Design and Implementation of a Flexible Location Directory Service for Tiered Sensor Networks

Sangeeta Bhattacharya, Chien-Liang Fok, Chenyang Lu, Gruia-Catalin Roman

Department of Computer Science and Engineering
Washington University in St. Louis

Abstract. Many emergent distributed sensing applications need to keep track of mobile entities across multiple sensor networks connected via an IP network. To simplify the realization of such applications, we present MLDS, a Multi-resolution Location Directory Service for tiered sensor networks. MLDS provides a rich set of spatial query services ranging from simple queries about entity location, to complex nearest neighbor queries. Furthermore, MLDS supports multiple query granularities which allow an application to achieve the desired tradeoff between query accuracy and communication cost. We implemented MLDS on Agimone, a unified middleware for sensor and IP networks. We then deployed and evaluated the service on a tiered testbed consisting of tmote nodes and base stations. Our experimental results show that, when compared to a centralized approach, MLDS achieves significant savings in communication cost while still providing a high degree of accuracy, both within a single sensor network and across multiple sensor networks.

1 Introduction

Many emerging distributed sensing applications require the capability of keeping track of a large number of mobile entities over a wide area that is covered by tiered sensor networks. Let's consider the specific example of co-ordinating doctors over multiple make-shift clinics, set up after a natural calamity. Such clinics are often short of doctors and so the doctors may move between the various clinics, depending on the need of the clinics. In such a scenario, there is often a need to keep track of the doctors, as they move within and between clinics, so that it is possible to find a particular doctor or the nearest available doctor. Existing infrastructure (e.g. phone lines and cell phone towers) is often destroyed or overloaded in such scenarios, requiring the deployment of sensor networks connected via ad hoc IP networks to achieve the objective. As another example, consider the tracking of tools that are shared between various workshops spread across a manufacturing facility. The tools are usually moved around within one or more workshops by the workers. Hence, it is very difficult to locate a particular tool when it is needed. In such a situation it would be helpful to keep track of the location of the tools as they are moved within and across workshops. This would

allow a worker to easily find the nearest available tool that he needs. Sensor networks help realize such applications by providing the capability to sense and identify the mobile entities. However, to fully realize such applications, it is essential to provide a location directory service that can efficiently maintain the location information of mobile entities as they move *across multiple sensor networks* as well as support a broad range of *spatial queries* concerning the mobile entities. Our goal is to realize exactly such a service. The primary contribution of our work is the design, implementation, and empirical evaluation of MLDS, the first Multi-resolution Location Directory Service for *tiered sensor networks*. The key contributions of our work include (1) Design of MLDS, which efficiently maintains location information of mobile entities across multiple sensor and IP networks *and* supports a rich set of multi-granular spatial queries; (2) Implementation of MLDS on tiered sensor networks composed of resource constrained sensor networks and IP networks and (3) Empirical evaluation of MLDS on a tiered testbed of 45 tmote nodes. Our empirical results show that MLDS can maintain a high degree of accuracy at low communication cost, both within a single sensor network and across multiple sensor networks.

2   Services

MLDS can support multiple sensor networks connected by IP networks. Each sensor network, consisting of stationary location-aware sensor nodes and a base station, is assumed to have a unique name that maps to the base station's IP address. We assume that the sensor networks track mobile entities in the physical environment using existing tracking algorithms [1–3] or RFID technology. Furthermore, in our implementation of MLDS, we assume that mobile entities are represented by mobile agents in the sensor network. A mobile agent is a software process that can migrate across nodes while maintaining its state. Mobile agents present a convenient way of representing mobile entities (e.g. cars, people and wild fire) in the sensor network [4]. For instance, in the make-shift clinic example described above, mobile agents may be created to shadow the doctors. Users can then query the locations of doctors by querying the locations of the corresponding mobile agents, through MLDS. Note that even though MLDS is implemented to work with mobile agents, it can be easily extended to work with other programming models for mobile entity tracking such as EnviroSuite [1] and others based on message passing [2, 3].

MLDS supports four types of flexible spatial queries that include (i) finding the location of a particular agent, (ii) finding the location of all agents, (iii) finding the number of agents and (iv) finding the agent that is closest to a particular location. To meet the needs of diverse applications, all of these queries support different scopes and granularities that can be specified by the application. MLDS supports two query scopes, (i) *local scope* i.e. within a single sensor network and (ii) *global scope* i.e. across all sensor networks. It supports three query granularities, *fine*, *coarse* and *network*. The query result of a fine query is based on the exact locations of the mobile agents while the query result of a coarse query is

based on the approximate locations of the mobile agents. The query result of a network query, on the other hand, is based only on the knowledge of the sensor networks that the agents are in. MLDS supports queries issued from both within a sensor network and from outside a sensor network (e.g. by an agent or user on the IP network). The scope and granularity of a query are set via parameters S and G, respectively. Queries can also be limited to a "class" of mobile agents through a parameter C. The API of the four spatial queries are as below:

1. **GetLocation**(id, S, G) returns the location of an agent with ID id.
2. **GetNum**(C, S) returns the number of class C agents.
3. **GetAll**(C, S, G) returns the location of all class C agents.
4. **GetNearest**(C, L, S, G) returns the location of the class C agent that is closest to the location L.

3 Design

MLDS is designed for common sensor network tracking applications like vehicle and personnel tracking for security, emergency care etc. Due to the high mobility of agents in these systems, the location information update rate is expected to be much higher than the query rate in these systems. Hence, MLDS is specifically tailored for systems in which the location information update rate is greater than the query rate. To optimize the operation of such systems, MLDS adopts a hierarchical architecture with multi-resolution information storage. As a result (1) it can support multi-granular spatial queries which enables applications to achieve the desired tradeoff between location information accuracy and communication cost, (2) location information update is not always propagated to the upper tiers of the hierarchy, which significantly reduces communication cost and (3) queries are answered at the closest tier of the hierarchy that meets the query scope and granularity requirements, thus reducing both communication cost and query latency. Note that while MLDS' hierarchical directory structure bears some resemblance to the Domain Name System (DNS) in the Internet and cellular networks, its novelty lies in the fact that it is specifically designed and implemented for tiered sensor networks consisting of resource constrained sensor platforms. In particular, our goal was to minimize communication cost without considerable loss in data accuracy. Moreover, MLDS provides a rich set of multi-granular spatial queries, which is not supported by the above systems.

3.1 Architecture

MLDS has a four tiered hierarchical architecture. The topmost tier of the hierarchy is a central registry that stores information about the different sensor networks. The base stations of the different sensor networks, that are connected by IP networks, form the second tier of the hierarchy. The other two tiers of the hierarchy lie within the sensor networks and are formed by a clustering algorithm that groups the sensor nodes into non-overlapping 1-hop clusters. The

clusterheads of these clusters form the third tier of the hierarchy while the cluster members form the fourth tier. Note that the system consists of heterogeneous nodes, with nodes at higher tiers having more resources than nodes at lower tiers. For example, the clusterheads are resource constrained sensor nodes; the base stations are more powerful computers such as PCs or stargates; while the registry, is stored at a server or server cluster.

MLDS stores location information at different resolutions, at different tiers of the hierarchy. Clusterheads store the exact location of the agents in their cluster while base stations store only the IDs of the clusters that the agents in their network belong to. The registry on the other hand stores the IDs of the networks (denoted by the network base station IP address) that all agents in the system belong to. A base station also maintains the location of the clusterhead and the minimum bounding rectangle (MBR) of each cluster in its network. While the registry also stores the MBR of all the connected sensor networks. The network and cluster MBRs are needed to answer nearest neighbor queries, as explained later in Section 3.3.

3.2   Location Information Maintenance

Since MLDS maintains less accurate information at higher tiers of the hierarchy, location information is not always propagated to the upper tiers, which significantly reduces communication cost. In the following we describe how MLDS maintains agent location information at different tiers of the hierarchy.

A node hosting an agent periodically sends location update messages to its clusterhead, at an interval ΔT. Note, periodic messages are required to maintain the directory in the face of node/agent failures. The location update messages contain the agent ID, class and location, which is set to the location of the host node. When a clusterhead receives a location update message, it first updates it's directory with the agent information. If the agent has just entered its cluster, it then sends a message to the base station containing the agent ID and class, and it's own ID, instead of the agent location. The base station in turn updates its directory on receiving this information and also updates the registry if an entry for the agent did not exist in its directory, previously.

Agent location information at a clusterhead expires after a period $2\Delta T$. Thus, if a clusterhead does not receive location update messages from an agent for a period $2\Delta T$, it assumes that the agent has left its cluster and hence deletes the agent from its directory. A clusterhead may therefore have stale information for a maximum period of $2\Delta T$. This design trades off accuracy for lower communication cost and was preferred over other options that provide higher accuracy but at a higher communication cost.

3.3   Query Processing

MLDS answers a query at the closest tier of the hierarchy that meets the query scope and granularity requirements. For queries issued from within(outside) a sensor network, the closest tier would be the lowest(highest) tier of the hierarchy

that meets the query scope and granularity requirements. This approach reduces both communication cost and query latency. All queries issued by an agent from within a sensor network are first sent to the clusterhead of the cluster that the agent is in. If the query type is GetLocation or GetNearest, the clusterhead checks if it can answer the query. If it can, it sends the query reply to the querying agent, otherwise it forwards the query to the base station. On the other hand, if the query type is GetAll or GetNum the query is directly forwarded to the base station. The base station processes the query and sends the reply to the clusterhead that sent the query, which in turn forwards the reply to the querying agent. Queries issued by an external agent or user on the IP network are sent to the relevant base stations that process the queries and route the result back to the querying agent/user.

We now explain how MLDS processes a query when the query is issued by an agent within a sensor network. Since a base station processes in-network-queries the same way that it processes out-of-network-queries, the later process can be derived from the description of the former, and hence is not explicitly described. Moreover, due to space limitations, we only describe the GetNearest and GetLocation query types in detail. The GetNum query is the simplest of all queries and just involves querying the base station, while the GetAll query is a simple extension of the GetLocation query. In the following discussion, we assume that the ID of the querying agent is q. We also assume that for any agent with ID i, C_i denotes the clusterhead of the cluster that agent i is in and B_i denotes the base station of the network that agent i is in.

GetLocation When clusterhead C_q receives a GetLocation(id, S, G) query from agent q, it checks if agent id is in its cluster. If the agent is in its cluster, it sends a query reply to agent q. If agent id is not in agent q's cluster, then C_q forwards the query to the base station B_q. On receiving this query, B_q checks if agent id is in its network. If the agent is in the network, B_q sends a reply containing either the location of the clusterhead C_{id}, if the query is coarse or the exact location of the agent, which it obtains from C_{id}, if the query is fine. In the case that agent id is not in the local network (i.e. $B_{id} \neq B_q$), B_q finds out B_{id} from the registry, and forwards the query to B_{id}. B_{id} processes the query as explained above and sends the result to B_q. B_q sends the query result to C_q, which forwards it to agent q.

GetNearest When C_q receives a GetNearest(C, L, S, G) query, it checks if there are class C agents in its cluster. If there are such agents, C_q finds the agent that is geographically closest to location L and sends a reply to agent q. If there are no class C agents in the cluster, C_q forwards the query to B_q.

Let's first see how B_q handles local queries. If the query is coarse, B_q just returns the location of the clusterhead, whose cluster contains class C agents and whose location is geographically closest to location L. However, if the query is fine, B_q finds the answer by using the branch and bound technique [5]. The intuition behind this technique is to query only those clusters that contain class

C agents whose locations could be closest to location L. These clusters are found by first obtaining a set of clusters that contain class C agents and then looking at the minimum and maximum distances of the MBRs of the clusters in this set, from L. Clusters whose minimum distances are greater than the maximum distance of the cluster that has the least minimum distance, are discarded. B_q queries the clusterheads of the remaining clusters and waits for a certain time period to hear from them. When B_q hears from all the clusterheads (before the end of the time period) or at the end of the time period, B_q computes the agent that is closest to location L based on the information obtained in the query replies and sends the reply to C_q. Note that although the MBR of a cluster does not accurately represent the cluster boundary, it is preferred over other complex methods like the convex hull due to its low computational complexity.

B_q handles global queries similarly, by first looking up the registry to find the networks that contain class C agents and then applying the above branch and bound technique at the network level. Note that by design, this query returns the approximate geographically closest agent. This design achieves lower communication cost by trading off accuracy.

3.4   Cost Benefits Over A Centralized Directory Approach

In this section we discuss key benefits of MLDS when compared to a centralized directory (CD) approach. In CD, location information is stored in a centralized directory maintained at the base stations of the sensor networks. Hence, all location information and queries in a sensor network are sent to the base station in CD. MLDS has the following key properties when compared to CD:

- MLDS has significantly lower location update cost, when compared to CD, especially when agents move locally (within a cluster) most of the time, which is common in many application scenarios.
- MLDS achieves a lower total communication cost compared to CD when the location update rate is higher than the query rate. Moreover, for a given update and query rate, the savings in communication cost increase with increasing locality of movement and also with increasing network size. Thus, MLDS is more scalable in comparison to CD.
- Coarse and network query cost in MLDS is close to the query cost in CD. The cost of fine queries is low, when answered locally (by the clusterhead), but high otherwise. Thus, applications that can tolerate coarse query results benefit the most from using MLDS.

A more in-depth theoretical comparison of MLDS and CD can be found in [6].

4   Implementation

We have implemented and integrated MLDS with Agimone, a unified middleware that integrates sensor and IP networks. In this section, we first give an overview of Agimone and then describe the implementation details of MLDS.
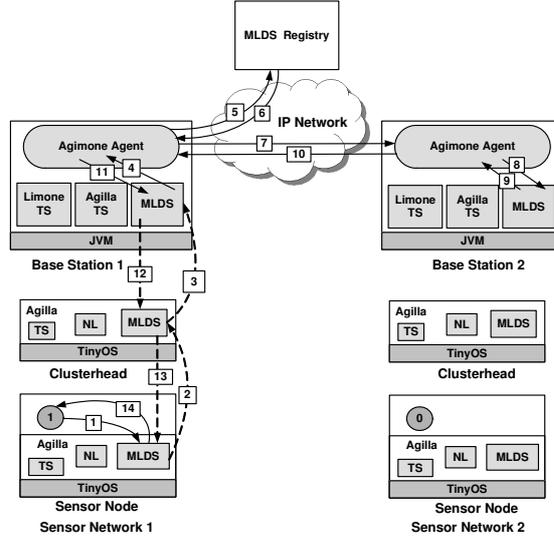
Fig. 1. Interaction between MLDS and Agimone modules when the **GetLocation(0, "global", "coarse")** query is issued by agent 1.(TS: Tuple Space, NL: Neighbor List)

Agimone [7] combines two mobile agent middlewares called Agilla [4] and Limone [8]. Agilla is optimized for resource-constrained sensor networks and is implemented in nesC on the TinyOS platform. Limone is designed for more powerful nodes (e.g. PDAs, stargates and laptops) connected by IP networks and is implemented in Java on standard Java Virtual Machines (JVMs). In Agimone, creation and deployment of mobile agents within a sensor network is done using Agilla, while migration of mobile agents across sensor networks via an IP network, is done using Limone. Agilla provides primitives for an agent to move and clone itself from sensor node to sensor node while carrying its code and state, effectively reprogramming the network. To facilitate inter-agent coordination within a sensor network, Agilla maintains a local tuple space and neighbor list on each sensor node. Multiple agents can communicate and coordinate through local or remote access to tuple spaces. In Agimone, the base stations communicate through Limone tuple spaces maintained at the base stations. Specific Limone agents called AgimoneAgents that reside at the base stations provide an interface between Agilla and Limone and enable the migration of Agilla agents across an IP network. Agimone maintains a central registry for the registration and discovery of sensor networks over the IP network.

MLDS is integrated with the Agimone modules that run on the sensor nodes and base stations. It is implemented in nesC on the sensor nodes and in Java on the base station. MLDS also extends the Limone registry to serve as the registry for its location directories. Figure 1 shows the interaction between the MLDS and Agimone modules at different tiers of the hierarchy when the **GetLocation(0, "global", "coarse")** query is issued by an agent with ID 1. Agent 1 is in sensor

network 1 while agent 0 is in sensor network 2, in the figure. Note that the agents are Agilla agents. Steps 1-3 in the figure show the query message being propagated up the hierarchy to the base station. Once it reaches the MLDS module at the base station, control is transfered to the AgimoneAgent (step 4), since agent 0 is not found in sensor network 1. The AgimoneAgent then queries the registry to find out which network agent 0 is in (steps 5-6). Once it finds that out, it sends the query to the AgimoneAgent at the base station in sensor network 2 (step 7). The AgimoneAgent in base station 2 queries the local MLDS module to obtain the result of the query (steps 8-9) and sends the result back to the AgimoneAgent in base station 1 (step 10). The AgimoneAgent in base station 1 then sends the query reply to the local MLDS module (step 11). After that, the query reply is forwarded down the hierarchy to agent 1 (steps 12-14).

MLDS adapts Agimone's sensor-network-discovery and neighborhood-maintenance mechanisms, to build and maintain its hierarchical structure. The upper two tiers of the hierarchy are formed via the sensor-network-discovery process, in which the base stations register themselves with the registry. The lower two tiers of the hierarchy that lie within individual sensor networks are formed via a simple clustering algorithm. MLDS uses Agimone's neighborhood-maintenance process to achieve clustering at minimum communication cost. Agimone maintains neighborhood information at each node through a periodic beaconing process. Each node periodically broadcasts *beacon messages* containing its ID and hop count to the base station. The hop count information is used for routing messages to the base station. Details of the clustering algorithm are left out due to space limitations but can be found in [6].

5   Experimental Results

We evaluated MLDS through two sets of experiments. The first set of experiments compares MLDS' performance to the centralized approach (CD) within a single sensor network. Recall that in CD, location information in a sensor network is stored only at the base station. All location information and queries are thus sent to the base station in CD. The second set of experiments evaluate MLDS' ability to keep track of mobile agents across sensor networks. In both experiments, tmotes were arranged in a grid, with the gateway node at one corner of the grid. The gateway node is the tmote that acts as a gateway between the sensor network and the PC which serves as the base station. Multi-hop communication between the nodes was achieved by setting a filter at the nodes, that accepted packets only from neighboring nodes on the grid. In order to collect trace data, all nodes in a sensor network were connected to a PC via USB ports.

We use the following four metrics to evaluate query performance, in our experiments. (1) **Success Ratio**: the ratio of the number of queries that returned the accurate result and the total number of queries issued. Network query results are considered accurate if they contain the correct network name; coarse query results are considered accurate if they contain the correct cluster information and fine query results are considered accurate if they contain the correct agent

(a) Success Ratio   (b) Average Error   (c) Average Query Latency
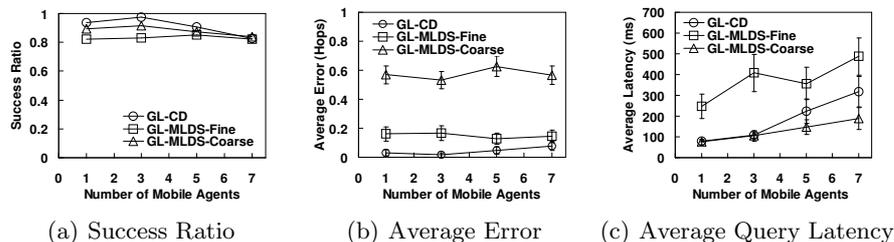
Fig. 2. Performance of local GetLocation queries.

location. (2) **Average Error**: the average error among all queries for which a query result is received, in term of hops. Fine query error is computed as the number of hops between the location returned in the query result and the actual location of the agent. Coarse and network query error is computed as the number of hops the agent is from the clusterhead and from the base station, respectively. (3) **Communication Cost** includes *Location Update Cost* and *Query Cost*. Location Update Cost is the total number of location information messages sent per experiment while Query Cost is the total number of query messages and query result messages sent per experiment. (4) **Average Query Latency**: the average query latency among all queries for which a query result is received. Query latency is the time interval between the issuance of a query and the arrival of the query result, at the querying node. We present 90% confidence intervals for both average error and query latency.

5.1   Single Sensor Network

This set of experiments was carried out on a testbed of 24 tmote nodes, arranged in a 6×4 grid, with a PC as the base station. In each of these experiments, we deployed one stationary agent two hops from the gateway, and n ($1 \leq n \leq 7$) mobile agents. The mobile agents were programmed to follow a random movement pattern over the sensor network at a speed of 1 hop every 5s. Queries were issued by the stationary agent at the rate of 0.2 queries/s. 200 queries were issued in each experiment. Note that by varying the number of mobile agents from 1 to 7 in the experiments, we vary the total location update rate from 0.2 updates/s to 1.4 updates/s and hence evaluate the performance of MLDS under varying network loads. We evaluate only the performance of the **GetLocation** and **GetNearest** queries in these experiments. Since the GetNum query is the same in both MLDS and CD by design and the GetAll query is just an extension of the GetLocation query, we do not evaluate them. We evaluate the performance of the **GetLocation** and **GetNearest** queries, at both fine and coarse granularities, in MLDS. However, only fine queries are evaluated in the centralized approach since it does not support coarse queries. We refer to the GetLocation query in the centralized approach as GL-CD, and the GetLocation fine and coarse queries in MLDS as

(a) Location Update Cost   (b) Query Cost   (c) Communication Cost
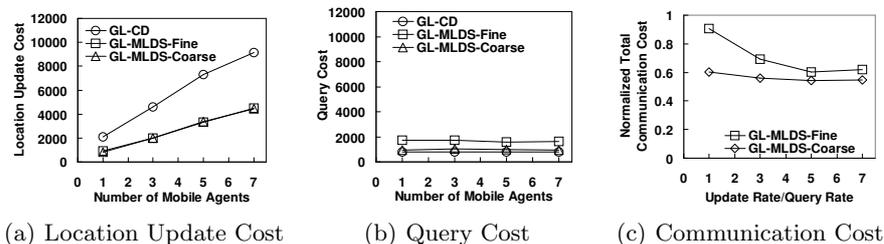
Fig. 3. Communication Cost of local GetLocation queries.

GL-MLDS-Fine and GL-MLDS-Coarse, respectively. Similarly, the GetNearest queries are referred to as GN-CD, GN-MLDS-Fine and GN-MLDS-Coarse.

GetLocation Query Results Figures 2 and 3 show the results obtained for the GetLocation query. From Figure 2(a) we see that the success ratio of GL-CD is higher than that of GL-MLDS-Fine, when there are fewer mobile agents in the network. GL-MLDS-Fine has a lower success ratio partly because a clusterhead retains outdated location information of an agent that has left its cluster, for a maximum time period $2\Delta T$. Interestingly, as the number of mobile agents increases, the success ratio of GL-CD decreases and approaches that of GL-MLDS-Fine. This is because, as the number of mobile agents increases, the number of location information messages also increases. Since all these messages are sent to the base station in CD, there is an increased number of collisions and message loss in the network, which lowers the success ratio of GL-CD. In contrast, the success ratio of GL-MLDS-Fine remains almost constant with the increase in the number of mobile agents, due to its hierarchical architecture. The success ratio of GL-MLDS-Coarse is higher than that of GL-MLDS-Fine and only slightly lower than that of GL-CD. However, since GL-MLDS-Coarse returns an approximate location, it's average error is higher than that of GL-MLDS-Fine, as shown in Figure 2(b). The query reply error of GL-MLDS-Coarse is mostly 1 hop, since MLDS constructs 1-hop clusters. Figure 2(c) displays the query latencies. As expected, GL-MLDS-Fine has the longest query latency since most queries and query results of this type take a longer path. The query latencies of GL-CD and GL-MLDS-Coarse are nearly the same when there are few mobile agents in the system. However, the query latency of GL-CD becomes higher than that of GL-MLDS-Coarse when the number of agents increases, as a result of increased network load.

Figures 3(a), 3(b) and 3(c) show the location update cost, query cost and total communication cost incurred by the GetLocation queries, respectively. From Figure 3(a) we see that MLDS achieves about 55% savings in location update cost when compared to CD. This is due to MLDS' hierarchical architecture, by virtue of which, a large number of location information messages are only sent to the clusterheads and are not forwarded to the base station. Figure 3(b) shows that the query cost of fine queries is higher in MLDS than in CD. This is

(a) Success Ratio   (b) Average Error   (c) Average Query Latency
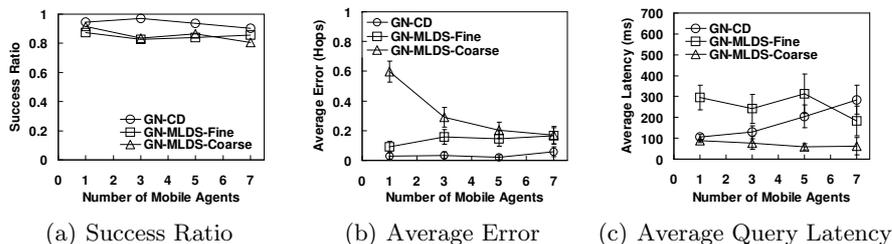
Fig. 4. Performance of local GetNearest queries.

because GL-MLDS-Fine queries that are routed to the base station, get further routed to a clusterhead. Likewise, the query results of these queries take a longer route to reach the querying agent. Comparatively, the query cost of GL-MLDS-Coarse is much lower and is close to that of GL-CD, since these queries are routed only up to the base station. Overall, MLDS achieves significantly lower total communication cost than CD as shown in Figure 3(c). The figure shows the total communication cost of MLDS normalized by the total communication cost of CD for varying ratios of total update rate and query rate. Note that the update rate increases due to the increase in the number of mobile agents in the system. From the figure, we see that the total communication cost of MLDS is lower than that of CD even when the update rate is the same as the query rate and decreases further as the update rate becomes higher than the query rate.

GetNearest Query Results Figure 4 shows the performance of the GetNearest queries. From Figure 4(a) we see that the success ratios of GN-MLDS-Fine and GN-MLDS-Coarse are almost the same, and remain above 80%, irrespective of the number of mobile agents in the network. The average errors of GN-CD and GN-MLDS-Fine reflect the same trend as their success ratios, as shown in Figure 4(b). What is interesting is the trend in the average error of GN-MLDS-Coarse. The average error of GN-MLDS-Coarse is higher than that of GN-MLDS-Fine when there are fewer mobile agents in the system. However, it decreases as the number of mobile agents in the system increases. This trend is due to the fact that as the number of mobile agents increases, the probability of a mobile agent being in the same cluster as the querying agent also increases and so more number of queries are answered directly by the clusterhead of the cluster that the querying agent is in. Thus, with the increase in the agent density, a higher percentage of the coarse query replies contain exact agent locations, which in-turn reduces the error.

The query latency of GN-MLDS-Fine is higher than that of GN-CD when there are few mobile agents in the network, as shown in Figure 4(c). However, as the number of mobile agents increases, the query latency of GN-CD increases considerably whereas the query latency of GN-MLDS-Fine decreases. The query latency of GN-MLDS-Fine becomes less that of GN-CD when there are 7 mobile agents in the network. The increase in the query latency of GN-CD with the

increase in the number of mobile agents is a result of increased network load. The reason for the decrease in query latency of GN-MLDS-Fine, with the increase in mobile agents in the network, is the increase in the percentage of queries that get answered locally, by the clusterhead. This same reason also causes the decrease in the query latency of GN-MLDS-Coarse as the number of mobile agents in the network increases. Thus, the GN query benefits significantly from local responses, made possible by MLDS' hierarchical architecture. The benefit is not only decreasing query latency but also decreasing query cost (not shown here), with increasing agent density.

In summary, MLDS consistently achieves success ratios above 80% in all our experiments, at significantly lower total communication cost than the centralized approach. In particular, coarse queries supported by MLDS achieved the lowest communication cost and query latency, while introducing an average error of less than 1 hop. Thus, applications that can tolerate a small amount of location error gain the most from using MLDS. Furthermore, MLDS' hierarchical architecture enables efficient execution (low cost and latency) of GetNearest queries, especially when the density of mobile agents is high.

5.2   Multiple Sensor Networks

We now evaluate MLDS' performance across multiple sensor networks. In these experiments, mobile agents move between three sensor networks via an IP network running over 100Mbps Ethernet. The IP network is private with a single Linksys WRT54G router and an 8 port switch. These experiments were carried out on a testbed of 45 tmote nodes, equally divided into three sensor networks arranged in a 5×3 grid. Each sensor network has a PC connected via USB to one of its corner motes that serves as a base station. These base station PCs are connected to each other via the IP network. A fourth PC on the IP network serves as the registry. Evaluating MLDS' performance requires comparing its results with the ground truth. The ground truth is obtained by connecting every mote except those directly attached to a base station to the registry PC via USB. The motes are programmed to send trace messages identifying key events like agent movement and query activities over their USB port. The registry PC monitors these connections for incoming trace data and saves them into a file. In addition, it also accepts trace messages over the IP network, which the base stations use to record trace messages generated by the motes they are attached to. The registry PC serves as a central aggregation point for the trace data. Each trace event is time stamped and saved for off-line analysis.

Like the single sensor network experiments, we evaluate only the performance of the GetLocation (GL) and the GetNearest (GN) queries. Both the network and coarse granularity versions of the queries are evaluated. In each of these experiments, the workload is varied by varying the number of mobile agents in the system from 1 to 21 in increments of 3. The mobile agents move 10 hops randomly in a sensor network before randomly migrating to another sensor network and repeating. Initially, the mobile agents are distributed evenly across

(a) Success Ratio   (b) Average Error   (c) Average Query Latency
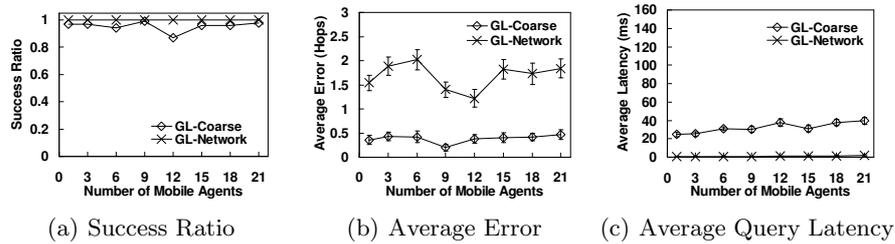
Fig. 5. Performance of global GetLocation queries.

the three sensor networks. The GetLocation and GetNearest queries are issued at a rate of 1 query every 5s by an external agent running on the registry PC. Note that this differs from the single network experiments where the querier was located within the sensor network. By placing the querier on the registry, the query messages only travel down the hierarchy. Scenarios where the query messages travel up the hierarchy were already evaluated in the single-network experiments. Each experiment is repeated 100 times.

GetLocation Query Results For these experiments, a querier located on the registry periodically issues a GL query for a particular mobile agent (termed the *target agent*) within the sensor networks. The success ratio of the GL query is shown in Figure 5(a). Both the coarse (GL-Coarse) and network (GL-Network) granularity versions of GL achieve nearly perfect success. GL-Coarse has a slightly lower success ratio because it attempts to return a more accurate location of the agent. However, it has approximately 3 times lower error, as shown in Figure 5(b). Notice that GL-Network has a higher average error variance and that GL-Coarse has an average error variance of less than one. This is because an agent may be multiple hops away from the network base station, but can be at most one hop away from its cluster head. GL-Coarse queries have significantly longer latency as shown in Figure 5(c). The latency of GL-Network queries is negligible since the querier is located on the registry and can query the registry locally to determine which network the target agent is in. For GL-Coarse queries, the agent must first lookup which network the agent is in, then query that network's base station to determine which cluster the agent is in. As the number of mobile agents increases, the latency also increases due to increased network congestion.

GetNearest Query Results The GN experiments are the same as the GL experiments except the target agent is a stationary agent that resides two hops away from the base station on one of the sensor networks. The querier on the registry periodically searches for the mobile agent closest to the target agent. The results are shown in Figure 6. As the number of mobile agents increase, the success ratio of GN-Coarse decreases due to network congestion preventing

(a) Success Ratio   (b) Average Error   (c) Average Query Latency
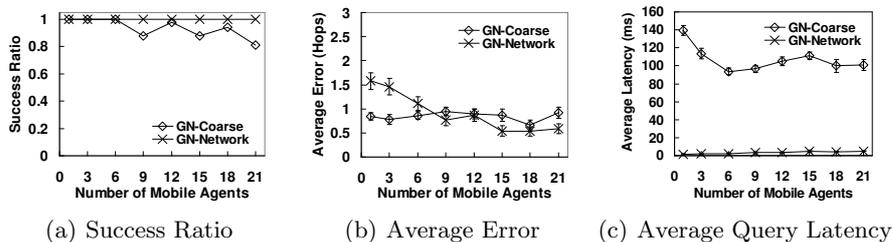
Fig. 6. Performance of global GetNearest queries.

updates from propagating up the hierarchy, as shown in Figure 6(a). On the other hand, GN-Network queries almost always succeed since it involves at most 1 call to the registry. The average error of GN-Coarse queries remains roughly less than 1 hop regardless of the number of agents as shown in Figure 6(b). This is expected since the cluster members are at most 1 hop away from the cluster head. The average error of GN-Network queries is also close to 1, but is dependent on the size of the network. As the number of mobile agents increase, the probability of finding an agent in the same network as the target increases, decreasing the latency of GN-Coarse queries, as shown in Figure 6(c). The latency of GN-Network queries is negligible because in our experiments, an agent is always present in the target agent's network and hence the queries are always answered locally by the base station.

6   Related Work

MLDS is related to data-centric storage (DCS) systems like GHT [9], DIFS [10], DIMENSIONS [11] and DIM [12]. GHT hashes data by name to nodes in the network and provides no index for accessing the data. Hence it is unsuitable for storing and accessing location information. DIFS leverages on GHT and maintains a hierarchical index of histograms to support multi-range queries. DIM, on the other hand, uses a locality-preserving hash function that maps a multi-attribute event to a geographic zone. It divides the network into zones and maintains a zone tree to resolve multi-dimensional range queries. However, the index of neither DIF nor DIM can efficiently support spatial queries. DIMENSIONS hashes sensor data to nodes in the network and maintains a multi-resolution hierarchical index that enables it to efficiently answer queries by drilling down to the appropriate nodes. However, DIMENSIONS was not designed for storing location information and hence does not support spatial queries such as **GetNearest**. Thus, the key differences between the above systems and MLDS are (1) the above systems store sensor data while MLDS is specifically tailored for storing location information of mobile entities, (2) MLDS supports a broad range of flexible spatial queries which cannot be supported efficiently by the above systems, and (3) MLDS builds a distributed directory over multiple sensor networks connected by an IP network, while the above systems systems are designed for

a single sensor network. TSAR [13] is another in-network storage architecture, which stores sensor data at a lower tier consisting of sensor nodes and stores only meta data at a higher tier consisting of a network of proxies. Unlike the above approaches, TSAR maintains a distributed index at the proxies. TSAR differs from MLDS in that it is not tailored for storing location information nor does it support spatial queries. Moreover, unlike MLDS, TSAR does not have an in-network tier that enables the system to take advantage of data locality while resolving queries. The comb-needle approach proposed in [14] also deals with in-network storage and retrieval of data and uses an adaptive push-pull technique to achieve this. Unlike MLDS, this approach does not maintain data locality and hence cannot efficiently support spatial queries. An analysis of this approach to other DCS approaches has been provided in [15].

Our work is also related to the protocols presented in [16] and [17], which address in-network processing of K-Nearest Neighbor (KNN) queries and are based on the branch-and-bound technique [5] which is also used in MLDS. Another related service is EASE [18], which keeps track of mobile entities within a single sensor network through in-network storage and supports multi-precision queries that fetch the location of a specified mobile entity. MLDS differs from the above protocols in the following three important ways: (1) MLDS presents an architecture for storing location information in sensor networks that enables efficient computation of nearest-neighbor as well as other multi-resolution spatial queries, (2) MLDS is the first location directory service that can keep track of mobile entities across multiple sensor networks, and (3) we implemented and integrated MLDS with a mobile agent middleware and present experimental results on a physical testbed. In contrast, the above protocols are only evaluated through simulations.

7   Conclusion

We have developed MLDS, a Multi-resolution Location Directory Service for tiered sensor networks comprising multiple sensor networks connected via IP networks. MLDS has several salient features: (1) it is the first system that maintains location information of mobile entities across sensor and IP networks, (2) it supports a range of multi-granular spatial queries that can span multiple sensor networks and (3) it has low communication cost. We integrated MLDS with Agimone, a mobile agent middleware for sensor and IP networks, and evaluated its performance on a testbed of tmote nodes. The empirical results obtained show that MLDS successfully keeps track of mobile agents across single and multiple sensor networks at significantly lower communication cost than a centralized approach. Most importantly, MLDS enables applications to achieve the desired tradeoff between accuracy and communication cost, which is particularly useful for resource constrained sensor networks. Currently, MLDS is optimized for systems that have a higher location update rate. As future work, we plan to extend MLDS such that it dynamically adapts to the query and the update rate and hence performs well under all conditions. We wish to achieve this by using

a push-pull strategy that dynamically adjusts the location and granularity of location information based on the query and update load.

References

1. Luo, L., Abdelzaher, T., He, T., Stankovic, J.A.: Envirosuite: An environmentally immersive programming framework for sensor networks. TECS (2006)
2. Liu, J., Reich, J., Zhao, F.: Collaborative in-network processing for target tracking. Journal of Applied Signal Processing (2003)
3. Pattem, S., Poduri, S., Krishnamachari, B.: Energy-quality tradeoffs for target tracking in wireless sensor networks. (In: IPSN'03)
4. Fok, C.L., Roman, G.C., Lu, C.: Rapid development and flexible deployment of adaptive wireless sensor network applications. (In: ICDCS'05)
5. Roussopoulos, N., Kelly, S., Vincent, F.: Nearest neighbor queries. (In: SIGMOD'95)
6. Bhattacharya, S., Fok, C.L., Lu, C., Roman, G.C.: Mlds: A flexible location directory service for tiered sensor networks. Technical Report WUCSE-2007-1, Dept. of Computer Science and Engineering, Washington University in St. Louis (2007)
7. Hackmann, G., Fok, C.L., Roman, G.C., Lu, C.: Agimone: Middleware support for seamless integration of sensor and ip networks. (In: DCOSS'06)
8. Fok, C.L., Roman, G.C., Hackmann, G.: A lightweight coordination middleware for mobile computing. (In: Coordination'04)
9. Ratnasamy, S., Karp, B., Yin, L., Yu, F.: GHT: A geographic hash table for data-centric storage. (In: WSNA'02)
10. Greenstein, B., Estrin, D., Govindan, R., Ratnasamy, S., Shenker, S.: DIFS: A distributed index for features in sensor networks. (In: SNPA'03)
11. Ganesan, D., Estrin, D., Heidemann, J.: DIMENSIONS: Why do we need a new data handling architecture for sensor networks? (In: HotNets-I'02)
12. Li, X., Kim, Y.J., Govindan, R., Hong, W.: Multi-dimensional range queries in sensor networks. (In: SenSys'03)
13. Desnoyers, P., Ganesan, D., Shenoy, P.: Tsar: A two tier storage architecture using interval skip graphs. (In: SenSys'05)
14. Liu, X., Huang, Q., Zhang, Y.: Combs, needles, haystacks: Balancing push and pull for discovery in large-scale sensor networks. (In: SenSys'04)
15. Kapadia, S., Krishnamachari, B.: Comparative analysis of push-pull query strategies for wireless sensor networks. (In: DCOSS'06)
16. Demirbas, M., Ferhatosmanoglu, H.: Peer-to-peer spatial queries in sensor networks. (In: P2P'03)
17. Winter, J., Xu, Y., Lee, W.C.: Energy efficient processing of k nearest neighbor queries in location-aware sensor networks. (In: Mobiquitous'05)
18. Xu, J., Tang, X., Lee, W.C.: EASE: An energy-efficient in-network storage scheme for object tracking in sensor networks. (In: SECON'05)