- (P2): N_{k-1} , ($m \geq k \geq 1$), lies on $\text{Seg}(N_k, c_{k-1})$; $N_{m \times}$ lies on $\text{Seg}(C, c_{m \times})$.
- (P3): $\text{dist}(p, c_k) < 2^{k-b}$, ($m \geq k \geq 1$) and $b \geq 1$ is a constant.

If ($\text{dist}(p, C) = 0$), trail_P is C ; and if ($0 \leq \text{dist}(p, C) < 1$), trail_P is $\text{Seg}(C, p)$. \square

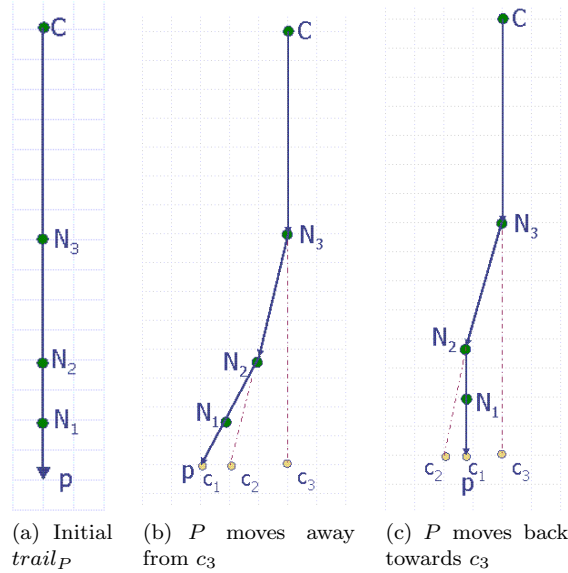


Fig. 1. Examples of Trail to an Object P

Observations about trail_P : From the definition of trail_P , we note that the auxiliary points $c_1..c_{m \times}$ are used to mark vertices $N_1..N_{m \times}$ of trail_P . P1 and P2 describe the relation between the auxiliary points and the vertices of trail_P . Given trail_P , points $c_1..c_{m \times}$ are uniquely determined using P1 and P2. Similarly given p and $c_1..c_{m \times}$, trail_P is uniquely determined. These properties are stated in the following Lemmas.

LEMMA 3.2. *Given trail_P , points $c_1..c_{m \times}$ are uniquely determined.*

PROOF. Extend $\text{Seg}(C, N_{m \times})$ of trail_P by a distance of $2^{m \times}$ to obtain $c_{m \times}$. Similarly extend $\text{Seg}(N_k, N_{k-1})$ by 2^{k-1} to obtain c_{k-1} for $0 < k \leq m \times$. Thus using properties P1 and P2 of trail_P , points $c_1..c_{m \times}$ are uniquely determined given $C, N_{m \times}, \dots, N_1, p$ of trail_P . \square

LEMMA 3.3. *Given $c_1, \dots, c_{m \times}$ and p , trail_P is uniquely determined.*

PROOF. N_{mx} lies on $Seg(C, c_{mx})$ such that $dist(c_{mx}, N_{mx}) = 2^{mx}$. Similarly N_{k-1} lies on $Seg(N_k, c_{k-1})$ such that $dist(c_{k-1}, N_{k-1}) = 2^{k-1}$ for $1 < k \leq mx$. Thus C, N_{mx}, \dots, N_1, p of $trail_P$ are uniquely determined. \square

By property P3, the maximum separation between p and any auxiliary point c_k decreases exponentially as k decreases from mx to 1. When an object moves a certain distance away from its current location, $trail_P$ has to be updated from the smallest index k such that property P3 holds at all levels. By changing parameter b in property P3, we can tune the rate at which the tracking structure is updated. We discuss these refinements in Section 5.

Note from the definition of $trail_P$ that mx is defined as $\lceil (\log_2(dist(C, p_o))) \rceil - 1$ where p_o was the location of the object when $trail_P$ was either created or updated from C . The value of mx which denotes the number of *trail segments* in $trail_P$, depends on the distance of P from C . When $trail_P$ is first created, c_1, \dots, c_{mx} are initialized to location p_o , the number of levels mx is initialized to $\lceil (\log_2(dist(C, p_o))) \rceil - 1$ and $trail_P$ is a straight line. The value of mx is updated when $trail_P$ has moved a sufficient distance to warrant an update of $trail_P$ all the way from C . The update (and create) procedure for $trail_P$ is described in more detail in the following subsection.

We now show 3 examples of the tracking structure in Fig. 1. In this figure, $b = 1$. Fig. 1(a) shows $trail_P$ when c_3, \dots, c_1 are collocated. When P moves away from this location, $trail_P$ is updated and Fig. 1(b) shows an example of $trail_P$ where c_2, c_1 are displaced from c_3 . In Fig. 1(b), $dist(c_3, c_2) = 2$ units, $dist(c_2, c_1) = 1$ unit, and p and c_1 are collocated. Moreover, N_3 lies on $Seg(C, c_3)$, N_2 lies on $Seg(N_3, c_2)$ and so on. In Fig. 1(c) we show an example of a zigzag trail to an object P , when P moves away from c_3 and then moves back in the opposite direction.


3.2 Updating the trail

We now describe a procedure to update the tracking structure when object P moves from location p' to p such that the properties of the tracking structure are maintained and the cost of update is distance sensitive.

Overview: When an object moves distance d away, if the distance $dist(c_1, p) \leq 1$, then the trail is updated by replacing $segment(N_1, p')$ with $segment(N_1, p)$. Otherwise, we find the minimal index m , along $trail_P$ such that $dist(p, c_j) < 2^{j-b}$ for all j such that $mx \geq j \geq m$ and $trail_P$ is updated starting from N_m . In order to update $trail_P$ starting from N_m , we find new vertices $N_{m-1} \dots N_1$ and a new set of auxiliary points $c_{m-1} \dots c_1$. Let $N'_{m-1} \dots N'_1$ and $c'_{m-1} \dots c'_1$ denote the old vertices and old auxiliary points respectively. Starting from N_m , we follow a recursive procedure to update $trail_P$. This procedure is stated below:

Update Algorithm:

- (1) If $dist(p, c_1) > 1$, then let m be the minimal index on the trail such that $dist(p, c_j) < 2^{j-b}$ for all j such that $mx \geq j \geq m$.
- (2) $k = m$
- (3) while $k > 1$

Fig. 2. Updating $trail_P$

- (a) $c_{k-1} = p$; Now obtain N_{k-1} using property P2 as follows: the point on segment N_k, c_{k-1} , that is 2^{k-1} away from c_{k-1} .
- (b) $k = k - 1$

If no indices exist such that $dist(c_k, p) < 2^{k-1}$, then the trail is created starting from C . This could happen if the object is new or if the object has moved a sufficiently large distance from its original position. In this case, mx is set to $(\lceil \log_2(dist(C, p)) \rceil) - 1$. c_{mx} is set to p . N_{mx} is marked on $Seg(C, p)$ at distance 2^{mx} from c_{mx} . Step 1 is executed with $k = mx$. \square

Fig. 2 illustrates an update operation, when $b = 1$. In Fig. 2a, $dist(p, p')$ is 2 units. Hence update starts at N_3 . Initially c_3, c'_2, c'_1 are at p' . We use the *update* algorithm to determine new c_2, c_1 and thereby the new N_2, N_1 . Using step (3a) of the update algorithm, the new c_2 and c_1 lie at p . The vertex N_2 then lies on $Seg(N_3, c_2)$ and N_1 lies on $Seg(N_2, c_1)$. In Fig. 2b, P moves further one unit. Hence update now starts at N_2 . Using step (3a) of the update algorithm, the new c_1 lies at p and N_1 lies on $Seg(N_2, c_1)$.

LEMMA 3.4. *The update algorithm for Trail yields a path that satisfies $trail_P$.*

PROOF. (1) Let m be the index at which *update* starts. By the condition in step 1, $dist(c_j, p) < 2^{j-b}$ for all $mx \geq j \geq m$. Now, for $m > j \geq 1$, $c_j = p$. Therefore for $m > j \geq 1$, $dist(c_j, p) < 2^{j-b}$. Thus property P3 is satisfied.

- (2) Properties P2 and P1 are satisfied because $m \geq k > 1$, we obtain N_{k-1} as the point on $Seg(N_k, c_{k-1})$, that is 2^{k-1} away from c_{k-1} .
- (3) mx is defined for $trail_P$, when $trail_P$ is created or updated starting from C . When mx is (re)defined for $trail_P$, c_{mx} is the position of the object and mx is set to $(\lceil \log_2(dist(C, p)) \rceil) - 1$. \square

Definition 3.5 Trail Stretch Factor. Given $trail_P$ to an object p , we define the trail stretch factor for any point x on $trail_P$ as the ratio of the length along $trail_P$ from x to p , to the Euclidean distance $dist(x, p)$. \square

LEMMA 3.6. *The maximum Trail Stretch Factor for any point along $trail_P$, denoted as TS_p is $sec(\alpha) * sec(\frac{\alpha}{2})$ where $\alpha = arcsin(\frac{1}{2^b})$.*

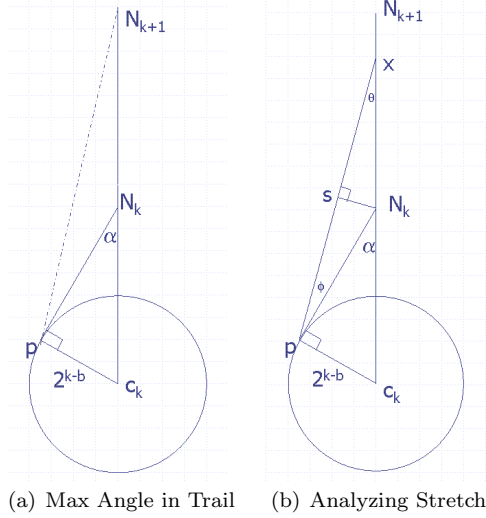


Fig. 3. Analyzing Trail Stretch Factor

PROOF. We prove Lemma 3.6 by using the following steps.

◇ (*Maximum angle ($\angle pN_k c_k$)*) Let the maximum angle formed by p and c_k at N_k in $trail_P$ for ($mx \geq k \geq 1$) be denoted as α . Refer Fig. 3(a)). Recall from properties of $trail_P$ that $dist(N_k, c_k) = 2^k$ and $dist(p, c_k) < 2^{k-b}$. Note that $\angle pN_k c_k$ is maximum when $Seg(N_k, p)$ is tangent to a circle of radius 2^{k-b} and center c_k . Therefore we have the following condition.

$$\alpha < \arcsin\left(\frac{1}{2^b}\right) \quad (1)$$

◇ (*Maximum stretch at a given level k*) Let x be any point on $trail_P$ which lies on $Seg(N_{k+1}, N_k)$. Refer Fig. 3(b). We note the following equation based on the geometry of Fig. 3(b).

$$\frac{(dist(x, N_k) + dist(N_k, p))}{dist(x, p)} = \frac{(\sin(\theta) + \sin(\phi))}{\sin(\theta + \phi)} \quad (2)$$

Also note that $(\theta + \phi) = \angle pN_k c_k$. Using this, we get Eq. 3

$$(\theta + \phi) \leq \alpha \quad (3)$$

Let $f(\theta, \phi)$ denote the following function.

$$f(\theta, \phi) = \frac{(\sin(\theta) + \sin(\phi))}{\sin(\theta + \phi)} \quad (4)$$

It can be shown that $f(\theta, \phi)$, where $\theta > 0$, $\phi > 0$ and $\theta + \phi \leq \alpha$ is maximum when $\theta = \phi = \frac{\alpha}{2}$. We state this as Proposition B.1 and prove it in Appendix B. Substituting $(\theta = \phi)$ in Eq. 2, we get the following condition.

$$\frac{(dist(x, N_k) + dist(N_k, p))}{dist(x, p)} \leq sec(\frac{\alpha}{2}) \quad (5)$$

Thus, we see that at a single level k the maximum stretch for $trail_P$ occurs when $\angle pN_k c_k$ is α for a point x on $Seg(N_{k+1}, N_k)$ such that $\angle pxN_k = \angle xpN_k$. Since $\angle pxN_k = \angle xpN_k$ when the maximum occurs, we also get the following equation.

$$\frac{(dist(x, N_k))}{dist(x, p)} = \frac{(dist(N_k, p))}{dist(x, p)} = \frac{sec(\frac{\alpha}{2})}{2} \quad (6)$$

◇ (*Maximum trail stretch factor from vertex N_k to p*) In order to find the maximum stretch factor over multiple levels, we consider $trail_P$ to be *split* at vertices N_k to N_2 in such a way that the stretch is maximized at each level. Thus we let $\angle pN_j c_j = \alpha$ and $\angle pN_j N_{j-1} = \angle N_j p N_{j-1}$ for all $k \geq j > 1$. In Fig. 4, we show one such *trail segment*, $Seg(N_j, N_{j-1})$ of $trail_P$.




Fig. 4. Analyzing Trail Stretch

Using Eq. 6, we get the following equations:

$$\frac{dist(N_{j-1}, p)}{dist(N_j, p)} = \frac{sec(\frac{\alpha}{2})}{2} \quad \forall j : k \geq j > 1 \quad (7)$$

$$\frac{dist(N_j, N_{j-1})}{dist(N_j, p)} = \frac{sec(\frac{\alpha}{2})}{2} \quad \forall j : k \geq j > 1 \quad (8)$$

When the above configuration is repeated at all levels of $trail_P$, we determine the ratio of the lengths of two successive *trail segments*, $Seg(N_{j-1}, N_{j-2})$ and $Seg(N_j, N_{j-1})$. Using Eq. 7 and Eq. 8, we get the following equation.

$$\frac{dist(N_{j-1}, N_{j-2})}{dist(N_j, N_{j-1})} = \frac{sec(\frac{\alpha}{2})}{2} \quad \forall j : k \geq j > 2 \quad (9)$$

Let L_k denote the length along $trail_P$ from vertex N_k . Using Eq. 9, we get:

$$L_k = dist(N_2, N_1) * \sum_{j=0}^{j=(k-2)} (2 * cos(\frac{\alpha}{2})^j) + dist(N_1, p) \quad (10)$$

Let R_k denote the trail stretch factor from N_k .

$$R_k = \frac{L_k}{dist(N_k, p)} \quad (11)$$

Upon simplification using Eq. 7, Eq. 8 and Eq. 10, we get:

$$R_k = \frac{1}{2 * cos(\frac{\alpha}{2}) - 1} + \frac{1}{(2 * cos(\frac{\alpha}{2}))^{k-1}} * (1 - \frac{1}{2 * cos(\frac{\alpha}{2}) - 1}) \quad (12)$$

Since, $\alpha < \frac{\pi}{6}$ for $b \geq 1$, $0 < 2 * cos(\frac{\alpha}{2}) - 1 \leq 1$. Therefore we get:

$$R_k \leq \frac{1}{2 * cos(\frac{\alpha}{2}) - 1} \quad (13)$$

Since, $cos(\alpha) = 2 * cos^2(\frac{\alpha}{2}) - 1$, $cos(\frac{\alpha}{2}) \leq 1$ and $0 < 2 * cos(\frac{\alpha}{2}) - 1 \leq 1$, we get:

$$R_k \leq \frac{1}{cos(\alpha)} \quad (14)$$

◇ (*Maximum trail stretch factor from any point in $trail_P$ to p*) Let x be any point on $trail_P$ which lies on a level $k + 1$ segment, i.e $Seg(N_{k+1}, N_k)$, but is not a *vertex* point. Let L_{xp} denote the length along $trail_P$ from x to p . Using Eq. 14 and Eq. 5, we have the following inequalities:

$$\begin{aligned} L_{xp} &\leq dist(x, N_k) + dist(N_k, p) * sec(\alpha) \\ &\leq (dist(x, N_k) + dist(N_k, p)) * sec(\alpha) \\ &\leq sec(\frac{\alpha}{2}) * sec(\alpha) * dist(x, p) \end{aligned}$$

Thus we have proved that the maximum *Trail Stretch Factor* for any point along $trail_P$, denoted as TS_p is $sec(\alpha) * sec(\frac{\alpha}{2})$ where $\alpha = arcsin(\frac{1}{2^b})$. □

LEMMA 3.7. *The length of $trail_P$ for an object P starting from a level $k(0 < k \leq mx)$ vertex, denoted as L_k is bounded by $(2^k + 2^{k-b}) * TS_p$.*

Proof Sketch: $dist(c_k, p) < 2^{k-b}$. Therefore, $dist(N_k, p) < 2^k + 2^{k-b}$. Then using Lemma 3.6, the result follows. □

THEOREM 3.8. *The upper bound on the amortized cost of updating $trail_P$ when object P moves distance $d_m(d_m > 1)$ is $4 * (2^b + 1) * TS_p * d_m * log(d_m)$.*

PROOF. Note that in *update* whenever $trail_P$ is updated starting at the level k vertex, we set $c_{k-1} = p$. P can now move a distance of 2^{k-1-b} before another

update starting at the level k vertex. Thus, between any two successive updates starting from a level k vertex, the object must have moved at least a distance of 2^{k-1-b} . The total cost to create a new path and delete the old path starting from a level k vertex costs at most $2 * L_k$.

When an object moves a distance d_m where $d_m > 1$, it could involve multiple updates at smaller distances. The object could be detected at multiple instances over this distance d_m . Therefore we calculate the upper bound on the amortized cost of *update* when the object moves distance d_m . We consider the minimum distance to trigger an update to be 1 unit. Note that between any two successive updates starting from a level k vertex, the object must have moved at least a distance of 2^{k-1-b} . Thus over a distance d_m , update can start at level $(b+1)$ vertex at most d_m times, update can start at level $(b+2)$ vertex can at most $\lfloor d_m/2 \rfloor$ times, and so on. The *update* can start at level $(\lfloor \log_2(d_m) \rfloor + b + 1)$ vertex at most once. Adding the total cost, the result follows. \square

b	Trail Stretch	Update Cost
1	1.2	$14 * d_m * \log d_m$
2	1.05	$20 * d_m * \log d_m$
> 3	Approaches 1	$4 * (2^b + 1) * d_m * \log d_m$

Fig. 5. Effect of b on *Update* cost

For illustration, we summarize the *Trail Stretch* factor and *update* costs for different values of b in Fig. 5. We explain the significance of the refinement of *Trail* by varying b in Section 5.

3.3 Basic Find Algorithm

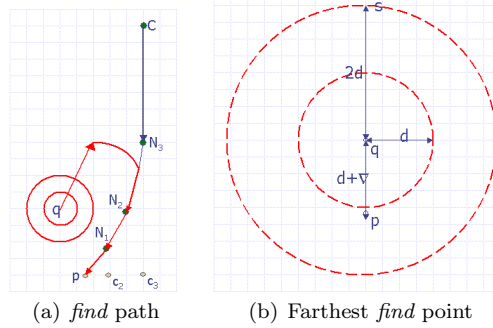
Given *trail_P* exists for an object P , we now describe a basic *find* algorithm that is initiated by object Q at point q on the plane. We use a basic ring search algorithm to intersect *trail_P* starting from Q in a distance sensitive manner. We then show from the properties of the *Trail* tracking structure that starting from this intersection point, the current location of P is reached in a distance sensitive manner.

Basic *find* Algorithm:

- (1) With center q , successively draw circles of radius $2^0, 2^1, \dots, 2^{\lfloor \log(\text{dist}(q,C)) \rfloor - 1}$, until *trail_P* is intersected.
- (2) If *trail_P* is intersected, follow it to reach object P ; else follow *trail_P* from C (note that if object exists, *trail_P* will start from C).

THEOREM 3.9. *The cost of finding an object P at point p from object Q at point q is $O(d_f)$ where d_f is $\text{dist}(p, q)$.*

PROOF. Note that as q is distance d_f away from p , a circle of radius $2^{\lfloor \log(d_f) \rfloor}$ will intersect *trail_P*. Hence the total length traveled along the circles before intersecting *trail_P* at point s is bounded by $2 * \pi * \sum_{j=1}^{\lfloor \log(d_f) \rfloor} 2^j$, i.e., $8 * \pi * d_f$. The total cost of connecting segments between the circles is bounded by $2 * d_f$.

Fig. 6. Basic find algorithm in *Trail*

Now, when the trail is intersected by the circle of radius $2^{\lceil \log(d_f) \rceil}$, the point s at which the trail is intersected can be at most $3 * d_f$ away from the object p . This is illustrated in Fig. 6(b). In this figure, q is $d_f + \nabla$ away from p . Hence the trail can be missed by circle of radius $2^{\log(d_f)}$. From Lemma 3.6, we have that distance along the trail from s to p is at most $3 * TS_p * d_f$. Thus, the cost of finding an object P at point p from object Q at point q is $O(d_f)$ where d_f is $dist(p, q)$. \square

Remark on update cost: Note that in Theorem 3.8 we have characterized the upper bound on the amortized cost of *update* when an object moves distance d_m . This is because over a distance of d_m , the object can be *detected* in multiple instances and consequently multiple *update* of the track can result over a distance d_m . We consider the minimum distance to trigger an update to be 1 unit. We add the total cost resulting from each *update* and the sum of the cost of all these possible updates results in an upper bound on the amortized cost.

But we note that, when an object makes a discrete or direct jump of distance d_m , the logarithmic factor in the cost disappears. In other words, when an object disappears at a certain location and is later detected at another location, the cost of individual updates need not be added. However there is an additional cost of exploring to find an existing trail because we do not assume memory of the previous location of an object. In the continuous setting we ignore the cost of finding an existing *trail_P* as it will exist within a distance of 1 unit. The cost of *update* when an object makes a discrete jump of distance d_m is summarized below.

THEOREM 3.10. *The cost of updating trail_P when object P makes a discrete jump of distance d_m is $O(d_m)$.*

PROOF. When P moves distance d_m away, an existing *trail_P* can be intersected at a cost of $8 * \pi * d_m$. This follows from the proof of Theorem 3.9. The highest level at which the update can start is level $(\lfloor \log_2(d_m) \rfloor + b + 1)$ vertex. The cost of modifying the trail pointers is then bounded by $4 * (2^b + 1) * TS_p * d_m$. The result follows. \square

4. IMPLEMENTING TRAIL IN A WSN

In this section, we describe how to implement the *Trail* protocol in a WSN that is a discrete plane, as opposed to a continuous plane as described in the previous section.

Trail can be implemented in any random deployment of a WSN aided by some approximation for routing along a circle. We describe one such implementation below.

System model: Consider any random deployment of nodes in the WSN. We impose a virtual grid on this deployment and snap each node to its nearest grid location (x, y) and assume that each node is aware of this location. We refer to unit distance as the one hop communication distance. $dist(i, j)$ now stands for distance between nodes i and j in these units. The separation along the grid is less than or equal to the unit distance. When the network is dense, the grid separation can be smaller. The neighbors of a node are a set of all nodes within unit distance of the node. Thus when the grid separation is unit distance, there are at most 4 neighbors for each node. We also assume the existence of an underlying geographic routing protocol such as GPSR [Karp and Kung 2000], aided by an underlying neighborhood service that maintains a list of neighbors at each node.

Note: We have implicitly assumed in the above model that each location on the grid is mapped to a unique node. This can be achieved by decreasing the size of the grid to be equal to the smallest separation between any 2 nodes in the network. However this assumption is not necessary. In other words, all nodes need not be necessarily assigned to some location on the grid. Nodes can take turns to play the role of a given location. But for ease of exposition, we have abstracted away these possibilities in our model.

Fault Model: We assume that nodes in the network can fail due to energy depletion or hardware faults, or there could be insufficient density at certain regions, thus leading to *holes* in the network. However, we assume that the network may not be partitioned; there exists a path between every pair of nodes in the network. A node may also transiently fail. But we assume that the failed nodes *return* in a clean or *null* state without any knowledge of tracks previously passing through them. We do not consider arbitrary state corruptions in the nodes.

When implementing on a WSN grid, *Trail* is affected by the following factors:(1) discretization of points to nearest grid location; (2) Overhead of routing between any two points on the grid; and (3) *holes* in the network. We discuss these issues in this section.

Routing stretch factor: When using geographic routing to route on a grid, the number of hops to communicate across a distance of d units will be more than d . We measure this stretch in terms of the routing stretch factor, defined as the ratio of the communication cost (number of transmissions) between any two grid locations, to the Euclidean distance d between two grid locations. It can be shown that the upper bound on the routing stretch factor for the WSN unit grid is $\sqrt{2}$. If we consider the grid to be of smaller separation than the communication range (denser grid), then the routing stretch factor will decrease as any straight line will now be approximated more closely when moving along the grid.

4.1 Implementing *find* on WSN Grid

We now describe how to implement the *find* algorithm in the WSN grid. As seen in Section 3, during a find, exploration is performed using circles of increasing radii

around the finder. However, in the grid model, we approximate this procedure and instead of exploring around a circle of radius r , we explore along a square of side $2 * r$. The perimeter of the square spans a distance $8 * r$ instead of $2 * \pi * r$. Tighter approximations of circle are possible, but approximating with a square is simpler. We characterize the upper bound on the *find* cost in the following Lemma, whose proof can be found in Appendix C.

LEMMA 4.1. *The upper bound on the cost of finding an object P at point p from object Q at point q is $(32 + 3 * \sec(\alpha) * \sec(\frac{\alpha}{2}) * \sqrt{2}) * d_f$ where d_f is $\text{dist}(p, q)$ and $\alpha = \arcsin(\frac{1}{2b})$.*

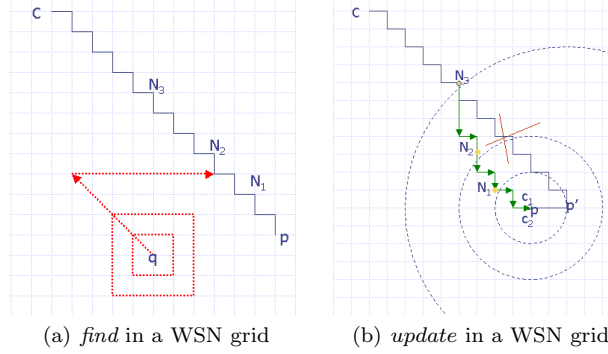


Fig. 7. Find and update algorithm in a WSN grid

4.2 Implementing *Update* on WSN Grid

For each object P in the network, $trail_P$ is maintained by parent / child pointers at each node in the network. Starting from C , following the child pointers for P would lead to the current location of P . Similarly, starting from p , following the parent pointers at each node will lead to C .

We use three types of messages in the update actions. Initially, when an object is detected at a node, it sends an *explore* message that travels in around the square perimeters of increasing levels until it meets $trail_P$ or it reaches the center. Note that if the object is updated continuously as it moves, then the *explore* message will intersect the trail within a 1 hop distance. As before, the trail update is started from the level m vertex node where m is the minimal index such that $\text{dist}(c_m, p) < 2^{m-1}$ for all j such that $mx \geq j \geq m$.

Starting from the level m node where update is started, a new track is created by sending a *grow* message towards c_{m-1} . Geographic routing is used to route the message towards c_{m-1} . On this route, the node closest to, but outside a circle of radius 2^{m-1} around c_{m-1} is marked as N_{m-1} . This node keeps memory of c_j for all levels $mx \geq j \geq m$, which is used to determine the smallest level at which an update should start. This procedure is then repeated at subsequent vertex nodes and the trail is updated. Fig. 7(b) shows how a trail is updated in the grid model with the grid spacing set equal to the unit communication distance. The vertex pointers N_3, \dots, N_1 are shown approximated on the boundary of the respective circles. Also, starting from the level k node where update is started, a *clear* message is used to delete the old path. We formally state the *update* and *find* algorithms in guarded command notation in Appendix H.

The heartbeat actions are sent by each node along $trail_P$ to its child. At any node r , if a heartbeat is not received from its child, a *search* message is sent using geographic routing tracing the boundary of the hole. $trail_P$ is reinforced starting from the first node where the *search* message intersects $trail_P$ using a *reinforce* message along the reverse path. (If the goal is to find $trail_P$ in the shortest time, the *search* should likely be enforced in both directions along the boundary of the hole).

Tolerating failure of C : The terminating point C provides a sense of direction to the trail and serves as a final landmark for the *find* operation. If C fails, the node that is closest to C will takeover the role of C . However, even in the transient stage when there is no C , the failure of C is tolerated locally. We describe this below.

Consider that C and all nodes in a contiguous region around C have failed. In this case, a *search* message will be initiated from the node closest to C that belongs to $trail_P$. Because a contiguous set of nodes surrounding C have failed, the *search* message eventually returns to the node initiating the *search* by following the boundary of the hole. Thus an existing $trail_P$ terminates at the node that belongs to $trail_P$ and is closest to C . Thus if a *find* message is unable to reach C , then routing along the boundary of the *hole* will intersect $trail_P$.

When a new node takes over the role of C , we do not assume that the state of the original C is transferred. The new C has no knowledge of tracks passing through it. Eventually, the maintenance actions for *Trail* will result in all tracks terminating at C .

5. REFINEMENTS OF TRAIL

In Sections 3 and 4, we have described the basic *Trail* protocol. In this section, we discuss two techniques to refine the basic *Trail* network protocol: (1) tuning how often to update a *Trail* tracking structure, and (2) tuning the shape of a *Trail* tracking structure. This yields a family of protocols.

In the first refinement we alter the parameter b to values greater than 1. By increasing b , we update the track of an object more often. This results in straighter tracks with smaller stretch factor. As tracks get straighter, the *find* pattern at higher levels of the search can follow a triangular pattern (as illustrated in Fig. 9) as opposed to complete circles. In fact, as b increases circular explorations can be avoided at more levels of the *find*. Thus the average find cost in the network decreases as b increases. In sum by increasing b , we increase the cost of update and decrease the cost of find. This refinement can be used when the rate of updates is small as compared to the rate of find. For example, as the speed of objects in the network decreases, we can increase the value of b .

In the second refinement we change the length of each segment in *Trail*. In the basic protocol, the segment at each level is a straight line to the next lower level. In this refinement, we modify the length of each segment to be a straight line to the next level plus an arc of length $x \times 2^k$. As x increases the amount of update at each level increases, but the *find* exploration can now be smaller. Specifically when

$x = 2 \times \pi$, the *find* is a straight line towards the center of the network. We call this particular parameterization, the *find-centric Trail* protocol.

5.1 Tightness of *Trail* Tracking Structure

The frequency at which $trail_P$ is updated depends on parameter constant b in property *P3* of $trail_P$. As seen in Section 3, for values of $b > 1$, $trail_P$ is updated more and more frequently, hence leading to larger update costs. However, $trail_P$ becomes tighter and tends to a straight line with the trail stretch factor approaching 1. We exploit this tightness of $trail_p$ to optimize the *find* strategy.

5.1.1 *Optimization of find.* The intuition behind this optimization is that since the trail to any object P originates at C , the angle formed by p with C and the higher level vertices is small and bounded. Hence as the levels of explorations increase in *find*, we can progressively decrease the size of exploration from full circles. We now describe the details of this optimization.

LEMMA 5.1. *Given $trail_P$, $(\angle C, p, N_k) < (mx - k + 1) * (\arcsin(\frac{1}{2^b}))$, where $(mx \geq k \geq 1)$. □*

The proof can be found in Appendix D.

From hereon, we let d_{pC} denote the distance of any object P from C . After the value of mx is defined for a $trail_P$, object P can move for a certain distance before mx is redefined. Therefore, given d_{pC} , the value of mx in $trail_P$ cannot be uniquely determined; however, we note that the value of mx can be bounded given d_{pC} and we define $m\hat{x}_p$ as the highest possible value of mx in $trail_P$, given d_{pC} . We now determine $m\hat{x}_p$.

Let R denote the network *radius*, defined as the maximum distance from C . Recall that mx denotes the number of levels in the track for an object P . mx is defined as $\lceil (\log(\text{dist}(C, p_o))) \rceil - 1$ where p_o is the position of the object when $trail_P$ was (re)created from C . Given network radius R , let \top be the highest number of levels possible for any object in the network. Thus in a given network, $\top = \lceil (\log(R)) \rceil - 1$.

LEMMA 5.2. *Given d_{pC} , $m\hat{x}_p = \text{minimum}(\lceil \log(d_{pC}) \rceil, \top)$. □*

The proof can be found in Appendix E.

Since given d_{pC} , the index of highest level mx in $trail_p$ cannot be uniquely determined, we state the maximum angle formed by $\angle CpN_k$ in terms of $m\hat{x}_p$ rather than mx . When the actual mx in $trail_P$ is lesser than $m\hat{x}_p$, then the actual maximum angles formed by $\angle C, p, N_k$ where $1 \leq k \leq mx$ is lower than the maximum angles stated in the following equation.

$$(\angle C, p, N_k) < (m\hat{x}_p - k + 1) * (\arcsin(\frac{1}{2^b})) \quad (15)$$

We note from the above properties that the angles formed by p with C and the higher level vertices are small and bounded. For example, the angle formed by p with C and N_{mx} is less than $(\arcsin(\frac{1}{2^b}))$. Because of this, at higher levels of the

search, a *find* operation does not have to explore in circles and yet can be guaranteed to intersect the track for an object. In the analysis below, we characterize the minimum size of exploration required at each level of exploration given the distance of finder object q from C . Only for ease of explanation, we assume that $b = 3$.

Let Q be the finder at distance d_{qC} from C . Thus $m\hat{x}_q = \text{minimum}(\lceil \log(d_{qC}) \rceil, \top)$. Let P be an object which should be found at the level k exploration. At level k of the exploration, $trail_P$ for any location of P within the circle of radius 2^k around q should be intersected. A circular exploration of radius 2^k around q is sufficient to achieve this. We now characterize the necessary exploration.

We show that, at levels of exploration k where $k \geq m\hat{x}_q - 7$, circular explorations can be avoided and instead a pattern of exploration along the base of an isosceles triangle with apex q and length of base determined by Fig. 10 is sufficient to intersect the trails of all objects at distance 2^k from Q . The base of the isosceles triangle is such that $segment(C, q)$ is the perpendicular and equal bisector of the base of the triangle. At levels of exploration $k < m\hat{x}_q - 7$, exploration along the entire circle is necessary. For reasons of exposition, we have moved the procedure to determine the necessary exploration at each level, to Appendix F.

Exploration level k	Length of triangle base	Height of triangle
$m\hat{x}_q - 1$	$2 * 2^k$	2^k
$m\hat{x}_q - 2$	$2.5 * 2^k$	2^k
$m\hat{x}_q - 3$	$3.1 * 2^k$	2^k
$m\hat{x}_q - 4$	$3.7 * 2^k$	2^k
$m\hat{x}_q - 5$	$4.3 * 2^k$	2^k
$m\hat{x}_q - 6$	$5 * 2^k$	2^k
$m\hat{x}_q - 7$	$6.2 * 2^k$	2^k

Fig. 10. Optimized *find*: pattern of exploration

Optimized find algorithm (b=3):

- Explore at levels k ranging from 0 to $(\lceil \log(d_{qC}) \rceil - 1)$. If $k < (m\hat{x}_q - 7)$, explore using a circle of radius 2^k around q . Else explore along the base of an isosceles triangle with apex q and length of base determined by Fig. 10. The base of the triangle is such that $segment(C, q)$ is the perpendicular and equal bisector of the base of the triangle. □

An example for the modified *find* algorithm is shown in Fig. 9. In this figure the object q is at distance 48 units from C . $m\hat{x}_q$ is 6. $\lceil \log(d_{qC}) \rceil - 1 = 4$. Therefore, the levels of exploration are in the range 0..4. Exploration is along the base of triangles at all levels. (This figure is not to scale but for illustration.)




Fig. 9. Optimized *find*: example

Impact of the Optimization: The optimization of *find* at higher levels is thus significant in that it yields: (1) smaller upper bounds for objects that are far away

from the finder; and (2) lower average cost of $find(p, q)$ over all possible locations of q and p .

As described earlier when $b = 3$, circular explorations are avoided at the highest 7 levels of the $find$ operation. As the value of b increases, the number of levels at which circular explorations can be avoided, increases. But by increasing b , we update the track of an object more often. Thus by increasing b , we increase the cost of update and decrease the cost of find.

We note that there are limits to tuning the frequency of updates, because for extreme values of b distance sensitivity may be violated. For example, for large values of b , that cause $dist(p, c_k) < y$ where y is a constant we end up with having to update the entire $trail_P$ when an object moves only a constant distance y . Similarly, for values of $b < 0$, the *Trail Stretch Factor* becomes unbounded with respect to distance from an object. Thus an object could be only δ away from a point on $trail_P$, yet the distance along $trail_P$ from this point to the p could travel across the network.

5.2 Modifying Trail Segments

The second refinement to *Trail* is by varying the shape of the tracking structure by generalizing property *P2* of $trail_P$. Instead of *trail segment* k between vertex N_k and N_{k-1} being a straight line, we relax the requirement on *trail segment* k to be of length at most $(2 * \pi + 1) * 2^k$. By publishing information of P along more points, the $find$ path can be more straight towards C . An extreme case is when trail segment k is a full circle of radius 2^k centered at c_k and $Seg(N_k, N_{k-1})$. We call this variation of *Trail* the *find-centric Trail*.

5.2.1 Find-centric Trail. In this refinement, the $find$ procedure eschews exploring the circles (thus traversing only straight line segments) at the expense of the update procedure doing more work. This alternative data structure is used when objects are static or when object updates are less frequent than that of $find$ queries in a system. Let $trail_P$ for object P consist of segments connecting C, N_{mx}, \dots, N_1, p as described before and, additionally, let all points on the circles $Circ_k$ of center c_k and radius 2^k contain pointers to their respective centers, where $mx \geq k > 0$.

Starting at q , the $find$ path now is a straight line towards the center. If a circle with information about object P is intersected then, starting from this point, a line is drawn towards the center of the circle. Upon intersecting the immediate inner circle (if there is one), information about its respective center is found, with which a line is drawn to this center. Object P is reached by following this procedure recursively. We characterize the upper bound on the $find$ cost in the following Lemma, whose proof can be found in Appendix G.




Fig. 11. Find-centric Trail

LEMMA 5.3. *In find-centric Trail, when $b = 1$, the total cost of finding an object P at point p from object Q at point q is $16 * d_f$ where $d_f = dist(p, q)$.*

We note that when events are static, the optimal publish structure is much smaller than publishing along circular tracks. We have studied optimal publish structures for querying in a static context in a related work [Demirbas et al. 2006].

Summary: In this section, we presented two refinements of Trail that lead to a family of protocols. In the first refinement we altered the rate of updates. In the second refinement, we altered the amount of updates at each level. One can choose an appropriate parameterization depending on the expected rate of updates and finds in the network.

As an example, given the expected rate of updates and expected rate of find operations in the network, we can use the value of b in refinement 1 that minimizes the sum of update costs and find cost over a given interval of time. We can compare this cost with that of *find-centric Trail* and then choose the most appropriate parameterization. The find centric version of trail is especially beneficial when the rate of updates is much smaller than that of finds.

6. PERFORMANCE EVALUATION

In this section, we first evaluate the performance of *Trail* using simulations in *JProwler* [Vanderbilt University] and then describe an experimental evaluation of *Trail* on a network of 105 Mica2 motes. The goals of our simulation are: (1) to study the effect of routing stretch and discretization errors on the trail stretch factor, (2) to study the effect of uniform node failures on the performance of *Trail*, (3) to compare the average costs for *find* and *update*, as opposed to the upper bounds we derived earlier, and (4) to analyze the performance of *Trail* when scaled in the number of objects. Our simulation involves a network of 8100 Mica2 motes arranged in a 90×90 grid. We also experimentally evaluate the performance of *Trail* on a network of 105 Mica2 motes in the Kansei testbed [Arora et al. 2006]. This implementation has been used to support a distributed intruder interceptor tracking application where the goal of the interceptor is to catch the intruders as far away from an asset as possible. The goals of the experimental evaluation are to study the performance of *Trail* on a real network and analyze the performance under different scaling factors such as the number of objects in the system, the frequency of queries, and the speed of the objects in the network.

6.1 Simulation

Our simulation involves a 90×90 Mica2 mote network arranged on a grid. The center of the network is placed at one corner, thus essentially simulating one quadrant of the network. This setup lets us test the protocol over larger distances without update and find operations reaching the center. We use *JProwler* as our simulation platform with a Gaussian radio fading model. The mean interference range is a grid area of 5×5 square units. Packet transmission time is set at 40 ms. We implement geographic routing on a grid to route messages in the network. In the presence of failures we use a left hand rule to route around the failure [Karp and

Kung 2000]. We assume an underlying link maintenance layer because of which the list of *up* neighbors is known at each node.

Performance of update operations: We determine the number of messages exchanged for object updates over different distances when an object moves continuously in the network. We consider the unit grid separation, where each node has at most 4 communication neighbors. The number of neighbors may be lesser due to failures. We calculate the amortized cost by moving an object in different directions and then observing the cumulative number of messages exchanged up to each distance from the original position to update the tracking structure. The results are shown in Fig. 12(a). The jumps visible at distances 4 and 8 show the impact of the $\log(d)$ factor in the amortized cost. At these distances, the updates have to be propagated to a higher level. We also study the effect of uniform failures in the network on the increase in update costs. We consider fault percentages up to 20. We see from the figure that even with failures the average communication cost increases log linearly with distance. This indicates that the failures are handled locally.

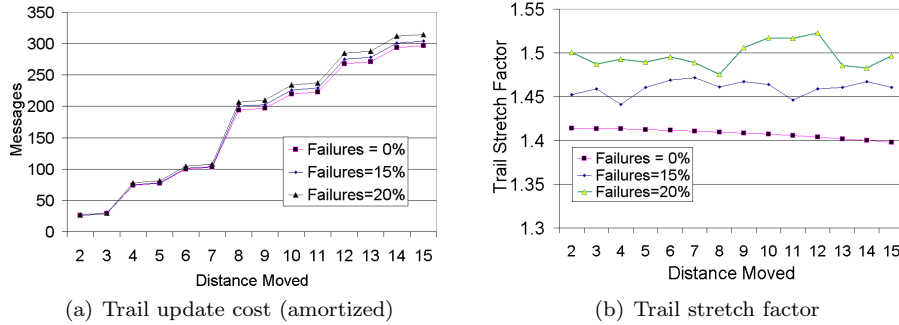


Fig. 12. Trail update costs and Trail stretch factor

Trail stretch factor: From Section 3, we note that in the continuous model, for an object P at distance d_{pC} from C , length of $trail_P$ is less than $1.2 * d_{pC}$. We now study the effect of routing overhead and the discretization factor on the length of the tracking structure that is created. We measure the trail length in terms of the number of hops along the structure. Fig. 12(b) shows the average ratio of distance from C to the length of the trail during updates over different distances from the original position. The parameter $b = 1$ in these simulations.

When the trail is first created, the trail stretch is equal to the routing stretch factor from C to the original location. In the absence of failures, we notice that the trail stretch is around 1.4 at updates of smaller distances and then starts decreasing. This can be explained by the fact the trail for an object starts bending more uniformly when the update is over a large distance. Even in the presence of failures, the trail stretch factor increases to only about 1.6 times the actual distance.

Performance of *find*: We first compare the average *find* costs of *Trail* with upper bounds derived. We study this in the presence of no interference, i.e. there is only one finder in the network. We fix the finder at distance 40 units from C . We vary

the distance of object being found from 2 to 16. We evaluate using the basic *find* algorithm with $b = 1$. In Fig. 13(a) and Fig. 13(b), we show the average number of messages and the average latency for the *find* operation respectively.

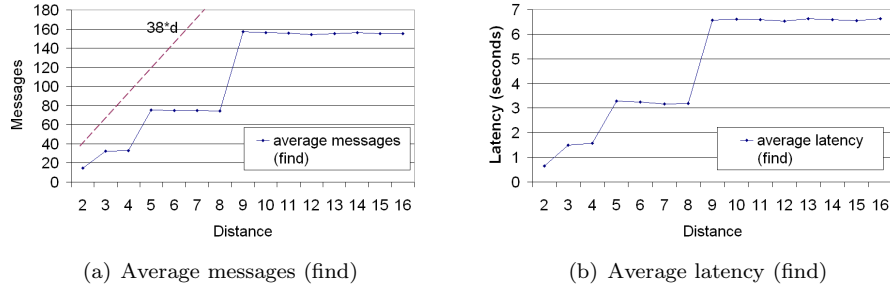


Fig. 13. Trail: find cost

The analytical upper bound $38*d$ (obtained from Lemma 4.1 for $b = 1$) is indicated using dotted lines in Fig. 13(a), and we see that the number of messages exchanged during *find* operations are significantly lower. When there is only one finder in the network, there is no interference. Therefore there are no re-transmissions except for the case when there is a loss due to probabilistic fading. Therefore the latencies are roughly equal to the number of messages times the message transmission time per hop. The jumps at distances 3, 5 and 9 are due to increase in levels of exploration at these distances. The results of the above simulations thus validate our theoretical bounds derived in Section 3 and 4.

Impact of interference: We now evaluate the effect of interference when multiple objects are present in the network. Note that *Trail* operates in a model where queries are generated in an asynchronous manner. We first evaluate the effect on *find* latency when objects are uniformly distributed across the network. We then evaluate the performance in a more severe environment where all the objects being found are collocated in the network.

In the presence of interference, messages are likely to be lost and we have implemented the following reliability mechanism to counter that. Forwarding a *find* message by the next hop is used as an implicit acknowledgment. *find* messages are retransmitted up to 4 times by each node. The interval to wait for an acknowledgment is doubled after every retransmission starting with 100 ms for the first retransmission. Note that 100 ms is a little more than twice the transmission time for each message. Also upon sensing traffic, the MAC layer randomly backs off the transmission within a window of 10 ms. The maximum number of MAC retries are set to 3.

In the first scenario we uniformly distribute the objects in a 50×50 area in the center of the quadrant being simulated. By distributing the objects in the middle of the quadrant being simulated, we avoid the decrease in *find* messages simply because an object is close to the boundary. We simulate 10, 20, 30, 40 and 50 objects in the network. We observe no significant increase in latency up to 30 objects. For the case of 40 and 50 objects in the network we observe increase in average latency

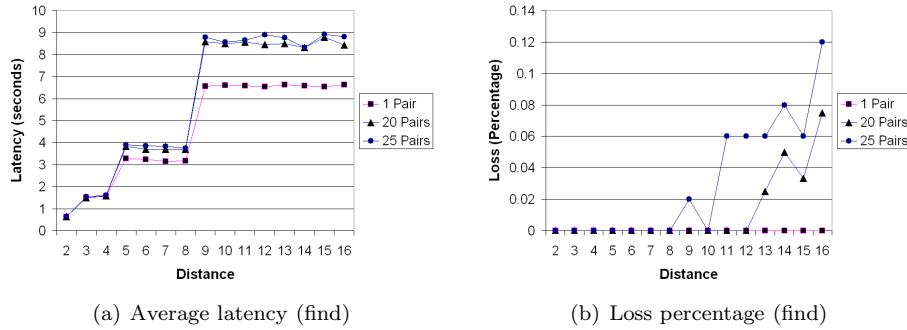


Fig. 14. Effect of interference on find cost

especially at larger distances. At larger distances, *find* messages from different objects interfere to a significant extent. Despite messages being re-transmitted up to 4 times, we also see losses during the find operation at distances greater than 12 units. The loss percentages at different distances are shown in Fig. 14(b)

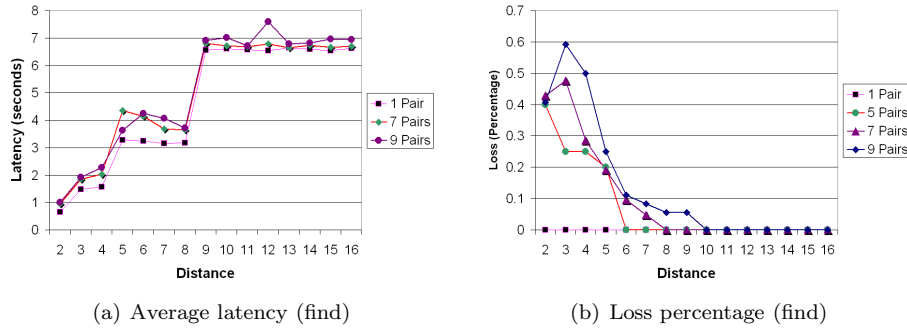


Fig. 15. Effect of interference on find cost: objects being found colocated

We now consider a more severe scenario where all the objects being found are at the same location. We compute the average latency for the find operation when objects issuing *find* query are uniformly distributed around this location, at different distances. As expected, we see from Fig. 15(a) and Fig. 15(b) that interference is severe at smaller distances. We see loss percentages as high as 60% when there are 9 pairs of objects at small distances.

Summary of evaluation: We observe from the above figures that *Trail* has a *find* time that grows linearly with distance. When scaled in the number of objects up to 50, with objects uniformly distributed in a 50×50 area and concurrently issuing queries, the query response time still does not increase substantially. However at a scale of 40 objects and distances of greater than 12 units we observe losses of around 10% during the find operation. This is because, at larger distances *find* messages from different objects interfere to a significant extent. In a potentially more severe scenario where all objects being found are at the same location and objects issuing *find* are distributed uniformly around that location, interference is

significant at smaller distances. We see loss percentages as high as 60% when there are 9 pairs of objects at small distances.

We now describe an experimental evaluation of *Trail* on a network of Mica2 nodes, summarize our observations and inferences and then discuss techniques to handle the effect of interference.

6.2 Experimental Evaluation on Real Network

We have implemented *Trail* for the special case of a long linear topology network for demonstrating an intruder interceptor tracking application. We have experimentally validated the performance of this version of *Trail* on 105 Mica2 nodes in the Kansei testbed [Arora et al. 2006] under different scaling factors such as the number of objects in the system, the frequency of queries, and the speed of the objects in the network. This implementation has been used to support a distributed intruder interceptor tracking application where the goal of the interceptor is to catch the intruders as far away from an asset as possible. This application was demonstrated at Richmond Field Station in Berkeley in 2005 as part of the DARPA NEST Program. In this section, we describe the results of these experiments.

Experimental setup: We use a network of 105 XSM-Stargate pairs in a 15×7 grid topology with 3 ft spacing in the Kansei testbed. The XSMs are a Mica2 family of nodes with the same Chipcon radio but with additional sensors mounted on the sensor board. The XSMs are connected to Stargate via serial port and the Stargates are connected via Ethernet in a star topology to a central PC. We are able to adjust the communication range by adjusting the power level and the XSMs can communicate reliably up to 6 ft at the lowest power level but the interference range could be higher. *Trail* operates asynchronously with no scheduling to prevent collisions. Hence, we implement an implicit acknowledgement mechanism at the communication layer for per hop reliability. The forwarding of a message acts as acknowledgment for the sender. If an acknowledgment is not received, then messages are retransmitted up to 3 times.

Object traces: We now describe how the object motion traces are obtained. Nodes were deployed in a grid topology with 10 m spacing at Richmond field station. Sensor traces were collected for objects moving through this network at different orientations. Based on these traces, tracks for the objects are formed using a technique described in [Arora et al. 2004]. These tracks are of the form (timestamp, location) on a $140\text{m} \times 60\text{m}$ network. These object tracks are then converted to tuples of the form (id, timestamp, location, grid position) where grid position is the node closest to the actual location on the 15×7 network and id is a unique identifier for each object. These detections are injected into the XSM in the testbed corresponding to the grid position via the Stargate at the appropriate time, using the inject framework in *Kansei*. Thus, using real object traces collected from the field and using the injector framework, we emulate the object detection and association layer to evaluate the performance of our network services.

Parameters: We evaluate the performance of *Trail* under different scaling factors such as increasing number of objects, higher speed of objects and higher query frequency in terms of the reliability and latency of the service. We run *Trail* with 2,

4, 6 and 10 mobile objects, in pairs. One object in each pair is the object issuing *find* query and the other object is the object being located. In each of these scenarios, we consider query frequency of 1 Hz, 0.5 Hz, 0.33 Hz and 0.25 Hz. The object speed affects the operation of *Trail* in terms of the rate at which *grow* and *clear* messages are generated. We consider 3 different object update rates, one in which objects generate an update every 1 second, every 2 seconds and every 3 seconds. Considering that the object traces were collected with humans walking across the network acting as objects with average speed of about 1 m/s, object update rates of 1 Hz and 0.5 Hz enable a tracking accuracy of 1m and 2m respectively. Note that each update can generate multiple *grow* and *clear* messages.

In the 4, 6 and 10 objects scenario, we consider a likely worst case distribution of the objects where all objects issuing *find* and all objects being found are within communication range. Moreover, as optimal pursuit control requirements suggest [Cao et al. 2006], the query frequencies depend on relative locations and are lesser when objects are far apart, but we consider all objects issuing queries at the same frequency. If the replies are not received before the query period elapsed, then the message is considered lost. The loss percentages are based on 100 *find* queries at every distance and the latencies are averaged over that many readings.

Scaling in number of objects: Fig. 16 shows the latency and loss for *find* operations as the number of objects increases with query frequency fixed at 0.33 Hz and object updates fixed at 0.5 Hz. Fig. 17 shows the latency and loss for *find* operations as the number of objects increases with query frequency fixed at 0.5 Hz and object updates fixed at 0.5 Hz.

Scaling in query frequency: Here we analyze how the latency and reliability of *Trail* are affected as the query frequency increases. In Fig. 18, we show the reliability and latency of *Trail* with 6 objects under query frequencies of 1 Hz, 0.5 Hz, 0.33 Hz and 0.25 Hz, with object update rate of 0.5 Hz.

Scaling in object speed: Fig. 19 shows the latency and loss for *find* operations with increasing object speeds that generate updates at 0.33 Hz, 0.5 Hz and 1 Hz. The query frequency is 0.5 Hz and the number of objects is 6.

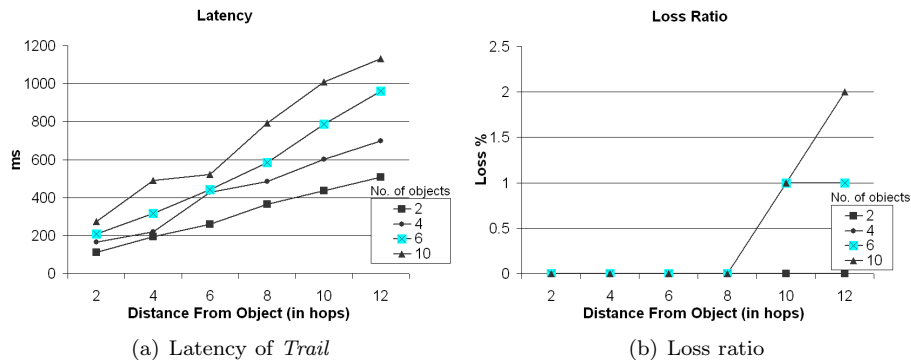


Fig. 16. Scaling in number of objects (query frequency 0.33 Hz, object update 0.5 Hz)

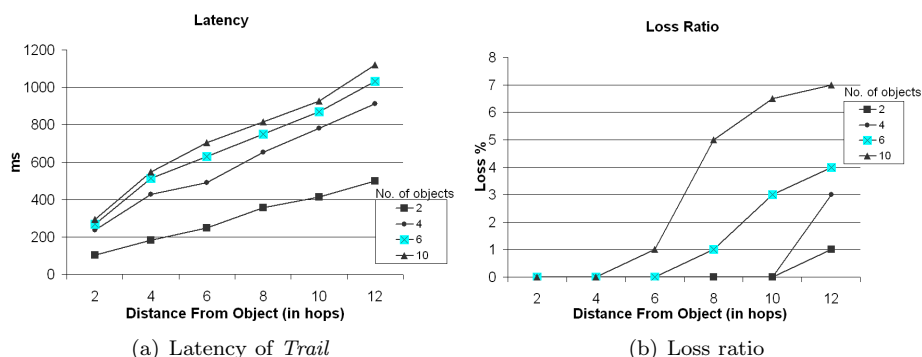


Fig. 17. Scaling in number of objects (query frequency 0.5 Hz, object update 0.5 Hz)

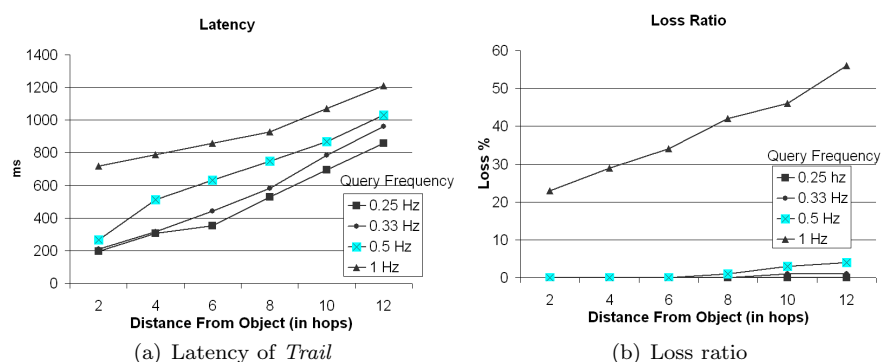


Fig. 18. Scaling in query frequency (6 objects, object update rate 0.5 Hz)

Summary of experimental evaluation: We observe from the above figures that *Trail* offers a query response time that grows linearly with the distance from an object. *Trail* operates in an environment where objects can generate updates and queries asynchronously and in such an environment, interference increases the response time. From our experiments, we observe that query latency and loss percentages increase with number of objects and speed of objects but the loss ratio is not severe. As seen in Fig. 16, scaling the number of objects up to 10 yields a loss rate of up to 7% with a query frequency 0.5 Hz and an object update rate of 0.5 Hz. As seen in Fig. 19, scaling the object speeds to generating 1 update per second results in a loss rate of up to 7% even with 6 objects in the system and query frequency of 0.5 Hz. Increasing query frequencies has a more severe impact on loss percentages especially with more objects in the network. In Fig. 18, we notice that with 6 objects in the network, loss increases substantially as the query frequency becomes 1 Hz; this happens due to higher interference leading to congestion.

Handling interference: To tolerate network interference, spatial and temporal correlations that exist in the application can be exploited in the following way. The rate at which information is needed by the pursuer is known. The network service can be notified of the next instant at which the state of the evader is needed and the

location where the query results need to be sent. The query results can then arrive *just in time*. Advance knowledge about the query and the deadline can be used to decrease the interference in the network when multiple pursuers query about evaders in the network and more so when the objects are densely located. Another possibility to deal with interference is a synchronous *push* version of the network tracking service where snapshots of objects are published to subscribers in a distance sensitive manner thereby avoiding interference. By the same token, applications should be aware of other extreme conditions (in terms of object number and speed) for effectively using the service. For example, applications may compensate for losses by increasing their query frequency, but this should account for extreme scenarios where the increased frequency itself results in higher interference.

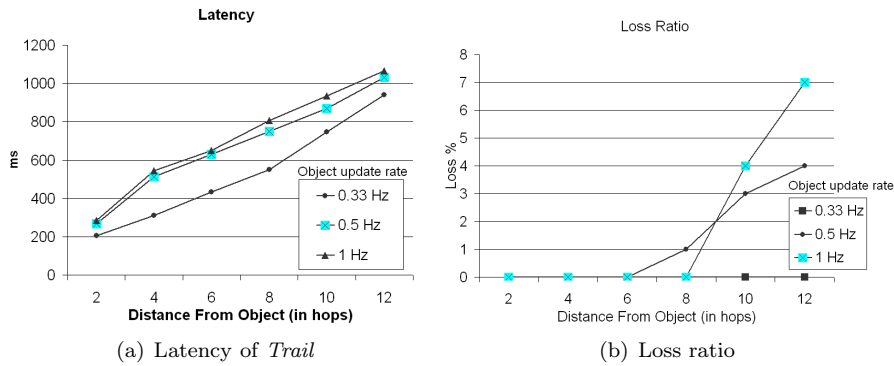


Fig. 19. Scaling in object speed (6 objects, query frequency 0.5 Hz)

7. DISCUSSION

In our solution, we made some design decisions like choosing a single point to terminate tracks from all points in the network and avoiding hierarchy in maintaining the tracks. In this section, we analyze these aspects of our solution and compare them with other possible approaches. We find that by avoiding hierarchy, we do not need to partition the network into clusters and maintain these clusters, we can be more locally fault-tolerant and we can obtain tighter tracks for any object. We also formally define the notion of terminating points, differentiate those from clusterheads of a hierarchy, and analyze the effect of more terminating points on the maximum *find* cost and maximum *update* cost in the network.

Terminating points vs clusterheads: There are some hierarchy based solutions [Demirbas et al. 2004; Funke et al. 2006] for the problem of object tracking in a distance sensitive manner, where the network is hierarchically partitioned into clusters and information of objects is maintained at clusterheads at each level. Even in these solutions, information about an object is published across the network to local clusterhead(s) at each level in the hierarchy, all the way up to one or more clusterheads at the highest level in the hierarchy. We call these points at the highest level of the hierarchy as terminating points.

Formally, a terminating set τ is a smallest set of points such that tracks of objects from every location in the network pass through at least one point in τ . The cardinality of a set τ is denoted as μ_τ . There can be one or more terminating sets, each with one or more terminating points. In *Stalk* [Demirbas et al. 2004] there is a unique clusterhead at the highest level, thus there is a single terminating set with a single terminating point. In *LLS* [Abraham et al. 2004] and *DSIB* [Funke et al. 2006], there are multiple clusterheads at the highest level. Thus there are multiple terminating sets, each with one terminating point. Tracks from every point in the network pass through each of those clusterheads. Thus each clusterhead at the highest level constitutes a terminating set by itself. In *Trail* there is a unique terminating set with a unique terminating point, namely C .

It is in the process of maintaining tracks from a terminating point that we have avoided hierarchy in *Trail*. In hierarchy based solutions, to maintain tracks and to answer queries, tracks from terminating points necessarily pass through these clusterheads, where as *Trail* avoids hierarchy by determining anchors for the tracking paths on-the-fly based on the motion of objects.

Merits of avoiding hierarchy: By avoiding hierarchical solutions we do not need either a distributed clustering service that partitions the network into clusterheads at different levels and maintains this clustering or a special (maybe centralized) allocation of infrastructure nodes. By avoiding a hierarchy of such special nodes *Trail* is also more locally fault tolerant. For example in the case of a *find* operation, failure to retrieve information from an information server at a given level would require the *find* to proceed to a server at the higher level [Funke et al. 2006]. This is particularly expensive at higher levels of the hierarchy. On the other hand in *Trail* a *find* operation redirects around a *hole* created by failed nodes using routing techniques such as the left hand rule [Karp and Kung 2000] and such faults can be handled, in a sense, proportional to the size of the fault. Similarly, tracks to existing objects can be repaired more efficiently. As the number of failures increase, there is only a proportional increase in *find* and *update* costs.

Moreover, avoiding hierarchies allows for minimizing the length of tracking paths given a terminating point. We analytically compare the performance of *Trail* with that of other hierarchy based solutions for tracking objects in Section 8 and we observe that *Trail* is more efficient than other solutions. *Trail* has about 7 times lower *update* costs at almost equal *find* costs. By using a tighter tracking structure, we are also able to decrease the upper bound *find* costs at larger distances and thereby decrease the average find cost across the network.

Choice of terminating points: In Appendix A, we have formally analyzed the choice of a unique terminating point for tracks from all points in the network and the tradeoffs associated with multiple terminating sets and multiple terminating points per set in terms of the maximum *find* and *update* costs in the network. We provide a summary of our analysis here.

We consider 2 cases: a single terminating set with multiple terminating points and multiple terminating sets each with one terminating point.

Single terminating set with multiple terminating points: Intuitively, there exists a possibility of decreasing the maximum track length in the network by dividing

the network into regions and having a local terminating point per region. The maximum track length and therefore the maximum *update* cost in the network thus depends on the size of the largest region. We are faced with the question of how small can these regions be. We show in Appendix A, that in order to maintain *find* distance sensitivity, the diameter of the largest region can only be a constant order less than the diameter of the network, i.e., at least $\Omega(N)$ in a $N \times N$ network. Thus, there can be only a constant order of cost decrease compared to having only one terminating point. Moreover, decreasing the maximum *update* cost by dividing the network into smaller regions results in proportionate increase in the maximum *find* cost. This is because if from any finder location, a track belonging to an object in any location in the network is to be found, then it must be the case that the *find* trajectory contains all points in the terminating set, thereby increasing the worst case *find* cost.

Multiple terminating sets: If there are multiple terminating sets then it is sufficient for *find* to traverse the terminating point in any such set. In this case there is a likelihood of decreasing the maximum *find* cost when compared to having only set of terminating points because tracks can be *found* by reaching a terminating point in any of the terminating sets. The maximum *find* cost in the network depends on the size of the largest region. However, we show in Appendix A that to maintain *update* distance sensitivity, the size of the largest region has to be at least $\Omega(N)$ in a $N \times N$ network. Also when the number of terminating sets is greater than 1, *update* has to traverse the terminating point in all terminating sets in the worst case. Thus the maximum *update* cost in the network increases.

Maintaining a track with respect to local terminating points could be advantageous if it is more likely that querying object and the object being found are closer. Thus, a *find* will never run into the scenario of having to traverse all regions in the network. Similarly, maintaining a track with respect to multiple terminating point sets could be advantageous if objects are likely to move within bounded regions within a network. In this paper we consider all distances between querying object and tracked object to be equally likely and do not restrict mobility of the objects. Hence we consider only the case where there is a unique terminating set with a single terminating point, namely C .

We note that it is also feasible to select a different terminating point for different *types* of objects. In this paper we describe how to maintain tracks for objects with respect to one terminating point and guarantee *find* and *update* distance sensitivity. A different terminating point can be chosen for each type of object based on hash functions and each *type* of finder object can choose the respective terminating points as a worst case landmark; but this concept is orthogonal to that of maintaining tracks with respect to a given terminating point.

Tolerating failure of the terminating point: In Section 4 we have shown that in *Trail*, we can locally tolerate the failure of even C . Our protocol actions are such that when C fails or a set of nodes in a contiguous region around C fail, track for an object will terminate at any node that is closest to the boundary formed by the *hole*. Thus during a *find* operation, a redirection rule as in GPSR is bound to intersect the track and once again the faults are tolerated locally.

Memory efficiency: Note that the tracks maintained in *Trail* only contain pointers to the current location of an object and not the state information of the object. Thus *Trail* is memory efficient even at C .

Handling bottleneck at C : The notion of C as a terminating point can be simply extended to construct a circle of constant radius δ around the center of the network and define each point on this circle as a terminating set. Thus tracks from any point in the network pass through all points on this circle. *find* can intersect any point on this circle to obtain a pointer to the object. The maximum update cost increases by a factor of δ and the maximum *find* cost decreases by a factor of δ , and still maintaining distance sensitivity. The responsibility of handling queries is now distributed evenly around C .

8. RELATED WORK

In this section, we discuss related work and also compare the performance of *Trail* with other protocols designed for distance sensitive tracking and querying.

Tracking: As mentioned earlier, mobile object tracking has received significant attention [Awerbuch and Peleg 1995; Demirbas et al. 2004; He et al. 2006] and we have focused our attention on WSN support for tracking. Some network tracking services [Dolev et al. 1995] have non-local updates, where update cost to a tracking structure may depend on the network size rather than distance moved. There are also solutions such as [Awerbuch and Peleg 1995; Demirbas et al. 2004; Abraham et al. 2004] that provide distance sensitive updates and location.

Locality Aware Location Services (*LLS*) [Abraham et al. 2004] is a distance sensitive location service designed for mobile ad-hoc networks. In *LLS*, the network is partitioned into hierarchies and object information is published in a spiral structure at well known locations around the object, thus resulting in larger update costs whenever an object moves. The upper bound on the update cost in *LLS* is $128 * d_m * \log(d_m)$, where d_m is the distance an object moves, as opposed to the $14 * d_m * \log(d_m)$ cost in *Trail*; the upper bounds on the *find* cost are almost equal. Moreover, as seen in Section 5, we can further reduce the upper bound on the *find* cost at higher levels in *Trail*.

The *Stalk* protocol [Demirbas et al. 2004] uses hierarchical partitioning of the network to track objects in a distance sensitive manner. The hierarchical partitioning can be created with different dilation factors ($r \geq 3$). For $r = 3$ and 8 neighbors at each level, at almost equal *find* costs, *Stalk* has an upper bound update cost of $96 * d * \log(d)$. This increase occurs because of having to query neighbors at increasing levels of the partition in order to establish *lateral* links for distance sensitivity [Demirbas et al. 2004].

Both *Stalk* and *LLS* use a partitioning of the network into hierarchical clusters which can be complex to implement in a WSN, whereas *Trail* is cluster-free. Moreover, in *Stalk*, the length of the tracking structure can span the entire network as the object keeps moving and, in *LLS*, the information about each object is published in a spiral structure across the network. In comparison, *Trail* maintains a tighter tracking structure (i.e., with more direct paths to the center) and is thus more efficient and locally fault-tolerant.

	Update Cost	Find Cost	Size
Trail	$14*d*logd$	$38*d$	$1.2*d_c$
LLS	$128*d*logd$	$36*d$	$16*D$
Stalk	$96*d*logd$	$39*d$	$3*D$
Awerbuch Peleg	$O(d*logd*logN)$	$16*d*log(2N)$	$4*D$

D – n/w Diameter ; N – No. of Nodes ; d_c – Distance from C

Fig. 20. Trail: analytical comparison

In [Awerbuch and Peleg 1995], a hierarchy of regional directories is constructed and the communication cost of a find for an object d_f away is $O(d_f * \log(N))$ and that of a move of distance d_m is $O(d_m * \log(D) * \log(N))$ (where N is the number of nodes and D is the network diameter). A topology change, such as a node failure, however, necessitates a global reset of the system since the regional directories depend on a non-local clustering program that constructs sparse covers.

Querying and storage: Querying for events of interest in WSNs has also received significant attention [Intanogonwiwat et al. 2003; Ratnasamy et al. 2002; Liu et al. 2004] and some of them focus on distance sensitive querying. We note that *Trail*, specifically the *find-centric* approach can also be used in the context of static events.

In [Liu et al. 2004], a balanced push-pull strategy is proposed that depends on the query frequency and event frequency; given a required query cost, the advertise operation is tuned to do as much work as required to satisfy the querying cost. In contrast, *Trail* assumes that query rates depend on each subscriber (and potentially on the relative locations of the publisher and subscriber), and it also provides distance sensitivity during find and move operations, which is not a goal of [Liu et al. 2004]. In directed diffusion [Intanogonwiwat et al. 2003], a tree of paths is created from all objects of interest to the tracker. All these paths are updated when any of the objects move. Also, a controller initiated change in assignment would require changing the paths. By way of contrast, in *Trail*, we impose a fixed tracking structure, and tracks to all objects are rooted at one point. Thus, updates to the structure are local. Rumor routing [Braginsky and Estrin 2002] is a probabilistic algorithm to provide query times proportional to distance; the goal of this work is not to prove a deterministic upper bound. Moreover, its algorithm does not describe how to update existing tracks locally and yet retain distance sensitive query time when objects move.

Geographic Hash tables [Ratnasamy et al. 2002] is a lightweight solution for the in-network-querying problem of static events. The basic *GHT* is not distance sensitive since it can hash the event information to a broker that is far away from a subscriber. The distance sensitivity problem of GHT can be alleviated to an extent by using geographically bounded hash functions at increasing levels of a hierarchical partitioning as used in DIFS protocol. Still, attempting such a solution suffers from a multi-level partitioning problem: a query event pair nearby in the network might be arbitrarily far away in the hierarchy. However, we do note that *GHT* provides load balancing across the network, especially when the types of events are known and this is not the goal of *Trail*.

Distance Sensitive Information Brokerage [Funke et al. 2006] protocol performs a hierarchical clustering of the network and information about an event is published to neighboring clusters at each level. Using *Find-centric Trail* we can query information about static events in a distance sensitive manner. We also note that when

events are static, the optimal publish structure is much smaller than publishing along circular tracks. We have studied optimal publish structures for querying in a static context in a related work [Demirbas et al. 2006].

Spatio-temporal query services Motivated by a class of applications in which mobile users are interested in continuously gathering information in real time from their vicinities, a network data service called *spatio temporal query* has been proposed in [Lu et al. 2005]. The spatial constraint for the network service comes from an energy efficiency point of view; only nodes relevant to a query area should be involved in contributing the query result. The temporal constraint is due to a requirement on data freshness for the query results. An approximate motion model for the mobile user is assumed. Specifically the motion can be predicted over a small interval of time. The query area at any time is a function of the current location of the mobile user. The key difference in *Trail* is that the query area in consideration is the entire network as opposed to a function of the querier location as in [Lu et al. 2005].

We note that, when *Trail* is used in the context of a distributed pursuer evader game, spatial and temporal correlations that exist in the application can be exploited using ideas presented in [Lu et al. 2005] to improve the performance of the application and the network. The rate at which information is needed by the pursuer is known. The network service can be notified of the next instant at which the state of the evader is needed and the location where the query results need to be sent. The query results can then arrive *just in time*. Constraints on evader speed can be exploited as follows. Subsequent queries for evader location can originate from the previous evader location and the results can be routed back to the pursuer.

9. CONCLUSIONS AND FUTURE WORK

We have presented *Trail*, a family of protocols for distance sensitive distributed object tracking in WSNs. *Trail* avoids the need for hierarchical partitioning by determining anchors for the tracking paths on-the-fly, and is more efficient than other hierarchy based solutions for tracking objects: it allows 7 times lower updates costs at almost equal *find* costs and can tolerate faults more locally as well.

Importantly, *Trail* maintains tracks from object locations to only one terminating point, the center of the network. Moreover, since its tracks are *almost* straight to the center with a stretch factor close to 1, *Trail* tends to achieve the lower bound on the total track length. By using a tight tracking structure, *Trail* is also able to decrease the upper bound *find* costs at larger distances and thereby decrease the average find cost across the network.

Trail is a family of protocols and we have shown that refinements of the basic *Trail* protocol are well suited for different network sizes and query frequency settings. We have validated the distance sensitivity and fault tolerance properties of *Trail* in a simulation of 90×90 network using *JProwler*. We have also successfully implemented and tested the *Trail* protocol in the context of a pursuer evader application for a medium size (over 100 node) mote network.

Trail operates in an environment where objects can generate updates and queries

asynchronously. We note that in such an environment, due to the occurrence of collisions, there can be an increase in the message complexity for querying and updates especially when the objects are densely located in the network. As future work, we are considering a *push* version of the network tracking service where snapshots of objects are published to subscribers in a distance sensitive manner, both in time and information, in order to increase the reliability and energy efficiency of the service when the density of objects in the network is high.

We also note that spatial and temporal correlations that exist in the application can be exploited to improve the performance of the application and the network. Two such examples are as follows. (1) The rate at which information is needed by the pursuer is known. The network service can be notified of the next instant at which the state of the evader is needed and the location where the query results need to be sent. The query results can then arrive *just in time*. Advance knowledge about the query and the deadline can be used to decrease the interference in the network when multiple pursuers query about evaders in the network and more so when the objects are densely located. (2) Constraints on evader speed can be exploited as follows. Subsequent queries for evader location can originate from the previous evader location and the results can be routed back to the pursuer. Thus the energy efficiency of the network can be improved.

REFERENCES

- ABRAHAM, I., DOLEV, D., AND MALKHI, D. 2004. LLS: A locality aware location service for mobile ad hoc networks. In *Joint Workshop on Foundations of Mobile Computing (DIALM-POMC)*. ACM, 75–84.
- ARORA, A., DUTTA, P., BAPAT, S., KULATHUMANI, V., ZHANG, H., AND NAIK, V. 2004. A line in the sand: A wireless sensor network for target detection, classification, and tracking. *Computer Networks, Special Issue on Military Communications Systems and Technologies* 46, 5 (July), 605–634.
- ARORA, A., ERTIN, E., RAMNATH, R., NESTERENKO, M., AND LEAL, W. 2006. Kansei: A high fidelity sensing testbed. *IEEE Internet Computing* 10, 2 (March), 35–47.
- ARORA, A. AND RAMNATH, R. 2004. Exscal: Elements of an extreme wireless sensor network. In *The 11th International Conference on Embedded and Real-Time Computing Systems and Applications*. 102–108.
- AWERBUCH, B. AND PELEG, D. 1995. Online tracking of mobile users. *Journal of the Association for Computing Machinery* 42, 1021–1058.
- BAPAT, S., KULATHUMANI, V., AND ARORA, A. 2005. Analyzing the yield of exscal, a large scale wireless sensor network experiment. In *13th IEEE International Conference on Network Protocols*.
- BRAGINSKY, D. AND ESTRIN, D. 2002. Rumor routing algorithm for sensor networks. In *First ACM Workshop on Wireless Sensor Networks and Applications*. ACM, 22–31.
- CAO, H., ERTIN, E., KULATHUMANI, V., SRIDHARAN, M., AND ARORA, A. 2006. Differential games in large scale sensor actuator networks. In *Information Processing in Sensor Networks (IPSN)*. ACM, 77–84.
- DEMIRBAS, M., ARORA, A., AND KULATHUMANI, V. 2006. Glance: A Light Weight Querying Service For Sensor Networks. In *International Conference on Principles of Distributed Systems (OPODIS)*. 242–257.
- DEMIRBAS, M., ARORA, A., NOLTE, T., AND LYNCH, N. 2004. A hierarchy-based fault-local stabilizing algorithm for tracking in sensor networks. In *International Conference on Principles of Distributed Systems (OPODIS)*. 143–162.
- DOLEV, S., PRADHAN, D., AND WELCH, J. 1995. Modified tree structure for location management in mobile environments. In *INFOCOM*. 530–537.

- FUNKE, S., GUIBAS, L., NGUYEN, A., AND ZHANG, Y. 2006. Distance sensitive information brokerage in sensor networks. In *International Conference on Distributed Computing in Sensor Systems DCOSS*. Springer-Verlag, 234–251.
- HE, T., KRISHNAMURTHY, S., AND STANKOVIC, J. 2006. Vigilnet: an integrated sensor network system for energy-efficient surveillance. *ACM Transactions on Sensor Networks* 2, 1, 1–38.
- INTANOGONWIAT, C., GOVINDAN, R., ESTRIN, D., HEIDEMANN, J., AND SILVA, F. 2003. Directed diffusion for wireless sensor networking. *IEEE Transactions on Networking* 11, 1, 2–16.
- KARP, B. AND KUNG, H. T. 2000. Greedy perimeter stateless routing for wireless networks. In *Proceedings of Mobile Computing and Networking, MobiCom*. 243–254.
- LIU, X., HUANG, Q., AND ZHANG, Y. 2004. Combs, needles, haystacks: Balancing push and pull for discovery in large-scale sensor networks. In *ACM Sensys*. ACM, 122–133.
- LU, C., XING, G., CHIPARA, O., FOK, C.-L., AND BHATTACHARYA, S. 2005. A Spatiotemporal Query Service for Mobile Users in Sensor Networks. In *ICDCS*. 381–390.
- RATNASAMY, S., KARP, B., YIN, L., YU, F., ESTRIN, D., GOVINDAN, R., AND SHENKER, S. 2002. GHT: A geographic hash table for data-centric storage. In *First ACM Workshop on Wireless Sensor Networks and Applications (WSNA)*. ACM, 12–21.
- SARKAR, R., ZHU, X., AND GAO, J. 2006. Double rulings for information brokerage in sensor networks. In *International Conference on Mobile Computing and Networking (MobiCOM)*. ACM, 286–297.
- SHIN, J., GUIBAS, L., AND ZHAO, F. 2003. A distributed algorithm for managing multi-target identities in wireless ad hoc networks. In *International Workshop on Information processing in Sensor Networks IPSN*. 223–238.
- SINOPOLI, B., SHARP, C., SCHENATO, L., SCHAFFERT, S., AND SASTRY, S. 2003. Distributed control applications within sensor networks. In *Proceedings of the IEEE*. Vol. 91. 1235–46.
- VANDERBILT UNIVERSITY. JProwler. <http://www.isis.vanderbilt.edu/Projects/nest/jproowler/index.html>.

A. TECHNICAL NOTE: ON TERMINATING POINTS FOR TRACKING MOBILE OBJECTS

In this section, we define the notion of *terminating points* for schemes that track mobile objects in a distance sensitive manner in terms of *update* and *find*. We also analyze the tradeoffs with respect to the choice of terminating points.

Let O be a set of mobile objects in a network of size $N \times N$. Mobile objects are of two types: *finder* objects (O_f) and *mover* objects (O_m). Thus O is a union of disjoint sets O_f and O_m . Let p denote the location of any object P . Let *Tracker* be a distance sensitive tracking scheme. *Tracker* maintains a track $track_P$ for every object P that belongs to O_m . $track_P$ is a set of points that contain information pertaining to P . This information could be the actual state of P or simply a pointer following which leads to the actual state of P . The length of $track_P$ is the length of the shortest curve connecting all points in $track_P$. *Tracker* offers two functions: $find(P, Q)$, that returns state of P to Q , where Q belongs to O_f and P belongs to O_m ; and $move(P, p', p)$ that updates $track_P$ when P moves from p' to p . *Tracker* satisfies property F (*find* distance sensitivity) and property U (*update* distance sensitivity) stated below:

Definition A.1 find distance sensitivity. *Tracker* satisfies property F if the cost of $find(P, Q)$ grows linearly with $dist(p, q)$.

Definition A.2 update distance sensitivity. *Tracker* satisfies property U if the cost of $move(P, p', p)$ grows linearly with $dist(p, p')$.

Note that in this document we consider only *discrete* moves of the objects. When the motion of the object is continuous, a subset of *Tracker* schemes may have a property that the cost of move is proportional to the distance of move in an amortized sense.

From here on we assume that there is only one *mover* in the network and track for that object is represented as *track* and we drop the identity subscript. We refer to *track* for the *mover* at location x in the network simply as *track* for point x .

Definition A.3 Terminating Set τ . A terminating set τ in *Tracker* is a smallest set of points such that *track* for every point in the network passes through at least one point in τ .

The cardinality of a terminating set τ is denoted as μ_τ . The points contained in a terminating set are called as terminating points of that set. Note that there could also be multiple terminating sets in the network with each set containing an equal but any number of points. In other words there are multiple smallest sets of points such that *track* for every point in the network passes through at least one point in each of those sets. For example in the case of *Trail* and *Stalk* [Demirbas et al. 2004], there is only terminating set and this set has exactly one point, namely the center of the network. Thus $\tau = \{C\}$. In *DSIB* [Funke et al. 2006] and *LLS* [Abraham et al. 2004], each set containing one publish location at the highest level constitutes a terminating set. Thus in those schemes there are multiple terminating sets but the cardinality of each set is 1. In the simple horizontal vertical double ruling scheme,

each horizontal line is a terminating set because every vertical track passes through at least one point in each of those sets.

LEMMA A.4. *A terminating set τ cannot be an empty set.*

PROOF. *track* for each point in the network is at least equal to the point itself. Thus τ cannot be empty. \square

In *Trail*, we have chosen a unique terminating set consisting of a unique terminating point. We now analyze the tradeoffs involved when the terminating set is not unique and when a terminating set contains more terminating points. We consider 2 cases: a single terminating set with multiple terminating points and multiple terminating sets each with one terminating point.

Case 1: $\mu_\tau \geq 1$ We note that it is possible to decrease the maximum track length in the network by dividing the network into regions and tracks being maintained with respect to a terminating point in each region. (The disproportionate updates that can be caused when an object keeps switching between boundaries can be avoided by making the regions overlap.) The question then arises as to how small the regions can be and yet maintain *distance sensitivity*. We now analyze the limits for decreasing the track length and its effect on maximum *find* cost.

Given any location f of a finder, let L_f denote the length of the *find* trajectory traversed from the finder location, after which the *track* for any point in the network is found. Thus L_f denotes the worst case *find* cost from location f . Let \hat{L}_f denote the maximum value of L_f in the network.

LEMMA A.5. *For F to hold, $\hat{L}_f = O(N)$.*

PROOF. Note that the maximum distance between any two points in the network is $O(N)$. The result follows. \square

LEMMA A.6. *\hat{L}_f is at least equal to the length of traversing all points in τ .*

PROOF. From the definition of a terminating set, τ is a minimum set of points through which tracks from all points pass through. \square

Using Lemma A.6, we state the following Theorem.

THEOREM A.7. *The maximum find cost in the network is minimized when $\mu_\tau = 1$.*

PROOF. In the worst case, a *find* trajectory has to traverse all terminating points in the terminating set. Compared to any configuration of terminating points when $\mu_\tau > 1$, there exists a configuration of the terminating point when $\mu_\tau = 1$, such that \hat{L}_f is lower. \square

Using Lemmas A.5 and A.6, we get the following Lemma.

LEMMA A.8. *All terminating points in τ must be traversable in $O(N)$.* \square

The above Lemma imposes a lower bound the maximum size of the regions in the following way. Let a region \mathcal{R}_t be a set of points that choose t as a terminating point. Let $\rho_{\mathcal{R}}$ denote the distance of the farthest point in the region from t . Let $\hat{\rho}_{\mathcal{R}}$ denote the maximum distance from a terminating point in any region of the network. We state the following Theorem.

THEOREM A.9. *In order to preserve F , $\hat{\rho}_{\mathcal{R}} = \Omega(N)$.*

PROOF. Recall that all terminating points in τ must be traversed in $O(N)$. If $\hat{\rho}_{\mathcal{R}} < \Omega(N)$, then the maximum diameter of any region in the network is less than $\Omega(N)$ in the $N \times N$ network and therefore all terminating points in τ cannot be traversed in $O(N)$. \square

Summary: The maximum track length in the network depends on the size of the largest region. We note the maximum track length can be decreased by dividing the network into regions and having local terminating points per region. However, the size of the largest region can only be a constant order less than the diameter of the network. For instance dividing the network into infinitesimally small regions or even regions of size $O(\sqrt{N})$ will violate F . Also when $\mu_{\tau} > 1$, *find* has to traverse all the points in τ in the worst case. Thus the maximum *find* cost in the network increases. \square

Case 2: Multiple terminating sets: Now we consider the case where there are multiple terminating sets each with one terminating point and analyze the tradeoffs in *find* and *update* cost. If there are multiple terminating sets then it is sufficient for *find* to traverse the terminating point in any such set. In this case there is a likelihood of decreasing the maximum *find* cost when compared to having only set of terminating points because tracks can be *found* by reaching a terminating point in any of the terminating sets. An example of this is the case where all tracks pass through a common set of points as opposed to just one common point. Thus the cardinality of each terminating point set is 1 because all tracks pass through each point, but there are multiple such points.

Similar to case 1, we can show that the maximum *find* cost can be decreased by dividing the network into regions and having one terminating set per region. We state the following Theorems whose proofs are similar to that of case 1.

THEOREM A.10. *The maximum update cost in the network is minimized when the number of terminating sets is 1. \square*

THEOREM A.11. *In order to preserve U , $\hat{\rho}_{\mathcal{R}} = \Omega(N)$, where $\hat{\rho}_{\mathcal{R}}$ is the maximum distance of any point in the network from the terminating point in its local terminating set. \square*

Summary: The maximum *find* cost in the network depends on the size of the largest region. We note the maximum *find* length can be decreased by dividing the network into regions and having local terminating points per region. However, the size of the largest region can only be a constant order less than the diameter of the network. For instance dividing the network into infinitesimally small regions or even

regions of size $O(\sqrt{N})$ will violate F . Also when the number of terminating sets is greater than 1, *update* has to traverse the terminating point in all terminating sets in the worst case. Thus the maximum *update* cost in the network increases. \square

Choice of unique terminating point: Maintaining a track with respect to local terminating points and a single terminating set could be advantageous if it is more likely that querying object and the object being found are closer and therefore it is unlikely that all terminating points in the set have to be traversed. Similarly maintaining a track with respect to multiple terminating sets could be advantageous if objects are likely to move within bounded regions within a network. In this paper we consider all distances between querying object and tracked object to be equally likely and do not restrict mobility of the objects. Hence we consider only the case where there is a unique terminating point, namely C .

Note that even if the network is divided into a constant number of regions, the concept of maintaining a track with respect to any terminating point is the same as in *Trail*. Note also that the tracks that are maintained in *Trail* with respect to the terminating point are *tight* with a stretch factor that is less than 1.2 times the shortest distance to the terminating point.

Multiple terminating sets with multiple terminating points: The cases that we have presented can be extended to the case of multiple terminating sets, each with multiple terminating points when there is more knowledge of *find* and *update* patterns within each region.

B. PROOF OF MAXIMA OF EXPRESSION IN EQ. 4

PROPOSITION B.1. *Let $f(\theta, \phi)$ be defined as in Eq. 16.*

$$f(\theta, \phi) = \frac{(\sin(\theta) + \sin(\phi))}{\sin(\theta + \phi)} \quad (16)$$

The maximum value of $f(\theta, \phi)$ where $\theta > 0$, $\phi > 0$, and $0 < (\theta + \phi) \leq \alpha$ is $\sec(\frac{\alpha}{\sqrt{2}})$ and occurs when $\theta = \phi = \frac{\alpha}{2}$.

PROOF. To simplify, we use the following transformation matrix.


$$\begin{aligned} x &= \frac{\theta + \phi}{\sqrt{2}} \\ y &= \frac{-\theta + \phi}{\sqrt{2}} \end{aligned}$$

where, $0 < x \leq \frac{\alpha}{\sqrt{2}}$ and $|y| \leq x$.

Based on the above transformation, $f(\theta, \phi)$ can be written in terms of x, y as follows:

$$g(x, y) = \frac{(\sin(\frac{x+y}{\sqrt{2}}) + \sin(\frac{x-y}{\sqrt{2}}))}{\sin(\sqrt{2} * x)} \quad (17)$$

Now we first find the value of y at which function $g(x, y)$ is maximum given a value

Fig. 21. Finding maxima for $f(\theta, \phi)$

of x . For this we differentiate $g(x, y)$ partially with respect to y and equate the result to 0.

$$\begin{aligned} \frac{\partial g(x, y)}{\partial y} &= 0 \\ \Rightarrow \frac{(\cos(\frac{x+y}{\sqrt{2}}) - \cos(\frac{x-y}{\sqrt{2}}))}{\sqrt{2} * \sin(\sqrt{2} * x)} &= 0 \\ \Rightarrow y &= 0 \end{aligned}$$

Thus we find that for function $g(x, y)$ at any given value of x , $y = 0$ is the only stationary point since $x > 0$. Differentiating $g(x, y)$ partially twice with respect to y , we note that result is less than 0 when $y = 0$.

$$\frac{\partial^2 g(x, y)}{\partial y^2} \Big|_{y=0} = -\frac{1}{\cos(\frac{x}{\sqrt{2}})} < 0$$

Thus for any given value of x , $g(x, y)$ is maximum when $y = 0$.

Also, when $y = 0$, $g(x, y) = \frac{1}{\cos(\frac{x}{\sqrt{2}})}$ which increases monotonically when $x > 0$ and since $x \leq \frac{\alpha}{\sqrt{2}}$, the maximum value is $\sec(\frac{\alpha}{\sqrt{2}})$ and the maximum occurs when $x = \frac{\alpha}{\sqrt{2}}$, i.e. $(\theta + \phi) = \alpha$.

□

C. PROOF OF LEMMA 4.1

In the WSN virtual grid, the total cost of exploring along squares up to level $2^{\lceil \log(d_f) \rceil}$ is given by $8 * \sum_{j=0}^{\lceil \log(d_f) \rceil} 2^j$, i.e. $32 * d_f$. Recall from the proof of Theorem 3.7 that when the trail is intersected by the circle of radius $2^{\lceil \log(d_f) \rceil}$, the point s at which the trail is intersected can be at most $3 * d_f$ away from the object p . The cost of reaching p from the point of intersection with $trail_P$ is bounded by $3 * d_f * TS_p * \sqrt{2}$ where TS_p is the maximum trail stretch factor possible for P . Note that there is an additional stretch of $\sqrt{2}$ because of routing along a grid. The result follows.

D. PROOF OF LEMMA 5.1

Recall from Eq. 1, $(\angle p, N_j, c_j) < (\arcsin(\frac{1}{2^b}))$, where N_j is any level j vertex where $mx \geq j \geq 1$. Since N_j, N_{j-1} and c_{j-1} form a straight line, recall that $\angle N_j p N_{j-1} + \angle p N_j N_{j-1} = \angle p N_j c_j$ for $mx > j \geq 1$. Similarly, since C, N_{mx} and c_{mx} form a straight line, also recall that $\angle C p N_{mx} + \angle p C N_{mx} = \angle p N_{mx} c_{mx}$. Using these we have the following equations.

$$(\angle N_j p N_{j-1}) < (\arcsin(\frac{1}{2^b})) \quad \forall j : (mx > j \geq 1) \quad (18)$$

$$(\angle C, p, N_{mx}) < (\arcsin(\frac{1}{2^b})) \quad (19)$$

Using Eq. 18 and Eq. 19, we obtain $\angle C p N_k$ by summing up as follows.

$$\begin{aligned} (\angle C, p, N_k) &= \angle C p N_{mx} + \sum_{j=k+1}^{mx} (\angle N_j p N_{j-1}) \\ &< (mx - k + 1) * (\arcsin(\frac{1}{2^b})) \end{aligned}$$

E. PROOF OF LEMMA 5.2

Let mx be the index of the highest level in $trail_P$. Using property *P3* we get that $dist(p, c_{mx}) < 2^{mx-b}$.

$$\begin{aligned} dist(C, c_{mx}) &\leq d_{pC} + dist(p, c_{mx}) \\ &< d_{pC} + 2^{mx-b} \end{aligned}$$

By the definition of mx , $2^{mx} < dist(C, c_{mx})$. Therefore we get the following equation.

$$2^{mx} < d_{pC} + 2^{mx-b}$$

Since $b \geq 1$, we get the following equation.

$$d_{pC} > 2^{mx-1}$$

Thus $mx \leq \lceil (\log(d_{pC})) \rceil$. Hence, $m\hat{x}_p = \text{minimum}(\lceil \log(d_{pC}) \rceil, \top)$. \square

F. ANALYSIS OF NECESSARY EXPLORATION FOR OPTIMIZED *FIND* ALGORITHM

Let Q be the finder at distance d_{qC} from C . Thus $m\hat{x}_q = \text{minimum}(\lceil \log(d_{qC}) \rceil, \top)$. In the basic *find* algorithm, the *find* operation will explore at all levels k where $0 \leq k \leq \lceil \log(d_{qC}) \rceil - 1$ and at each level in circles of radius 2^k . In terms of $m\hat{x}_q$, the highest level of exploration for any location of q in the network is $m\hat{x}_q - 1$. This is because when $m\hat{x}_q = \top$, it follows that $m\hat{x}_q = \lceil \log(d_{qC}) \rceil$ and the highest level

of exploration in $m\hat{x}_q - 1$. When $m\hat{x}_q = \lceil \log(d_{qC}) \rceil$, the highest level of exploration is $m\hat{x}_q - 2$. Thus in either case, the level of exploration is bounded by $m\hat{x}_q - 1$.

Let P be an object which should be found at the level k exploration. At level k of the exploration, $trail_P$ for any location of P within the circle of radius 2^k around q should be intersected. A circular exploration of radius 2^k around q is sufficient to achieve this. We now determine the necessary exploration.

Since $dist(p, q) \leq 2^k$, $d_{pC} \leq d_{qC} + 2^k$. Thus at any level k of the exploration, $m\hat{x}_p \leq \lceil \log(d_{qC} + 2^k) \rceil$. Note that $d_{qC} \leq m\hat{x}_q$ and $k \leq (m\hat{x}_q - 1)$. Therefore $m\hat{x}_p \leq m\hat{x}_q + 1$.

We now outline our procedure for level of exploration $k = m\hat{x}_q - 2$.

Level of Exploration $k = m\hat{x}_q - 2$ Refer to Fig. 22. Let α_b denote the value of $\arcsin(\frac{1}{2^b})$ for a given b . Since in our analysis we consider $b = 3$, using Eq. 15 we get that the maximum angle $\angle CpN_k$ is $(m\hat{x}_p - k + 1) * (\alpha_3)$. Since the finder object Q is unaware of $m\hat{x}_p$, the worst case estimate for $m\hat{x}_p$ is used, i.e. $m\hat{x}_p = m\hat{x}_q + 1$. Thus, the maximum angle $\angle CpN_k = 3 * \alpha_3$.

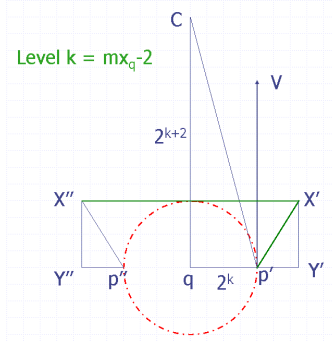


Fig. 22. Level of exploration $k = m\hat{x}_q - 2$

Given this angle, we are interested in determining the smallest segment (X'', X') that will intersect $trail_P$ for any location of P within the dotted circle. This is obtained by drawing a segment from point p' and p'' at angle $3 * \alpha_3$ with segment (C, p') and segment (C, p'') respectively. Point X' is obtained by extending this segment such that $dist(X', Y') = 2^k$. Point X'' is obtained similarly. Now segment (X', X'') will intersect trails of all objects at distance 2^k from q , where $k = m\hat{x}_q - 2$.

From Fig. 22, we note that $\angle qCp' = \arctan(\frac{1}{4})$ and $\angle Cp'X' = 3 * \alpha_3$. Plugging in the values we obtain the minimum required exploration at level $k = m\hat{x}_q - 2$ as follows (approximated to one decimal place):

$$dist(X', X'') = 2.5 * 2^k$$

Similarly, we determine the necessary pattern of exploration at levels $0, \dots, m\hat{x}_q - 1$. Finally we show that, at levels of exploration k where $k \geq m\hat{x}_q - 7$, circular explorations can be avoided and instead a pattern of exploration along the base of an isosceles triangle with apex q and length of base determined by Fig. 10 is sufficient to intersect the trails of all objects at distance 2^k from Q . The base of the

isosceles triangle is such that $segment(C, q)$ is the perpendicular and equal bisector of the base of the triangle. At levels of exploration $k < m\hat{x}_q - 7$, exploration along the entire circle is necessary.

Note: All distances d_{qC} in the range $2^{m\hat{x}_q-1} < d_{qC} \leq 2^{m\hat{x}_q}$, result in the same value of $m\hat{x}_q$. But in the above analysis, we assumed $d_{qC} = 2^{m\hat{x}_q}$. This results in finding the maximum exploration needed because when d_{qC} is smaller, $\angle Cp'X'$ increases, thus decreasing $\angle Vp'X'$ and lowering the length of exploration. \square

G. PROOF OF LEMMA 5.3

Let Q lie between circles of level k and $k - 1$ of the *find-centric* trail for P . The worst case find cost occurs when q is just outside the level $k - 1$ circle. Note that $dist(p, c_{k-1}) < 2^{k-2}$ and therefore $dist(p, q) > 2^{k-2}$

Now, q can travel distance $2 * 2^k$ to reach circle k . Let the point of intersection of the *find* path from q and circle k be t . The cost of following pointers from t to the centers of inner circles recursively and reaching P is given by $(2^0 + 2^1 + \dots + 2^k)$, i.e., $2 * 2^k$.

The ratio of find cost to the distance is thus less than $\frac{4 * 2^k}{2^{k-2}}$, i.e 16. Hence if $dist(p, q) = d$, then the maximum cost of finding object P is $16 * d$. \square

H. TRAIL PROTOCOL ACTIONS IN GUARDED COMMAND

<p>Protocol Trail at mote j</p> <p>Var</p> <ul style="list-style-type: none">$j.child_p$: child pointer for object p$j.prnt_p$: parent pointer$j.detect_p$: boolean$j.level_p$: level of mote j$j.vertex_p$: boolean indicating if j is a vertex$j.cent_p$: location of P when j was last updatednh : temporary variable to store the next hop for any message

Fig. 23. Trail: State at Mote j

Track Update Actions

$\langle U_1 \rangle :: ((j.detect_p) \wedge (j.child_p \neq j)) \longrightarrow$
 $j.child_p = j;$
 $nh = \text{nexthop of exploration};$
 $send_{(j,nh)}(explore(p,j));$

\square

$\langle U_2 \rangle :: recv_{k,j}(explore(p,m)) \longrightarrow$
if $(j == C)$
 $nh = \text{nexthop towards } m;$
 $w = \lceil \log(dist(C,m)) \rceil - 1;$
 $send_{(j,nh)}(grow(p,m,w));$
else
if $(\neg j.child_p)$
 $nh = \text{nexthop of exploration};$
 $send_{(j,nh)}(explore(p,m));$
else
if $((j.vertex_p) \wedge (dist(j.cent_p, m) < 2^{j.level_p-1}))$
 $send_{(j,j.child_p)}(clear(p));$
 $nh = \text{nexthop towards } m;$
 $send_{(j,nh)}(grow(p,m,j.level_p-1));$
 $j.child_p = nh;$
else
 $send_{(j,prnt_p)}(explore(p,m));$
fi
fi
fi

\square

$\langle U_3 \rangle :: recv_{k,j}(grow(p,m,w)) \longrightarrow$
if $(j == m)$
 $j.prnt_p = k;$
else
 $nh = \text{nexthop towards } m;$
if $(dist(j.cent_p, m) > 2^{w-1}) \wedge (dist(j.cent_p, nh) \leq 2^{w-1})$
 $j.vertex_p, j.cent_p, j.level_p, j.child_p = true, m, w, nh;$
 $send_{(j,nexthop)}(grow(p,m,w-1));$
else
 $j.prnt_p, j.level_p = k, w;$
 $send_{(j,nh)}(grow(p,m,w));$
 $j.child_p = nh;$
fi
fi

\square

$\langle U_4 \rangle :: recv_{k,j}(clear(p)) \longrightarrow$
 $j.child_p = \perp;$
if $((j.child_p \neq j))$
 $send_{(j,j.prnt_p)}(clear(p));$
fi

Fig. 24. Trail: Update Actions

Actions for Finding Object	
$\langle F_1 \rangle ::$	$recv_{k,j}(find(p,q)) \longrightarrow$ if ($j.child_p == \perp$) $nh = \text{nexthop of exploration};$ $send_{j,nh}(find(p,q));$ \square $(j.child_p \neq j) \wedge (j.child_p \neq \perp)$ $send_{j,j.c_p}(find(p,q));$ \square $(j.child_p = j);$ $nh = \text{nexthop towards } p;$ $send_{j,nh}(found(p,q));$ fi
\square $\langle F_1 \rangle ::$	$recv_{k,j}(found(p,q)); \longrightarrow$ if ($j \neq q$) $nh = \text{nexthop towards } p;$ $send_{j,nh}(found(p,q));$ fi

Fig. 25. Trail: Find Actions

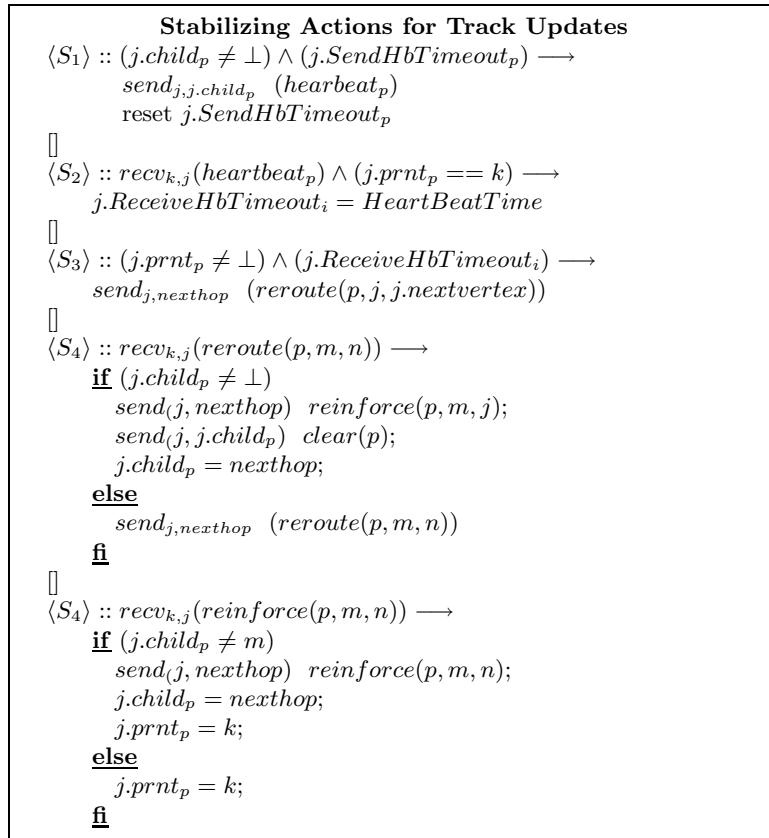


Fig. 26. Trail: Stabilizing Actions