

Bidding Protocols for Deploying Mobile Sensors

Guiling (Grace) Wang, Guohong Cao, Piotr Berman, and Thomas F. La Porta, *Fellow, IEEE*

Abstract—Constructing a sensor network with a mix of mobile and static sensors can achieve a balance between sensor coverage and sensor cost. In this paper, we design two bidding protocols to guide the movement of mobile sensors in such sensor networks to increase the coverage to a desirable level. In the protocols, static sensors detect coverage holes locally by using Voronoi diagrams and bid mobile sensors to move. Mobile sensors accept the highest bids and heal the largest holes. Simulation results show that our protocols achieve suitable trade-off between coverage and sensor cost.

Index Terms—Mobile sensor networks, sensor deployment, distributed algorithm, bidding protocol.

1 INTRODUCTION

RECENT advances in micro-electro-mechanics, micro-processors, and wireless communication have enabled the design of small-size, low-cost sensor nodes. After being deployed into the target field, these nodes can self-organize into a multihop wireless sensor network [13], [5], [4], [31]. Recently, wireless sensor networks have been adopted to a vast number of military and civil applications.

For many applications, a desired distribution of sensors in the target field is difficult to achieve because manual deployment is nearly impossible and the deployment may be affected by uncontrollable factors such as wind and obstacles. In most early research, sensor nodes are assumed to be static and a large number of redundant nodes are deployed to achieve a desired level of coverage. This may introduce high cost and still cannot guarantee coverage.

Recently, researchers have started to consider sensors that are capable of a controlled mobility [33]. With added mobility, sensors can move to provide the required coverage. Various algorithms and protocols [14], [34], [21], [22] have been proposed to assist mobile sensors moving from densely covered areas to sparsely covered areas to achieve balanced coverage. However, to equip every sensor with a motion base increases the network cost and is unnecessary when the coverage requirement is not very strict, or if sensors can be scattered in the target field relatively uniformly. We propose to deploy a mixture of mobile and static sensors to construct sensor networks such that a balance between sensor cost and coverage can be achieved.

In this paper, we design two distributed bidding protocols for the placement of mobile sensors in a sensor network composed of both mobile and static sensors: a basic bidding protocol and a proxy-based bidding protocol, which is an improvement on the basic bidding protocol. In the protocols, mobile sensors are treated as *servers* to heal *coverage holes*.

Coverage holes are locations not covered by any sensor. Each mobile sensor has a *base price*, which is related to the size of any new hole generated by its movement. This represents the cost of its movement in terms of coverage. Static sensors detect coverage holes locally and estimate their sizes as *bids*. The static sensors *bid* the mobile sensors that have a base price lower than the hole to be covered. In the *basic bidding protocol*, mobile sensors choose the highest bids and thus move to heal the largest coverage holes. Using this process, sensors will only move to cover holes larger than those generated by their movements. After moving to the holes, mobile sensors raise their base prices to reflect the new coverage cost and re-enter the bidding process. This process iterates until no static sensor can give a bid higher than the base price of any mobile sensor. Simulation results show that the basic protocol can significantly increase the coverage.

To reduce the moving distances of mobile sensors, the *proxy-based* bidding protocol proposes that mobile sensors perform virtual movements from small holes to large holes and only perform physical movements after the final destinations are identified. Compared with the basic protocol, the proxy-based bidding protocol can save about 50 percent of the moving distance, without sacrificing coverage.

The rest of the paper is organized as follows: Section 2 introduces some preliminaries on our assumptions, technical background, and theoretical analysis. We present the basic bidding protocol in Section 3 and present the proxy-based bidding protocol in Section 4. Section 5 evaluates the performance of the proposed protocols. We conclude the paper in Section 7.

2 PRELIMINARIES

2.1 Assumptions

We assume an isotropic sensing model in which the sensing area of each node is represented by a circle with the same radius [26], [27], [14]. All sensor nodes know their locations. There are mature techniques for a wireless sensor network to determine the location of each sensor in both indoor and outdoor applications [1], [28], [32], [9], [23], and we assume at least one of these is available in our network.

Because we are dealing with mobile sensors, path planning is an important consideration. We assume sensors can plan paths from their current position to a desired destination using one of several existing techniques [8], [24],

• G. Wang is with the Department of Computer Science, New Jersey Institute of Technology, GITC Building, Room 4400, University Heights, Newark, NJ 07102. E-mail: gwang@njit.edu.

• G. Cao, P. Berman, and T.F. La Porta are with the Department of Computer Science and Engineering, The Pennsylvania State University, 111 IST Building, University Park, PA 16802. E-mail: {gcao, berman, tlp}@cse.psu.edu.

Manuscript received 23 Jan. 2006; revised 24 May 2006; accepted 27 June 2006; published online 7 Feb. 2007.

For information on obtaining reprints of this article, please send e-mail to: tmc@computer.org, and reference IEEECS Log Number TMC-0145-0506. Digital Object Identifier no. 10.1109/TMC.2007.1022.

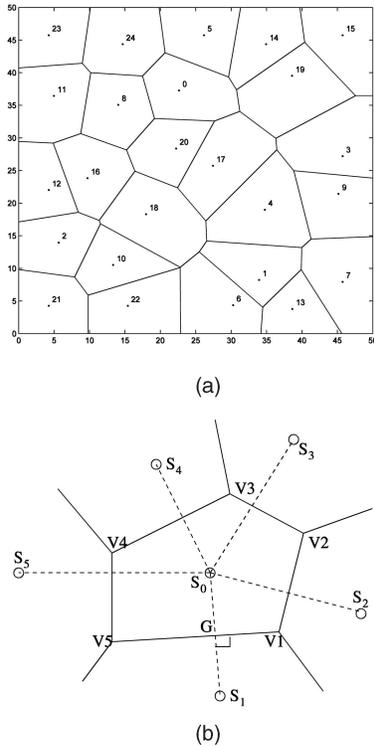


Fig. 1. Voronoi diagram. (a) Voronoi diagram. (b) Voronoi cell G_0 of s_0 .

[17], [25]. We comment more on the impact of this assumption in Section 5.

2.2 Technical Preliminary: Voronoi Diagram

A Voronoi diagram [6], [12], [24] represents the proximity information about a set of geometric nodes. One example is shown in Fig. 1a. The Voronoi diagram of a set of sensors partitions the space into cells. Every point in a given cell is closer to the sensor in this cell than to any other sensor. Thus, if this sensor cannot detect the expected phenomenon in its Voronoi cell, no other sensor can detect it. Therefore, to examine coverage holes, each sensor only needs to check its own Voronoi cell. The calculation of Voronoi cells in a distributed fashion is presented in Section 3.2.

In the paper, we use the following notations: We define the Voronoi cell of sensor s_0 as $G_0 = \langle \mathcal{V}_0, \mathcal{E}_0 \rangle$, where \mathcal{V}_0 is the set of Voronoi vertices of s_0 and \mathcal{E}_0 is the set of Voronoi edges. As shown in Fig. 1b, $\mathcal{V}_0 = \{V_1, V_2, V_3, V_4, V_5\}$ and $\mathcal{E}_0 = \{V_1V_2, V_2V_3, V_3V_4, V_4V_5, V_5V_1\}$. We use \mathcal{N}_0 to denote the set of Voronoi neighbors of s_0 . In Fig. 1b, $\mathcal{N}_0 = \{s_1, s_2, s_3, s_4, s_5\}$. The Voronoi edges of s_0 are the vertical bisectors of the line passing s_0 and its Voronoi neighbors, e.g., V_1V_5 is s_0s_1 's bisector.

2.3 Theoretical Analysis

When a portion of deployed sensors are mobile, the deployment problem can be described as follows: Given a target field covered by a number of circles (the sensing circles of the static sensors), but still having some uncovered areas, how can we place a certain number of additional circles (the sensing circle of the mobile sensors) to maximize the overall coverage?

This problem is an NP-hard problem, which can reduce to the vertex covering problem [29]. The detailed proof of NP-completeness is shown in Appendix A.

Although our problem is a fundamentally difficult problem and there is no optimal solution, we can still find some practical solutions to approximate the optimal solution based on heuristics. Similarly to the greedy algorithm, which is a commonly used heuristic for the vertex covering problem, we can place mobile sensors at the largest coverage holes.

3 BASIC BIDDING PROTOCOL

In this section, we present the basic bidding protocol. We evaluate its performance in terms of coverage, energy consumption, and deployment time in Section 5. The description of this protocol provides a basic understanding and insight into our solution. We improve upon this protocol with optimizations described in Section 4 to reduce the required moving distance for each mobile sensor.

3.1 Bidding Protocol Overview

According to the greedy heuristic for this NP-hard problem, mobile sensors should move to the area where the most additional coverage can be obtained. After a mobile sensor leaves its original location to cover (heal) another coverage hole, it may generate a new hole in its original location. Thus, a mobile sensor only moves to heal another hole if its leaving will not generate a larger hole than that to be healed. However, due to lack of global information, mobile sensors may not know where a coverage hole exists. Even with the location of the coverage hole, it is still a big challenge to find the target position inside the coverage hole which can bring the most additional coverage when a mobile sensor is placed there compared to other positions. We propose letting the static sensors detect the coverage holes locally, estimate the size of these holes, and determine the target position inside the hole. Based on the properties of the Voronoi diagram, static sensors can find the coverage holes locally and provide a good way to estimate the target location of the mobile sensors.

The roles of mobile and static sensors motivate us to design a bidding protocol to assist the movement of the mobile sensors. We view a mobile sensor as a hole healing server. Its service has a certain base price, which is the estimate of any generated coverage hole after it leaves the current place. Static sensors are the bidders of the coverage hole healing services. Their bids are the estimated sizes of the holes they detect. Static sensors bid mobile sensors that have a base price lower than their bid. Mobile sensors choose the highest bid and move to the target locations provided by the static sensors.

The bidding protocol runs round by round after the initialization period. During the initialization period, all static sensors broadcast their locations and identities locally. We choose the broadcast radius to be two hops with which sensors can construct the Voronoi diagram in most cases. After the initialization period, static sensors broadcast this information again only when new mobile sensors arrive and need this information to construct their own Voronoi cells.

Each round consists of three phases: *service advertisement*, *bidding*, and *servicing*. In the advertisement phase, mobile sensors broadcast their base prices and locations in a local area. The base price is set to be zero initially. By the end of the service advertisement phase, each static sensor has a

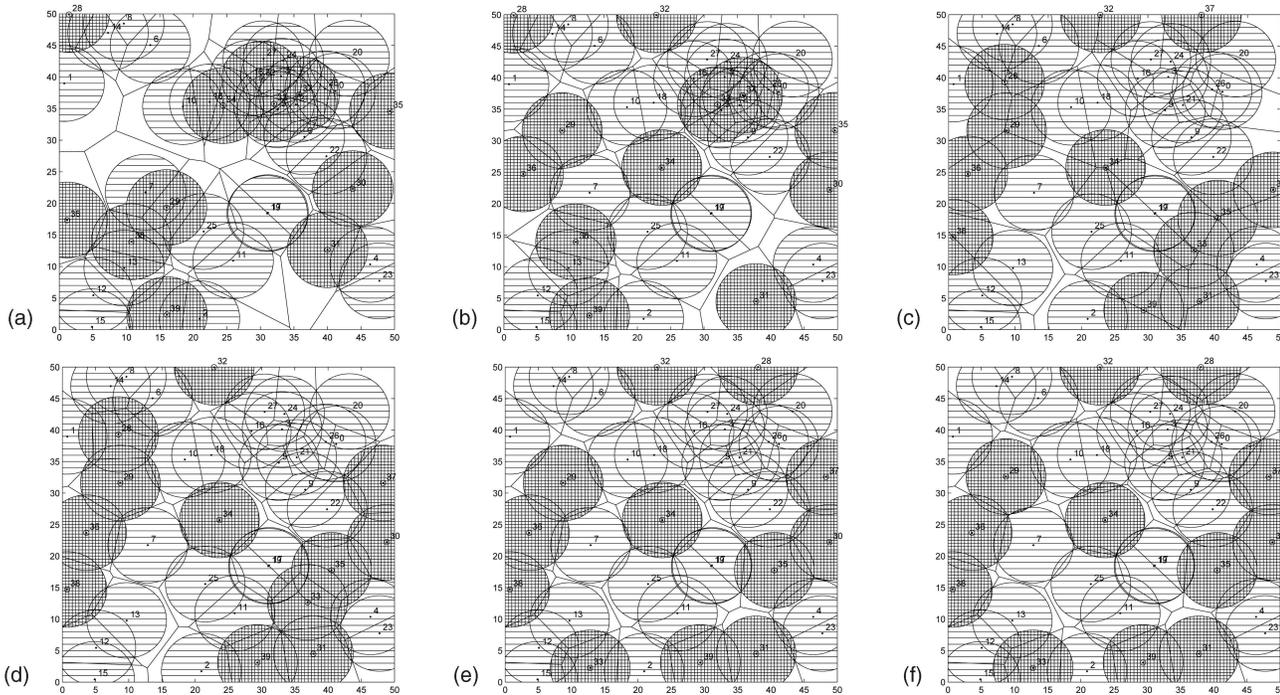


Fig. 2. Snapshot of the execution of the bidding protocol. (a) Original. (b) Round 1. (c) Round 2. (d) Round 3. (e) Round 4. (f) Round 5.

service list, which is a list of mobile sensor IDs along with their location and base price. In the bidding phase, static sensors detect coverage holes locally by examining their Voronoi cells. If such holes exist, they calculate the bids and the target locations for the mobile sensors. Examining the service list, the static sensor chooses a mobile sensor whose base price is lower than its bid and sends a bidding message to this mobile sensor. We will present how to determine the mobile sensor to bid if there are multiple mobile sensors whose base price is lower than the bid of the static sensor. In the serving phase, the mobile sensor chooses the highest bid and moves to heal that coverage hole. The accepted bid will become the new base price of the mobile sensor. After the serving phase, the mobile sensors broadcast their new locations and new base prices and a new round begins. Because the base price increases monotonically, when no static sensors can give out a bid higher than the base price of the mobile sensors, the protocol terminates.

Before getting into the technical details of the bidding protocol, we first use an example to show how the protocol works. As shown in Fig. 2, the circles with a striped shadow represent the sensing coverage of the static sensors, and the circles with grid shadow are that of the mobile sensors. Initially, 40 sensors are randomly placed in a 50 m × 50 m flat field, among which 30 percent are mobile sensors. The initial coverage is 82 percent. The protocol terminates in the fifth round when the coverage reaches 93 percent. The sixth round has the same topology as the fifth round.

3.2 Distributed Calculation of the Voronoi Cell

It is difficult to compute Voronoi diagrams [6]. However, to detect and calculate the coverage hole, each sensor only needs to know its own Voronoi cell, whose calculation can be simplified as follows: We take sensor s_0 as an example. Initially, as shown in Fig. 3a, the Voronoi cell of s_0 is set to be a large rectangle. After receiving the *hello* message from

sensor s_1 , s_0 knows the location of s_1 and computes the bisector line of s_1 and itself. This line is added to the original graph and two cells are generated. Shown in Fig. 3b, the cell including s_0 becomes the new view of s_0 's Voronoi cell.

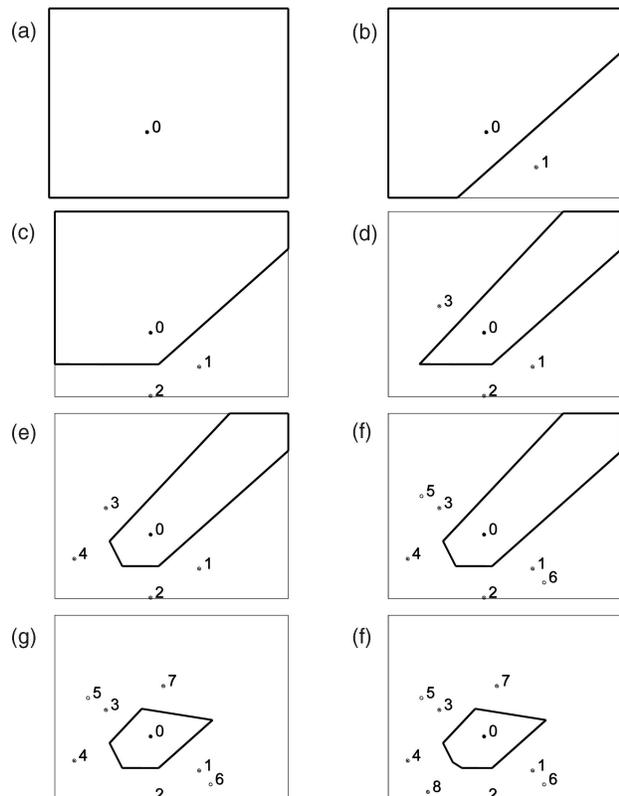


Fig. 3. Computing the Voronoi cell. (a) Original. (b) s_1 is discovered. (c) s_2 is discovered. (d) s_3 is discovered. (e) s_4 is discovered. (f) s_5 and s_6 are discovered. (g) s_7 is discovered. (h) s_8 is discovered.

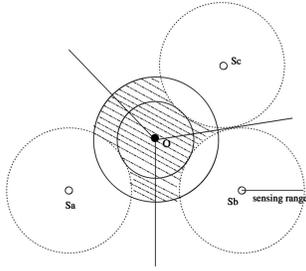


Fig. 4. Bid estimation.

Later, after s_0 receives the *hello* messages from s_2 , s_3 , and s_4 , its Voronoi cell changes from Fig. 3c to Fig. 3e accordingly. The Voronoi cell will not change if the computed bisector line has no intersection with it. As shown in Fig. 3f, knowing s_5 and s_6 does not affect s_0 's Voronoi cell. Finally, the true Voronoi cell is generated after s_0 knows the existence of s_7 and s_8 .

Static sensors construct Voronoi cells considering only static neighbors and mobile neighbors which are not likely to move. These mobile sensors are detected by examining their base prices. If the base price of a mobile sensor is zero, this mobile sensor has not moved yet and most likely it will move to heal some coverage hole. Thus, when detecting coverage holes, static sensors do not consider mobile sensors which are likely to leave. To find out if a coverage hole exists, a static sensor checks whether its distance to the farthest Voronoi vertex is longer than the sensing range. If yes, then some coverage hole exists and this sensor should prepare to bid some mobile sensor to heal it.

Voronoi cells calculated in this way will not be accurate when Voronoi neighbors are far away from each other and cannot communicate with each other. The accurate calculation of Voronoi cells is not required in these cases because the coverage holes will be large. The algorithm will not mis-detect coverage holes.

3.3 Bid Estimation

In the bidding message, static sensors provide the estimated coverage hole size as the bid and the target location to which the mobile sensor should move. This information is calculated based on their Voronoi cells. If there exists a coverage hole, the static sensor chooses the farthest Voronoi vertex as the target location of the coming mobile sensor. Inside one coverage hole, there are many positions at which a mobile sensor can be located. If the mobile sensor is placed at the position farthest from any nearby sensors, the gained coverage is the highest since the overlap of the sensing circles between this new coming mobile sensor and existing sensors is the lowest. As shown in Fig. 4, sensor s_a chooses its farthest Voronoi vertex O as the target location of the mobile sensor for which it bids.

From the global point of view, using the greedy heuristic to choose the largest coverage hole may not be optimal in some cases. As shown in Fig. 5, A is the farthest Voronoi vertex of s_a . Although a high additional coverage can be obtained by placing a mobile sensor at A , it is not globally optimal since it leaves some scattered coverage holes which are hard to cover by placing additional mobile sensors. To deal with this problem, we propose an optimization which puts a limit on the maximum distance between the calculated target location and the bidder. As shown in Fig. 5,

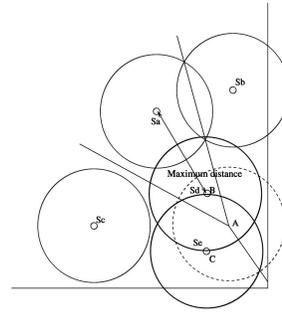


Fig. 5. Optimizing the greedy heuristic.

by setting this maximum distance, a mobile sensor will be placed at B so that another mobile sensor can move to point C to achieve better coverage. This maximum distance, denoted by d_{limit} , is a function of sensing range. We choose d_{limit} to be $\sqrt{3} * sensing_range$. In a large uncovered space, to place sensing circles in a hexagonal relative position will minimize overlapping and maximize the coverage. Under these conditions, the distance between the centers of the sensing circles is $\sqrt{3} * sensing_range$.

Having determined the target location of the mobile sensor it bids, static sensors calculate the bid as

$$\pi * (d - sensing_range)^2,$$

where d is the distance between the bidder and the target location. As shown in Fig. 4, s_a 's bid is the area of the inner circle centered at O , which is not the actual additional coverage to be obtained. The actual additional coverage is the shadow area, which is difficult to calculate since it involves the union of circles. Using the inner circle as the bid simplifies the calculation and can be used to approximate the actual additional coverage, which is the sensing circle minus the overlapping area of the sensing circles. The larger the overlapping area, the smaller the inner circle. Thus, the bid used can represent the relative size of the coverage holes.

Note that the maximum base price (or bid) is

$$\pi * (d_{limit} - sensing_range)^2,$$

which is $\sqrt{3} * \pi * sensing_range^2 / 2$.

The property of the Voronoi diagram guarantees that the shadow area is always the additional coverage. This can be explained as follows: The points inside one Voronoi cell are closest to the sensor in this cell. The points in the Voronoi edge are closest to these two sensors besides this edge. The Voronoi vertex is the point closest to the sensors which contribute to the existence of this vertex. The sensing circle centered at the Voronoi vertex must only overlap with the sensing circles of the sensors which contribute to the construction of this Voronoi vertex. Thus, we guarantee that the shadow area shown in Fig. 4 is always the additional coverage brought by placing a mobile sensor at O .

In addition to the static sensors, mobile sensors with a base price larger than zero also act as bidders. This is necessary because mobile sensors with a relatively larger price are essentially acting as static sensors. At this point, they can assist the movement of other mobile sensors.

3.4 Criteria of Choosing Mobile Sensors to Bid

After the service advertisement phase, each sensor has a list of mobile sensors, their locations, and their base prices. A

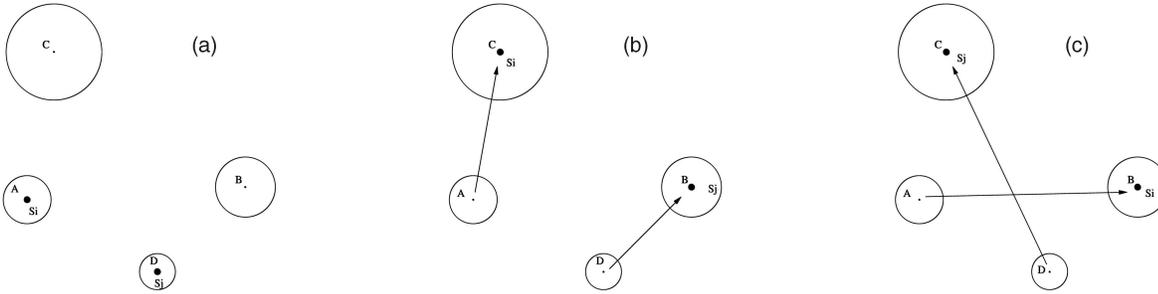


Fig. 6. Distance-based versus price-based. (a) Original situation. (b) Distance-based. (c) Price-based.

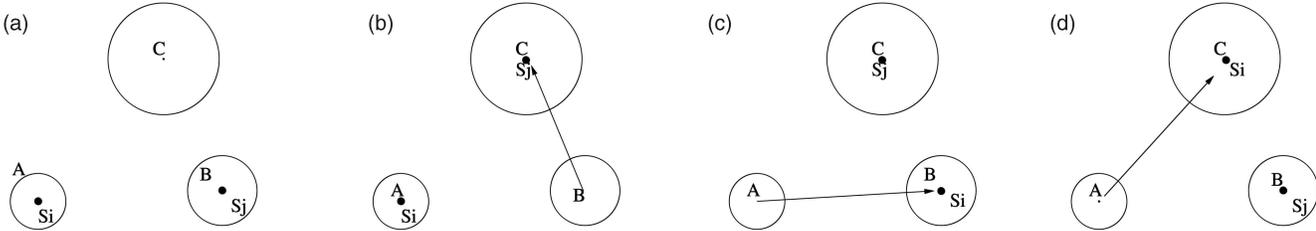


Fig. 7. Distance-based versus price-based. (a) Original situation. (b) Distance-based. (c) Distance-based. (d) Price-based.

bidder needs to determine which mobile sensor to bid among those having a lower base price than its bid. We propose two criteria for choosing mobile sensors: *distance-based* and *price-based*. In the distance-based approach, a bidder chooses the closest mobile sensor to bid; in the price-based approach, a bidder chooses the cheapest mobile sensor to bid. The advantage of the distance-based approach is shown in Fig. 6. We use a dashed circle to represent the coverage hole. The center of the circle is the target position of the mobile sensor to heal this hole. Initially, s_i is located in hole A and s_j is in hole D . The base price of s_i is higher than s_j since the size of A is larger than D . In the distance-based approach, hole C will bid s_i and hole B will bid s_j . The movement of s_i and s_j is shown in Fig. 6b. In the price-based approach, both holes C and B will bid s_j and hole C wins. Then, hole B bids s_i . Their movement is shown in Fig. 6c. As can be seen, the average moving distance of s_i and s_j is shorter in the distance-based approach because the distance-based approach helps sensors move to their closest holes.

The advantage of the price-based approach is shown in Fig. 7. With the price-based approach, hole C bids s_i since it has a lower base price. s_i moves once and no other sensor needs to move. But, with the distance-based approach, hole C bids s_j since it is closer. After sensor s_j moves to hole C , hole B needs a sensor and it will bid s_i . In this way, both s_i and s_j have to move.

3.5 Multiple Healing Detection

Due to the limited service advertisement radius, static sensors may have different knowledge about the mobile sensors. Therefore, it is possible that several static sensors independently bid different mobile sensors for the same coverage hole since the cheapest mobile sensor or the closest mobile sensor in their views is different. If more than one succeeds in bidding, multiple mobile sensors will move to heal the same hole, which is not necessary. Fig. 8a shows one example. A is the farthest Voronoi vertex of s_a and B is the farthest Voronoi vertex of s_b . Both s_a and s_b bid mobile sensors to their farthest Voronoi vertices. When both biddings are accepted, a multiple healing occurs.

We propose a self-detection algorithm for mobile sensors to solve this problem. A mobile sensor has a knowledge of the locations and base prices of other mobile sensors in its neighborhood after the service advertisement phase. If it finds out that some other mobile sensors have a higher base price than its own, it will run the detection algorithm to check whether a multiple healing has occurred. If yes, the mobile sensor will lower its base price to zero and most likely some sensor will bid it to cover a different hole.

In the detection algorithm, the detecting mobile sensor calculates a **detecting threshold** equal to

$$\pi * (d_{min} - sensing_range)^2,$$

where d_{min} is the distance to its closest neighbor. If the detecting threshold is smaller than its new base price, or d_{min} is smaller than the sensing range, a multiple healing has occurred since, without multiple healing, the calculated value should be the same as its new base price. As shown in Fig. 8b, s_e and s_f , located in A and B , respectively, are the mobile sensors bid by s_a and s_b . s_f 's new base price, the bid put forward by s_b , is calculated without considering s_e , which is $\pi * (d_{b,f} - sensing_range)^2$, where $d_{b,f}$ is the distance between s_b and s_f . Without a multiple healing, $d_{b,f}$ is just d_{min} and the calculated detecting threshold should be the same as the new base price. If multiple healing has occurred, $d_{e,f}$ is d_{min} , which is

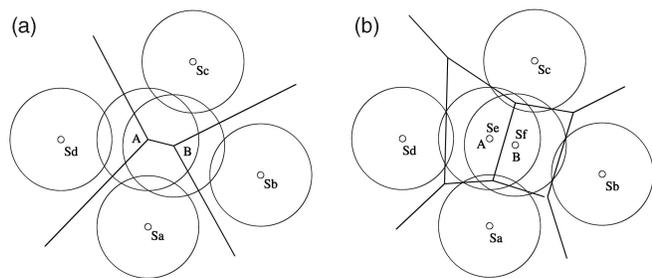


Fig. 8. Duplicate healing. (a) The duplicate healing problem. (b) Fixing the duplicate healing problem.

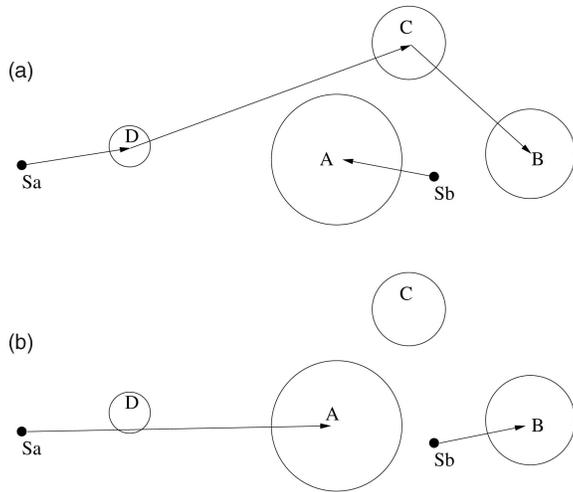


Fig. 9. Motivation of proxy-based bidding protocol. (a) The characteristic of the basic bidding protocol. (b) Ideal solution.

smaller than $d_{b,f}$, and the detecting threshold is smaller than the new base price.

4 PROXY-BASED BIDDING PROTOCOL

In this section, we present the proxy-based bidding protocol. This protocol improves the performance of the basic bidding protocol in terms of energy efficiency and load balance. In this protocol, sensors only move after their final location is determined; all calculations with respect to multiple healing detection and optimization are carried out through the exchange of messages before movement. The key trade-off is the increased number of messages versus the decreases in required movement. Because movement is typically much more expensive than exchanging messages, this protocol provides a more efficient solution than the basic bidding protocol.

4.1 General Idea: Logical Movement

Although the basic bidding protocol can achieve a high coverage, there is still room for improvement in terms of energy efficiency and load balance. In the basic bidding protocol, mobile sensors move iteratively to heal larger and larger holes. Most likely, mobile sensors will move in an irregular pattern, which consumes more energy than moving directly from their initial location to the final destination. Also, in the basic bidding protocol, some sensors are penalized by being required to move a long distance. These phenomena are illustrated by the following example, shown in Fig. 9a: In the first round, holes A, B, and C bid for mobile sensor s_b and hole D bids for s_a . Hole A and hole D win due to their large size, and these two sensors move. In the second round, hole C bids for sensor s_a ; hole B does not bid in this round since it does not know the existence of s_a due to the limited advertisement radius. In round three, hole B knows of s_a and bids for it. s_a moves the third time to reach its final location, resulting in a much longer moving distance than s_b . Ideally, s_a shall move to heal hole A and s_b moves to heal hole B, as shown in Fig. 9b. The comparison between the basic bidding protocol and the ideal solution motivates us to propose the proxy-based bidding protocol to better allocate mobile

sensors to coverage holes such that the overall moving distance is shortened and no sensor is penalized.

Following the same bidding framework, the proxy-based bidding protocol deploys the idea of *virtual movement*. Instead of moving physically in each round, mobile sensors perform virtual movements once they accept a bid. They only perform physical movements after they determine their final destinations. In this way, mobile sensors will not move in an irregular pattern. Also, virtual movement enables the possibility for mobile sensors to exchange their coverage holes to further shorten the moving distance since it does not matter which sensor heals which hole when all the largest holes are to be healed. For example, as shown in Fig. 9, through logical movement, s_a identifies hole B as its final destination and s_b identifies hole A. Before they perform the physical movements, they can exchange their destinations, i.e., s_a moves to hole B and s_b moves to hole A, such that an ideal allocation of mobile sensors to holes is obtained. In addition, with virtual movement, we can do multiple-healing detection before sensors physically move, and the vain movements of the sensors involved can be saved. In the following sections, we present the details of this protocol.

4.2 Proxy Sensor

To implement virtual movement, the first problem to be addressed is how to advertise services to the neighborhood of those virtual positions when mobile sensors do not move. One intuitive solution is to perform a network-wide broadcast. However, this may significantly increase the communication overhead. To keep the same communication overhead and let the sensors in the neighborhood of the virtual position of the mobile sensor receive the advertisement messages, we propose to use *proxy sensors*, which are static sensors located closest to the virtual positions of the mobile sensors to advertise the services and process the bidding messages for those mobile sensors.

The sensor closest to a mobile sensor's virtual position should be the bidder who detects the coverage hole and bids this mobile sensor since sensors detect coverage holes locally by checking their Voronoi cells. Therefore, we choose the winning bidder as the proxy of the mobile sensor who accepts its bid.

The proxy of a mobile sensor is not fixed during the lifetime of the mobile sensor. When a mobile sensor accepts a bid in the first round, it sends a *delegate* message to the bidder and the bidder becomes the first proxy of this mobile sensor. In the next round, the bidder (proxy) advertises the virtual position and the new base price of the mobile sensor. In the view of other sensors, the mobile sensor has moved to its virtual position and the Voronoi diagram is computed based on the new virtual position of the mobile sensor. Based on the new base price of the mobile sensor, static sensors can still bid for the mobile sensor. Their bidding messages, if any, will be sent to the proxy instead of the real mobile sensor. Based on the received bidding messages, the proxy determines which new hole should be healed. If a new bid is accepted, the proxy delegates the proxy role to that bidder who will become the new proxy of the mobile sensor. In this way, the physical movement of the mobile sensor is replaced by delegating the role of proxies between static sensors, thus realizing virtual bidding movement. When a proxy sensor does not receive any bidding messages for a *waiting*

threshold of n rounds, it will notify the mobile sensor to perform physical movement. Intuitively, this is a good value: If a sensor is not bid for in two rounds and is bid for in the third round, most likely, there are other sensors around it that have received many bids in consecutive rounds; this is highly unlikely. We experimentally determined that $n = 2$ provides good results.

In addition to virtual movement, by using a proxy sensor, *multiple healing* can be detected before it happens in many situations. After the service advertisement, proxy sensors have a service list which contains the information of the virtual positions of mobile sensors. A proxy sensor can act as the mobile sensor it represents and detect whether a *multiple healing* would happen by examining its service list with the same method of *multiple healing* detection presented in Section 3.5. A proxy sensor calculates the Voronoi cell without considering its mobile sensor, as if its bid in the previous round had failed. Then, it checks whether the original coverage hole remains; if the same hole exists, no multiple healing has occurred since its mobile sensor is required to heal the hole; otherwise, some neighbor has bid for a mobile sensor to heal the same hole and a multiple healing has occurred. If the proxy discovers that a multiple healing has occurred, it reduces the base price of its delegated mobile sensor to zero and readvertises the new service in the subsequent rounds.

To avoid all proxies from detecting the same multiple-healing and reducing the base prices of their delegated mobile sensors to zero, the proxies check whether the moving distance of its delegate from its current position to this hole is the shortest among those mobile sensors that heal the same hole. If not, it reduces the base price of its delegate to zero; otherwise, it waits for other mobile sensors to leave.

4.3 Coverage Hole Exchange

Coverage hole exchange is proposed to reduce the overall moving distance and to reduce the chance that an individual sensor is penalized by moving a long distance. It is performed by proxy sensors. A proxy sensor checks the service list obtained after the service advertisement phase and determines with which mobile sensor to exchange the virtual position of its mobile sensor. The exchange criteria will be described in depth in the next paragraph. If an exchange is necessary, the proxy sensor sends a request to the proxy of the mobile sensor with which it wants to exchange position. A proxy sensor which receives multiple exchange requests chooses one by the same criteria and sends back a confirm message. Then, these two proxy sensors exchange delegation of their mobile sensors, and the two mobile sensors exchange the proxies and their associated coverage holes.

Before presenting the exchange criteria, we introduce the following notations: We use d_i to represent the moving distance of s_i before exchange, and \hat{d}_i the moving distance after exchange. d_{max} is a maximum moving distance threshold. All exchanges between s_i and s_j must satisfy the following prerequisites; otherwise, the exchange will not be performed:

$$\begin{cases} \hat{d}_i + \hat{d}_j \leq d_i + d_j \\ \hat{d}_i \leq \max[d_i, d_{max}] \\ \hat{d}_j \leq \max[d_j, d_{max}] \end{cases} \quad (1)$$

As shown in (1), all exchanges must reduce the overall distances. Also, the exchanges must not increase the moving distance of a single sensor to be longer than d_{max} if it is not so before the exchange and must not further increase the moving distance of a single sensor if its moving distance is already longer than d_{max} .

Among the exchanges which satisfy the prerequisites shown in (1), we give higher priority to those which can release or mitigate node penalization. We first check exchanges, in which one or both involved sensors have to move longer than d_{max} before the exchange, and choose the one which can reduce the overall moving distance the most.

Formally, the exchange is chosen as follows:

$$[s_i, s_j] = \operatorname{argmin}_{\hat{d}_k + \hat{d}_l - d_k - d_l} \{[s_k, s_l] : d_k \geq d_{max} \vee d_l \geq d_{max}\}. \quad (2)$$

Here, $[s_i, s_j]$ indicates mobile sensor s_i exchanges its virtual position with s_j . If there is no such exchange, we choose the exchange which can reduce the overall moving distance the most. That is,

$$[s_i, s_j] = \operatorname{argmin}_{\hat{d}_k + \hat{d}_l - d_k - d_l} \{[s_k, s_l] : d_k \leq d_{max} \wedge d_l \leq d_{max}\}. \quad (3)$$

Without hole exchange, proxy sensors can notify mobile sensors to move if they do not receive bidding messages for the *waiting threshold* of n rounds. Hole exchange complicates the decision of when to tell a sensor to move. As shown in Fig. 9, if s_b moves physically in the third round, s_a has no sensor with which to exchange its virtual position after it virtually moves to hole B. To solve this problem, s_b should wait for more rounds before movement. In general, a mobile sensor that gets a high base price in the first two rounds should wait for additional rounds before physically moving so that other sensors have an opportunity to perform hole exchange. Through extensive experiments, we determined that $n = 5$ for sensors that receive high bid prices in the first two rounds,¹ and $n = 2$ for other sensors, yields good results.

There is an exception to this general principle. For very large holes, i.e., holes bigger than the sensing range of a single sensor, as shown in Fig. 5, two mobile sensors (or more) are needed for healing. In Fig. 5, static sensor s_a bids for mobile sensor s_d to move, and it is s_b which bids for another mobile sensor s_e to heal the same hole. Normally, sensors that move first to heal the hole act as bidders in the next rounds to bid for more sensors to heal the same hole. These sensors, like s_d , which move first and have the maximum base price $\sqrt{(3)} * \pi * \text{sensing-range}^2 / 2$ (described in Section 3.3), should move immediately because they will act as bidders.

4.4 Protocol Specification

As with the basic bidding protocol, the proxy-based bidding protocol runs round by round until mobile sensors obtain their final locations and move there directly. Each round consists of four phases: service advertisement, bidding, virtual movement, and hole-exchange.

1. For a higher value, the latency is unnecessarily increased; for a lower value, the problem shown in Fig. 9 will frequently occur. For different configurations, different values of n will sometimes lead to better results, but on average, we found this value to be best.

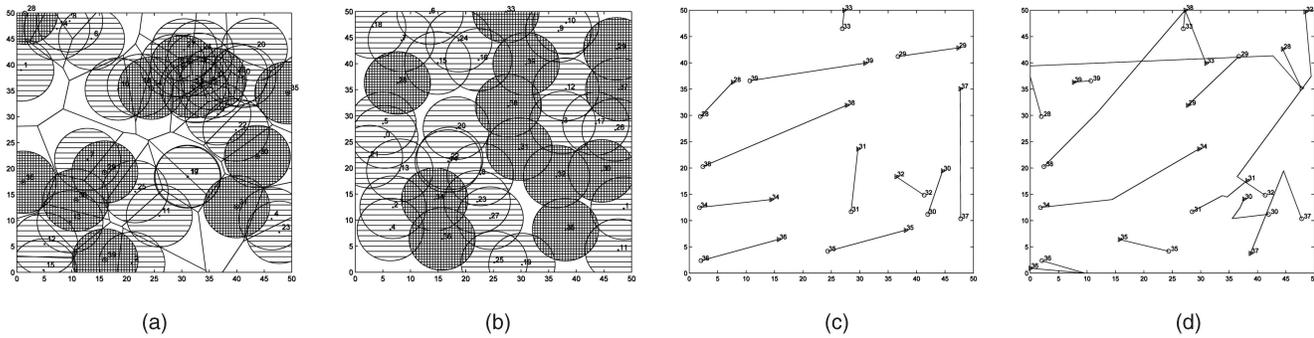


Fig. 10. An operational example. (a) Initial distribution. (b) After deployment. (c) Trace of the proxy-based protocol. (d) Trace of the basic bidding protocol.

1. In the service advertisement phase, proxy sensors advertise the virtual locations, physical locations, and base prices for their delegated mobile sensors. In the first round, a mobile sensor does not have a proxy and advertises its physical location and base price by itself.
2. In the bidding phase, static sensors calculate their Voronoi cells based on the virtual positions of mobile sensors. They detect coverage holes by examining the Voronoi cells, estimate the hole size, choose the closest or cheapest mobile sensor, and send bidding messages to its proxy or the mobile sensor itself if the mobile sensor has no proxy.
3. In the virtual movement phase, proxy sensors (or mobile sensors without a proxy) choose the highest bid and send a delegate message to the bidder. The bidder becomes the new proxy. The base price of mobile sensors is updated by their new proxies. Also, proxy sensors need to check whether hole-exchange is needed. If yes, they choose the mobile sensor suitable for exchange and send out an exchange request to the proxy of that mobile sensor.
4. In the hole-exchange phase, proxy sensors check the received requests, choose one with the highest priority, and return the confirm message to the requester. Then, the mobile sensors delegated by these two proxy sensors exchange the hole to heal.

The protocol terminates naturally when all the largest holes are healed and no more hole exchanges are necessary. Through the bidding process, when no sensors can raise a bid higher than the lowest base price of mobile sensors, all the largest holes are healed. This process terminates naturally as presented in Section 3. For hole exchange, we require that all the exchanges must reduce the overall moving distance. There is a lower bound of the overall moving distance, and hole exchange will finish naturally. Through this iterative 4-phase process, proxy sensors notify mobile sensors to move and the deployment process terminates. We show the formal algorithm in Appendix B.

We show an operational example to illustrate the advantage of the proxy-based protocol over the basic bidding protocol. Forty sensors, of which 30 percent are mobile, are randomly distributed in a 50 m * 50 m field. The initial distribution is shown in Fig. 10a; the distribution after deployment is shown in Fig. 10b. In this example (and most others), the proxy-based bidding protocol and the basic

bidding protocol get the same distribution of sensors after deployment. Fig. 10c shows the moving trace of mobile sensors in the proxy-based bidding protocol. The mobile sensors move 13.65 m on average. Sensor 38 moves the longest distance 27.85 m. Fig. 10d shows the moving trace of mobile sensors in the basic bidding protocol. The average moving distance is 23.77 m. Sensor 28 has the longest moving distance. It moves five times for a total distance of 68.68 m. From this example, we can see that the proxy-based protocol is more energy-efficient and load-balanced.

5 PERFORMANCE EVALUATIONS

5.1 Objectives, Metrics, and Methodology

Our deployment protocols are implemented in the ns-2 (version 2.1b9a). We have three objectives in conducting this evaluation: first, justifying our proposal of constructing sensor networks with both mobile and static sensors to balance cost and sensing coverage; second, testing the effectiveness of our bidding protocols in providing high coverage; finally, comparing the basic bidding protocol and the proxy-based bidding protocol and giving some insight on choosing deployment protocols.

The performance of our schemes is evaluated from three aspects: *sensor cost*, *deployment quality*, and *energy consumption*. Sensor cost is measured by the money used to construct the network. Deployment quality and energy consumption are measured by the same metrics as in [14]. In particular, deployment quality is measured by the sensor coverage and the time (number of rounds) to reach this coverage. Energy consumption is measured by the message complexity, which is an indicator of the energy consumption in communication, and moving distance and the number of movements, which are indicators of the energy consumption in mechanical movement.

We run simulations for different compositions of sensor networks, and determine the coverage that can be reached. In a 60 m * 60 m flat field, we randomly distribute 60 sensors. Among these sensors, we assign a percentage of sensors to be mobile. This percentage varies from 10 percent to 50 percent, with an increment of 10 percent. The mobile sensors are chosen randomly. To evaluate each metric under different parameter settings, we run 50 experiments based on different initial distributions and calculate the average results.

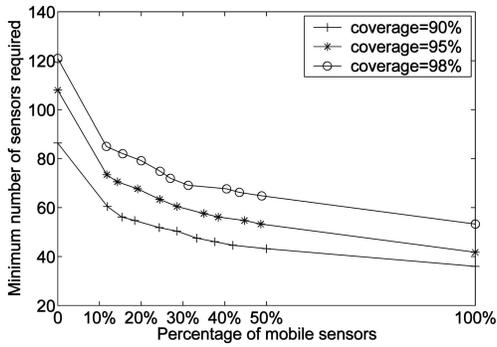


Fig. 11. The number of sensors needed to reach certain coverage under different mobile percentage.

We choose 802.11 as the MAC layer protocol and DSDV as the routing protocol. The transmission range is set to be 20 meters. Based on the information from [2], we set the *sensing range* to be 6 meters. This is consistent with other current sensor prototypes, such as Smart Dust (UC Berkeley), CTOS dust, and Wins (Rockwell) [3].

In the following sections, we show the simulation results.

5.2 Trade-Off between Cost and Coverage

In order to evaluate the trade-off between sensor cost and coverage, we consider three cases of network composition: all the sensors are mobile, all the sensors are static, and a percentage of sensors are mobile. When all the sensors are static, random deployment is used. When all the sensors are mobile, the VOR protocol [14] is used for sensor deployment. When a percentage of sensors are mobile, our basic bidding protocol (using the distance-based criteria) is used. Fig. 11 shows the total number of sensors needed to reach certain coverage with under different network compositions. 0 percent of mobile sensors means that all sensors are static.

As shown in Fig. 11, to reach a certain coverage, random deployment of static sensors uses the most number of sensors; as the percentage of mobile sensors increases, the required number of sensors to reach a certain coverage decreases; a deployment of 100 percent mobile sensors requires the fewest sensors. However, the cost of mobile sensors may be high.

Compared to random deployment, the basic bidding protocol can significantly reduce the number of sensors required to reach a certain coverage. For example, to reach a 90 percent coverage with only 10 percent of mobile sensors, the basic protocol needs 30 percent fewer sensors; when 50 percent of the sensors are mobile, the required number of sensors is reduced by 50 percent.

Compared to the case in which 100 percent of the sensors are mobile, to reach 90 percent coverage, the basic bidding protocol requires 40 percent fewer mobile sensors in the case in which 50 percent sensors are mobile. Note that the cost of mobile sensors is higher than static sensors, so the overall cost of using a percentage of mobile sensors may be reduced even though more sensors in total are used.

Fig. 12 shows the sensor cost of these three protocols to reach a certain sensor coverage. Based on the *cost ratio* between the mobile sensor and the static sensor, the overall sensor cost of these three protocols may be different. Intuitively, if the cost ratio is low (e.g., 1.5), increasing the percentage of mobile sensors can reduce the overall sensor cost. On the other hand, if the mobile sensors are very expensive, using only static sensors may have the lowest sensor cost (not shown in the figure). When the cost ratio is somewhere in the middle, the basic protocol, which has a mix of mobile and static sensors, can achieve the lowest sensor cost. For example, when the cost ratio is 3.5, to reach 95 percent coverage, the basic protocol has the lowest cost when 10 percent sensors are mobile. Currently, the cost of a static sensor prototype Motes is about \$100 and the cost of a mobile sensor prototype is about \$200 [10]. The ratio is expected to increase under mass production. Based on this figure, we can see that there is a trade-off between cost and coverage. The basic protocol can achieve a balance between these two most of the time.

5.3 Comparing the Protocols

We consider four cases: Both the proxy and basic bidding protocols using both distance and price-based criteria. In the figures showing the simulation results, we use “Proxy-distance,” “Proxy-price,” “Basic-distance,” and “Basic-price” to represent them, respectively.

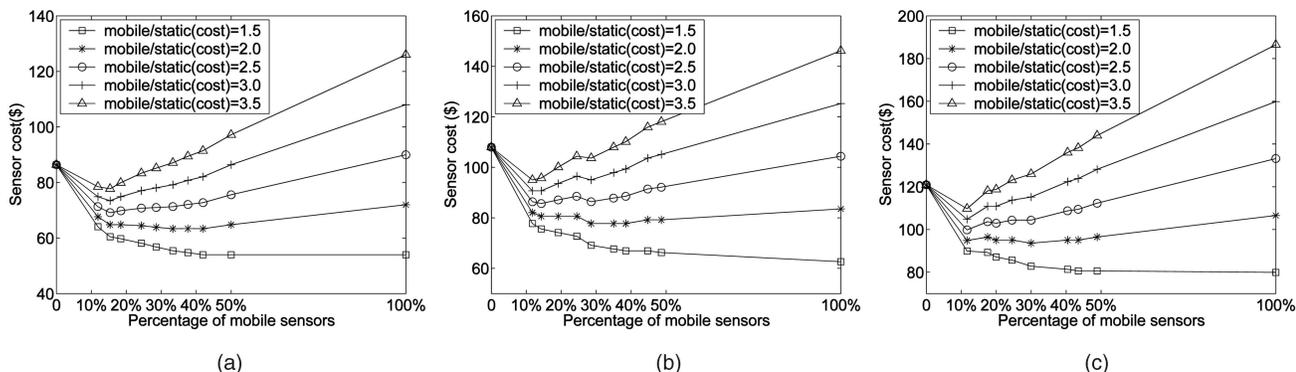


Fig. 12. The cost of sensors to reach certain coverage. (a) To reach 90 percent coverage. (b) To reach 95 percent coverage. (c) To reach 98 percent coverage.

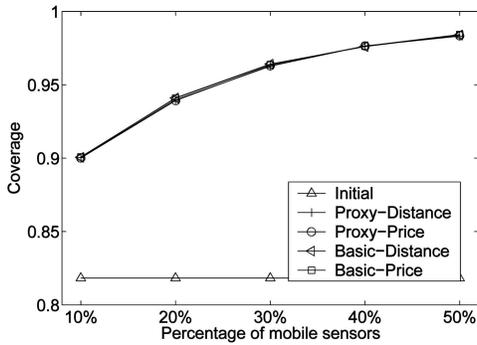


Fig. 13. Coverage.

5.3.1 Coverage

Fig. 13 shows the coverage obtained by our protocols under different mobile sensor percentage. We can make two observations from the figure. One is that our bidding protocols can increase the coverage significantly. The other is that all four cases we consider achieve very similar coverage. All the four cases follow the same bidding framework and heal the largest holes. In terms of coverage, there is no preference between the basic-bidding protocol and the proxy-based bidding protocol; there is no preference between distance-based criteria and price-based criteria to choose mobile sensors.

5.3.2 Termination

When all the largest holes are healed and no sensor can give a higher bid than the lowest base price of mobile sensors, the protocols terminate. Fig. 14 shows the number of rounds that the protocols have run when the protocols terminate. As expected, the proxy-based bidding protocol requires more rounds to terminate. In the proxy-based bidding protocol, each sensor waits several rounds before physical movement and sensors spend a number of rounds on exchanging holes. However, because physical movements will likely dominate the recursion time and the proxy-based protocol reduces movement, it may still terminate in the shortest time. In both the proxy-based protocol and the basic protocol, using distance-based criteria or price-based criteria does not significantly affect the termination time.

The deployment rounds are increased when the mobile sensor percentage increases. With more mobile sensors available, the allocation of mobile sensors to coverage holes is more complicated and needs more rounds.

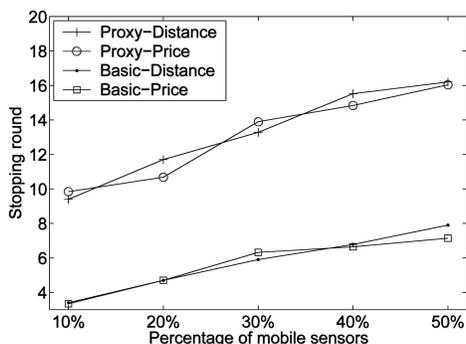


Fig. 14. Termination.

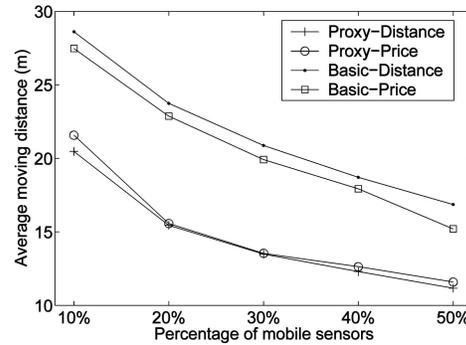


Fig. 15. Moving distance.

5.3.3 Energy Consumption

Energy consumption includes two parts, mechanical movement and communication. We use message complexity to measure the energy consumed in communication; we use the number of movements and moving distance to measure the energy consumption in movement. We first show the performance of our protocols in these three metrics. Then, we show a unified energy consumption considering all these metrics.

Fig. 15 shows the moving distance. As expected, the moving distance is much lower when the proxy-based bidding protocol is used. Between the distance-based criteria and price-based criteria, the moving distance is quite similar when using the proxy-based bidding protocol, and it is shorter when the latter criteria is used in the basic bidding protocol. The figure tells us that the phenomena shown in Fig. 7 are dominant compared to those shown in Fig. 6. In the proxy-based bidding protocol, these two criteria achieve similar performance. For most cases, the hole exchange and virtual movement change the situations illustrated in Fig. 7 and Fig. 6 to an ideal case. Therefore, these two criteria achieve a similar performance.

When considering the number of movements versus the percentage of mobile sensors (not shown), we find that the number of movements required does not change as the percentage of mobile sensors increases. In addition, both the distance-based criteria and price-based criteria perform the same when using the proxy-based protocol (about 1.1). When using the basic bidding protocol, the price-based criteria achieves a smaller number of movements (about 1.45) than the distance-based criteria (about 1.6) for the same reason as presented in the above paragraph.

Fig. 16 shows the message complexity. The proxy-based protocol has higher message complexity than the basic protocol since it needs more rounds to terminate and needs to negotiate how to exchange holes. As mobile sensor percentage increases, the number of rounds increases, and message complexity increases accordingly.

To get a clear picture of energy consumption, we normalize the moving distance and the number of movements into message complexity. From the Robomote specification [33], approximately, to move a sensor one meter consumes a similar amount of energy as transmitting 300 messages. We set the energy consumption in starting/braking to be the same as that in moving one meter. Fig. 17 shows the unified energy consumption. As expected, the proxy-based bidding protocol consumes much less energy

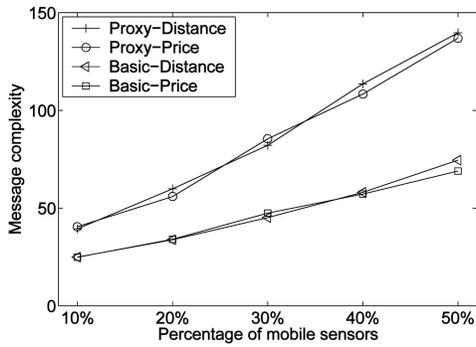


Fig. 16. Message complexity.

than the basic bidding protocol. Though sensors spend more energy in communication, they save much energy in movement. Mechanical movement is the dominant factor in energy consumption. Thus, the proxy-based protocol is much more energy efficient than the basic bidding protocol.

5.4 Load Balance

The maximum moving distance among the mobile sensors is an indication of whether individual sensors are penalized in terms of moving distance. Our simulations show that the maximum moving distance is about 39 m ~ 42 m in the proxy-based protocol, which is much shorter than the 60 m ~ 80 m in the basic bidding protocol.

6 RELATED WORK

In this section, we introduce related work in sensor coverage, static sensor deployment, mobile sensor deployment, and relay node placement in sensor networks.

6.1 Coverage

Meguerdichian et al. presented several interpretations of coverage in sensor networks, including deterministic coverage and stochastic coverage [27]. Also, the authors proposed a centralized polynomial time algorithm for coverage calculation. Another metric of sensor coverage, exposure, was defined in [26]. The authors also designed a centralized algorithm for calculating the minimal exposure paths.

6.2 Sensor Deployment

All previous work on sensor deployment either assumes all sensors are static or assumes all sensors are mobile. In the following, we first introduce papers on static sensor deployment followed by papers on mobile sensor deployment. Deployment of static sensor networks has been addressed in [7], [11]. Clouqueur et al. proposed to deploy sensors in several steps and assumed random deployment in each step [7]. The number of sensors in each step and the cost of deployment were used as a cost function. The authors proposed algorithms to determine the number of steps of sensor deployment such that the cost is low and the desired distribution is obtained. Dhillon et al. proposed a centralized polynomial-time algorithm to determine sensor distribution such that a minimum number of sensors are deployed and a minimum amount of data is transmitted [11].

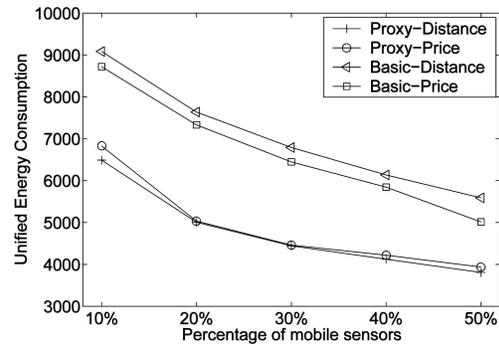


Fig. 17. Unified energy consumption.

Deployment of mobile sensors has been addressed in [14], [21], [34], [15], [16]. The work in [34] assumes that a cluster head is available to collect the sensor location and determine the target location of the mobile sensors. Howard et al. proposed an algorithm to deploy mobile sensors into a building from outside, in which sensors are deployed iteratively one by one, utilizing the location information obtained from the previous deployment [21]. The same authors proposed algorithms based on potential field to maximize the monitoring field in [22]. Wang et al. proposed three algorithms, VEC, VOR, Minimax, and two protocols to deploy mobile sensors to increase the coverage considering energy efficiency and deployment time. The authors gave insight on how to choose the algorithm and the protocol under different system requirements [14].

The only work, to our knowledge, that addressed a mixed of mobile and static sensors is our preliminary result of [15], [16].

6.3 Other Related Work

Other related work includes the study of heterogeneous networks in which not all sensors are the same, for example, networks that have both sensor nodes and relay nodes, which only have communication capability. Hou et al. proposed a centralized polynomial-time heuristic algorithm for relay node placement to increase network lifetime [20]. Patel et al. designed centralized deployment strategies for sensor nodes, relay nodes, and base stations considering connectivity and coverage [30].

7 CONCLUSIONS

In this paper, we proposed using a mix of mobile and static sensors to construct sensor networks to balance cost and sensor coverage. We identified the problem of deploying mobile sensors in a mixed sensor network as an NP-complete problem and designed bidding protocols to tackle this problem in a distributed fashion, in which static sensors act as bidders and mobile sensors act as hole-healing servers. Intensive simulation justified our idea of deploying both mobile and static sensors. Users can determine the percentage of mobile sensors to get the most economical deployment of sensors to construct a network satisfying the coverage requirement. Simulation results also showed the efficiency and effectiveness of our proxy-based bidding protocol in placing mobile sensors to achieve high coverage.

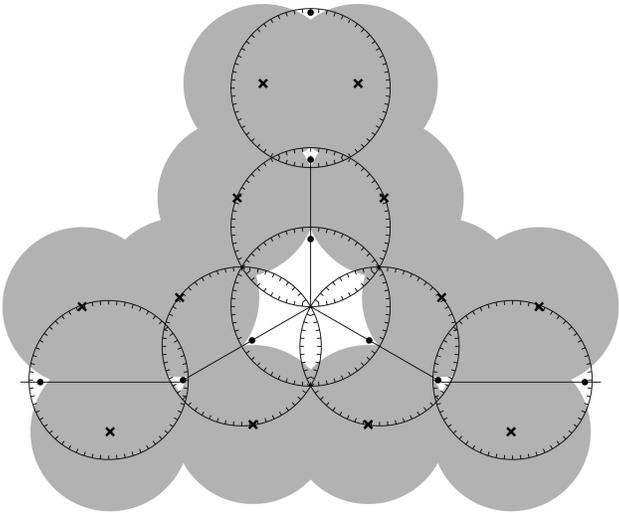


Fig. 18. Gadget for NP-completeness proof. Dots indicate the special points used in the proof; crosses are the locations of already placed sensors.

In the future, we will work on the deployment of mobile sensors for nonuniform coverage requirements, or for purposes other than coverage. In many applications, some locations are more important than others and may require more sensors for coverage. The bidding protocols presented here can be adapted to this scenario by modifying the rules of assigning base and bid price. In addition, we believe the bidding protocol can be used in many other applications, such as distributed resource allocation.

APPENDIX A

To prove that this problem is NP-hard, we will reduce the following question:

Given a cubic planar graph G and integer k , does G have an independent set of size k ?

to

Given a square target field partially covered with a number of unit circles (i.e., of radius 1) and integer k , can we obtain a complete coverage of the target field by adding k unit circles?

The first question is NP-complete [19], [18].

The second question can be clearly reduced to our optimization problem.

The first stage in translating the problem is to draw the given graph G on an integer grid, which is shown in Fig. 18. We will request that each node has both coordinates divisible by, say, 10.

For each node of the original graph, we have a point $(10i, 10j)$; near that point, we center a unit circle and inside we create an uncovered area as shown in the figure below; note that this area has three special points and that nearby are another three points, each in distance 1 from a corresponding special point; we call them *outer special points*. Outer special points will be located at lines in which at least one coordinate is divisible by 10.

Each edge of the original graph corresponds to a line in which points have at least one coordinate which is an integer divisible by 10. We cover this path with points that

are in distance, say, between 1.5 and 1.75 from each other, so that such a "trail of points" starts at one of the special points of a node gadget and ends at an outer special node of another gadget. The trails of points of each edge must be disjoint, and each must have an odd number of points.

We create a little uncovered area around each point on our trails. We finish the construction by covering all areas that we explicitly did not wish to leave uncovered. It follows from the figure that it can be done.

Now, suppose that we had n nodes and m edges in the original graphs and the trails of points of the edges together contain $2K - m$ points. Then, we ask if we can achieve the complete coverage by adding $K + n - k$ circles.

Suppose that the original question has answer "YES," i.e., there exists an independent set with k nodes and, thus, a vertex cover with $n - k$ nodes. In gadgets corresponding to the vertices of the vertex cover, we cover the central area (with the three inner special nodes) with a single circle. In gadgets corresponding to the vertices of the independent sets, we use three circles to cover the central area and the areas of the outer special points. Now, consider an edge; it corresponds to a trail of, say, $2h - 1$ points, two of them being outer special points. One of these points is covered by a circle used because of the independent set, so we have $2h - 2$ points left and the uncovered areas around these points are placed in such a way that we can cover pairs of them in one circle. Thus, we use $h - 1$ circles, and to these circles, we can add the circle placed by the independent set rule, so we attribute the use of h circles to this edge. If we add together all circles used that way, we get K circles. Hence, we covered the entire area with $K + n - k$ circles and the new question has answer "YES."

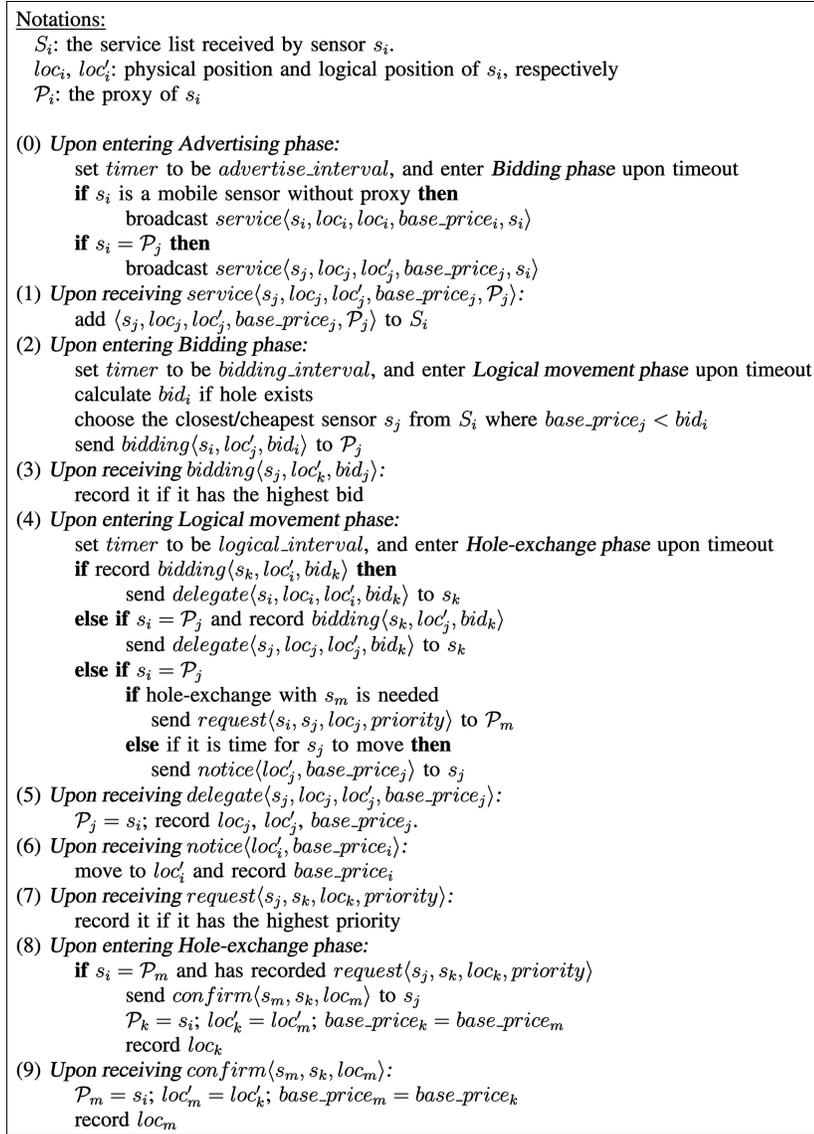
Now, suppose that the new question has answer "YES," so that we obtain a complete coverage by adding $K + n - k$ circles. We will change the placement of the new circles without changing their number to assure some good properties of the placement.

Consider a vertex gadget and its three inner special points.

Suppose that two of these points are covered with a single circle; then, this circle cannot cover any of the outer special points. We move this circle so it covers the entire central area, in particular, all three inner points. We had to have three circles with outer special points; we move them so they cover the area of these points as well as the areas of the adjacent points on their respective trails. We call such a vertex a *cover vertex*.

Suppose that the inner special points were covered with three different circles. We move these circles so they cover the central area and the areas of the outer special points. Consider a trail of points of an edge, say, with $2h - 1$ points. Suppose that these points are covered with $h + 1$ (or more) circles. Then, we remove all the circles that cover the areas of these points and we cover them together with the entire inner area and the areas of the outer special points. We call such a vertex an *independent vertex*.

Now, consider a trail of points of an edge $\{u, v\}$, say, with $2h - 1$ points. Suppose both u and v are independent; in this case, both of the outer special points that are at the ends of this trail are covered together with their respective inner points, leaving $2h - 3$ points to cover. We remove the

Fig. 19. The proxy-based protocol at sensor s_i .

circles that cover these points as well as the outer special point of v that is on the trail; because the latter was covered together with its respective inner point, we are removing the cover of at least $2h - 1$ points, and they are placed in such a way that we surely remove at least h circles. Now, we cover $2h - 2$ points on the trail with $h - 1$ circles and we use one more circle to cover three inner points of the gadget of v . As a result, we changed the classification of v to *cover vertex* without increasing the number of circles.

Now, independent vertices form an independent set and cover vertices form a vertex cover. When we consider a trail of points of an edge with $2h - 1$ points, we cover them with h circles, together with an inner special point of the incident independent point (if any). Thus, we cover the trails of points of edges together with the gadgets of the independent points with K circles, and the remaining $n - k$ circles cover gadgets of the cover vertices, which, in turn, form a vertex cover. Hence, the answer to the original question is "YES," we do have an independent set of k nodes.

APPENDIX B

The proxy-based protocol at sensor s_i ; see Fig. 19.

REFERENCES

- [1] "US Naval Observatory (USNO) GPS Operations," <http://tycho.usno.navy.mil/gps.html>, Apr. 2001.
- [2] "Berkeley Sensor and Actuator Center," <http://www-bsac.eecs.berkeley.edu>, 2004.
- [3] "Wireless Sensing Networks," <http://wins.rsc.rockwell.com>, 2004.
- [4] I.F. Akyildiz, W. Su, Y. Sankarasubramaniam, and E. Cayirci, "A Survey on Sensor Networks," *IEEE Comm. Magazine*, pp. 102-114, Aug. 2002.
- [5] I. Akyildiz, W. Su, Y. Sankarasubramaniam, and E. Cayirci, "Wireless Sensor Networks: A Survey," *Computer Networks*, vol. 38, no. 4, Mar. 2002.
- [6] F. Aurenhammer, "Voronoi Diagrams—A Survey of a Fundamental Geometric Data Structure," *ACM Computing Surveys*, vol. 23, pp. 345-405, 1991.
- [7] T. Clouqueur, V. Phipatanasuphorn, P. Ramanathan, and K.K. Saluja, "Sensor Deployment Strategy for Target Detection," *Proc. Int'l Workshop Wireless Sensor Networks and Applications*, 2002.

- [8] D. Koditschek, "Planning and Control via Potential Functions," *Robotics Review I*, pp. 349-367, 1989.
- [9] D. Niculescu and B. Nath, "Ad Hoc Positioning Systems (APS) Using AoA," *Proc. IEEE INFOCOM*, 2003.
- [10] K. Dantu, M. Rahimi, H. Shah, S. Babel, A. Dhariwal, and G.S. Sukhatme, "Robomote: Enabling Mobility in Sensor Network," Technical Report CRES-04-006, CRES, 2004.
- [11] S. Dhillon, K. Chakrabarty, and S. Iyengar, "Sensor Placement for Grid Coverage under Imprecise Detections," *Proc. Int'l Conf. Information Fusion*, 2002.
- [12] S. Fortune, D. Du, and F. Hwang, "Voronoi Diagrams and Delaunay Triangulations," *Euclidean Geometry and Computers*, 1992.
- [13] G. Wang, "Dependable Sensor Networks," PhD dissertation, Dept. of Computer Science and Eng., Pennsylvania State Univ., May 2006.
- [14] G. Wang, G. Cao, and T.F. La Porta, "Sensor Deployment Protocols," *IEEE Trans. Mobile Computing*, to appear.
- [15] G. Wang, G. Cao, and T.F. La Porta, "A Bidding Protocol for Deploying Mobile Sensors," *Proc. 11th IEEE Int'l Conf. Network Protocols (ICNP '03)*, Nov. 2003.
- [16] G. Wang, G. Cao, and T.F. La Porta, "Proxy-Based Sensor Deployment for Mobile Sensor Networks," *Proc. IEEE Int'l Conf. Mobile Ad-Hoc and Sensor Systems (MASS '04)*, Oct. 2004.
- [17] G. Wang, M.J. Irwin, P. Berman, H. Fu, and T.F. La Porta, "Optimizing Sensor Movement for Energy Efficiency," *Proc. Int'l Symp. Low Power Electronics and Design (ISPLED '05)*, Aug. 2005.
- [18] M.R. Garey and D.S. Johnson, *Computers and Intractability*. W.H. Freeman & Co., 1979.
- [19] M.R. Garey, D.S. Johnson, and L. Stockmeyer, "Some Simplified NP-Complete Graph Problems," *Theoretical Computer Science*, G. Ausiello and D. Sannella, eds., Elsevier, 1976.
- [20] Y.T. Hou, Y. Shi, H.D. Sherali, and S.F. Midkiff, "Prolonging Sensor Network Lifetime with Energy Provisioning and Relay Node Placement," *Proc. IEEE Comm. Soc. Conf. Sensor and Ad Hoc Comm. and Networks (SECON '05)*, 2005.
- [21] A. Howard, M.J. Mataric, and G.S. Sukhatme, "An Incremental Self-Deployment Algorithm for Mobile Sensor Networks," *Autonomous Robots*, special issue on intelligent embedded systems, 2002.
- [22] A. Howard, M.J. Mataric, and G.S. Sukhatme, "Mobile Sensor Networks Deployment Using Potential Fields: A Distributed, Scalable Solution to the Area Coverage Problem," *Proc. Sixth Int'l Symp. Distributed Autonomous Robotics Systems*, June 2002.
- [23] *Handbook of Sensor Networks: Algorithms and Architectures*, I. Stojmenović, ed., IEEE Computer Society, Oct. 2005.
- [24] J. Lengyel, M. Reichert, B. Donald, and D. Greenberg, "Real-Time Robot Motion Planning Using Rasterizing Computer Graphics Hardware," *Proc. SIGGRAPH*, 1990.
- [25] Q. Li, M. De Rosa, and D. Rus, "Distributed Algorithms for Guiding Navigation across a Sensor Network," *Proc. ACM MobiCom*, 2003.
- [26] S. Meguerdichian, F. Koushanfar, G. Qu, and M. Potkonjak, "Exposure in Wireless Ad-Hoc Sensor Networks," *Proc. ACM MobiCom*, 2001.
- [27] S. Meguerdichian, F. Koushanfar, M. Potkonjak, and M.B. Srivastava, "Coverage Problems in Wireless Ad-Hoc Sensor Network," *Proc. IEEE INFOCOM*, 2001.
- [28] N. Patwari and A. Hero III, "Using Proximity and Quantized RSS for Sensor Location in Wireless Location in Wireless Networks," *Proc. Workshop Wireless Sensor Networks and Applications*, 2003.
- [29] C.H. Papadimitriou and K. Steiglitz, *Combinatorial Optimization: Algorithms and Complexity*. Dover Publications, 1998.
- [30] M. Patel, R. Chandrasekaran, and S. Venkatesan, "Energy Efficient Sensor, Relay and Base Station Placements for Coverage, Connectivity and Routing," *Proc. 24th IEEE Int'l Performance, Computing and Comm. Conf.*, 2005.
- [31] G.J. Pottie, "Wireless Sensor Networks," *Proc. Information Theory Workshop*, June 1998.
- [32] A. Savvides, C. Han, and M.B. Srivastava, "Dynamic Fine-Grained Localization in Ad-Hoc Networks of Sensors," *Proc. ACM MobiCom*, 2001.
- [33] G.T. Sibley, M.H. Rahimi, and G.S. Sukhatme, "Robomote: A Tiny Mobile Robot Platform for Large-Scale Sensor Networks," *Proc. IEEE Int'l Conf. Robotics and Automation (ICRA '02)*, 2002.
- [34] Y. Zou and K. Chakrabarty, "Sensor Deployment and Target Localization Based on Virtual Forces," *Proc. IEEE INFOCOM*, 2003.



Guiling (Grace) Wang received the BS degree from Nankai University, China. She received the PhD degree in computer science and engineering from Pennsylvania State University in May 2006. She joined the New Jersey Institute of Technology as an assistant professor after that. Her research interests include distributed systems, wireless networks, and mobile computing with a focus on wireless sensor networks.



Guohong Cao received the BS degree from Xian Jiaotong University, Xian, China, and the MS and PhD degrees in computer science from the Ohio State University in 1997 and 1999, respectively. Since then, he has been with the Department of Computer Science and Engineering at Pennsylvania State University, where he is currently an associate professor. His research interests are wireless networks and mobile computing. He has published 100 papers in the areas of sensor

networks, cache management, data dissemination, resource management, wireless network security, and distributed fault-tolerant computing. He is an editor of the *IEEE Transactions on Mobile Computing* and *IEEE Transactions on Wireless Communications*, a co-guest editor of the *ACM/Kluwer Mobile Networking and Applications* special issue on heterogeneous wireless networks, and has served on the program committees of many conferences. He was a recipient of the Presidential Fellowship at the Ohio State University in 1999 and a recipient of the US National Science Foundation CAREER award in 2001.



Piotr Berman received the master of mathematics degree from Warsaw University in 1978 and the PhD degree in applied mathematics from the Massachusetts Institute of Technology (MIT) in 1985. Afterward, he became a faculty member in the Department of Computer Science and Engineering at Pennsylvania State University. His research interest is in proving NP-hardness.



Thomas F. La Porta received the PhD degree in electrical engineering from Columbia University, New York. He joined the Computer Science and Engineering Department at Pennsylvania State University in 2002 as a full professor. He is currently the director of the Networking and Security Research Center at Penn State. Prior to that, he was with Bell Laboratories since 1986, where he was the director of the Mobile Networking Research Department leading various projects in wireless and mobile networking. He was the founding editor-in-chief of the *IEEE Transactions on Mobile Computing*. He also served as editor-in-chief of *IEEE Personal Communications Magazine* for three years. He has published more than 50 technical papers and holds 28 patents. He was an adjunct member of the faculty at Columbia University for seven years, where he taught courses on mobile networking and protocol design. He received the Bell Labs Distinguished Technical Staff Award in 1996 and is a fellow of the IEEE and Bell Labs.

► For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.