Title

Actuation Techniques for Sensing Uncertainty Reduction

Authors
Aman Kansal
William J Kaiser
Gregory J Pottie
et al.

Publication Date
2005

Actuation Techniques for Sensing Uncertainty Reduction

AMAN KANSAL, WILLIAM J KAISER, GREGORY J POTTIE, and MANI B SRIVASTAVA
University of California, Los Angeles

The information acquisition performance of a sensor network is critical to all applications based on it. This performance depends on factors which cannot be completely known at design or deployment time: sensing medium characteristics and the phenomenon distribution. Simplifying assumptions such as the homogeneous nature of sensing media do not hold in most practical scenarios due to the presence of sensing obstacles. Further, the medium and phenomena may change over time. We propose to use controlled mobility to enhance coverage at run time in an autonomous manner. However, extensive robotic capabilities and supporting services such as precise navigation may be infeasible in large scale sensor networks. We present feasible alternatives for physical reconfiguration using low complexity and low energy actuation. The key contribution of the paper is to show that even small degrees of actuation can lead to a significant coverage advantage. We also compare this approach to conventional means for achieving equivalent coverage by increasing node density without actuation. Further, we discuss the relevant trade-offs which affect the use of mobility in terms of the time required for actuation.

Categories and Subject Descriptors: C.4 [**Computer Systems Organization**]: Performance of Systems

General Terms: Design, Reliability, Measurement

Additional Key Words and Phrases: Mobile or actuator systems, Coverage, Sensing, Actuation

1. INTRODUCTION

In this paper, we address sensing performance concerns that arise in realistic sensor network deployments. Sensing performance is an important attribute for any sensor network. Several deployment specific factors lead to uncertainty in the quality of sensor data. This uncertainty prone data is the basis for all the applications based on the sensor network, and hence directly affects the utility delivered to the user. The problem of interest then is to help reduce the uncertainty in collected data.

The sensor density required to reliably cover a region depends on the application used for processing the collected data. In certain situations this density may be too high in terms of cost of the sensors, the cost of installation or the induced interference with the existing environment. The problem is further exacerbated by

Contact: kansal@ee.ucla.edu

the fact that the sensing medium is typically not homogeneous and the network must deal with obstacles and other anisotropies which reduce the region covered by each sensor. These medium characteristics may change over time due to the movement of environmental entities, such as, due to growth of foliage in outdoor environments or movement of people and objects in indoor ones. Apart from the medium, the coverage quality depends on the phenomenon distribution itself. This distribution is unlikely to be known at design time, and in many cases even at deployment time. Even when known, manually installing a system to match this distribution is not a preferred alternative in large scale deployments. Further, the phenomenon distribution itself may change over time and a static deployment may be ineffective.

1.1 Key Contributions

To overcome the above problems we propose to use mobile or actuated sensor nodes, enabling the network to reconfigure itself in response to its environment. The sensor network could learn its medium and the phenomenon distribution, and adaptively position its nodes to maximize phenomenon sensitivity. Mobility also enables the network configuration to evolve with the medium characteristics and phenomenon distribution over time.

However, in many deployments, the use of mobile robotic nodes instead of static ones is not feasible, because robotic mobility has a high resource and energy overhead. It requires significant terrain sensing, navigational and management support to function autonomously. When the sensors are attached to a wired energy or communication infrastructure, as is the case with high bandwidth sensors, such motion is not applicable. We resolve the above conflict by utilizing a specialized form of mobility, known as *motility*, which relies on reduced complexity motion primitives to reconfigure the network in response to its environment. Motility has low hardware and energy overheads enabling it to be easily incorporated into embedded systems.

It is intuitive to expect that motility will enhance coverage, but it is not immediately obvious when such an approach is better than other alternatives, such as using a higher density of static sensors or long range mobility. The design choice depends on an understanding of the expected gains due to motility, especially when a network of such motile nodes is considered as a whole. Trade-offs based on time required for motion also need to be considered. We show that despite its simplicity, motility can significantly improve the quality of sensing, within reasonable delay constraints on motion. We evaluate this improvement for realistic sensors and practical deployment scenarios. We also contrast this approach against the alternatives, and help recognize the space where motility may be preferred.

2. A SYSTEM MODEL TO SUPPORT SENSOR NETWORK RECONFIGURATION USING MOBILITY

Mobility has many advantages for sensor networks, such as in communications [Grossglauser and Tse 2002; Shah et al. 2003; Goldenberg et al. 2004; Somasundara et al. 2004], security [Capkun et al. 2003], localization [Scott and Hazas 2003], energy transfer [Rahimi et al. 2003], resource replenishment [LaMarca et al. 2002] and coverage repair [Merrill et al. 2002]. We are considering the coverage advantage alone. Any form of mobility helps improve sensing performance in three ways:

(1) Increased Sensor Range: Mobile sensors can cover a larger volume, depending on the acceptable patrolling delay.
(2) Medium Adaptivity: Mobile network components can reposition themselves to overcome obstacles and other anisotropies in the medium.
(3) Phenomenon Adaptivity: Sensors can move in order to sample the phenomenon where highest resolution coverage is required.

In order to understand the benefits of motility, we must consider its use in the above three contexts. As an example, consider a network of cameras deployed to capture all the faces of people present in a shopping mall. One alternative is to deploy a sufficient number of cameras which cover the entire mall volume at the resolution required to recognize a face. A second alternative is to use motile cameras, such as those with pan, tilt and zoom capabilities. Here, a much smaller number of cameras can cover the entire region. Motile cameras can adapt their poses to provide higher resolution in specific regions where more people are present, such as a food court at meal times or a box office before show times. If the cameras know the locations of obstacles in the medium, they can further adapt their poses to cover larger regions by orienting away from an obstacle. Our aim is to determine which of the two alternatives is superior and the relevant trade-offs is using each.

We assume the model shown in Figure 1 for a motile sensor network, which is capable of exploiting the three mobility advantages. Multiple sensors move in a coordinated manner, represented by block III, to collaboratively cover the maximum

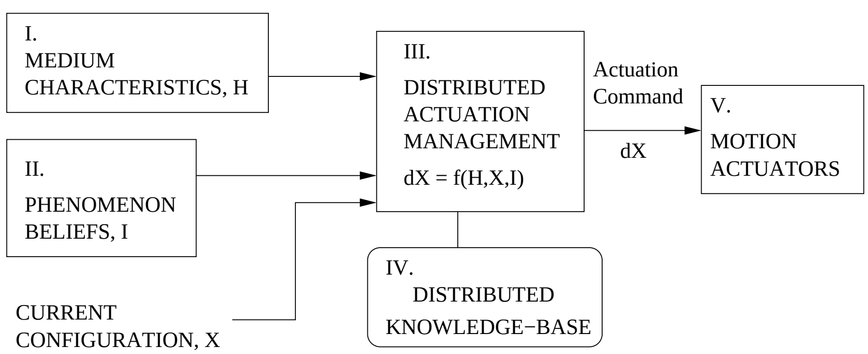

Fig. 1. A framework for managing actuation for sensing uncertainty control.

possible area. The system may have means to learn the locations of the obstacles, block I, in which case this knowledge is utilized in determining the motion strategy. Such methods may be based on the use of range sensors [Harle and Hopper 2003] or in the case of cameras, the use of stereo-vision. Additionally, the system may have means to discover events or phenomena of interest, shown as block II, in which case this information affects the motion strategy. This may be realized using a high density deployment of low power sensors (Eg. Passive Infra-Red sensors) for detecting human motion, or using the cameras themselves at low resolution to detect an event at a large distance, before zooming in to cover the event at the resolution required by the application. The behavior of the medium and the phenomenon may be learnt over time to build up statistical models in a distributed manner (block IV).

The planned motion strategy is limited by the capabilities of the available motion actuators (block V). While each of the blocks shown in Figure 1 presents interesting research challenges, the focus of this paper is restricted to block V, objective being to determine the potential of limited motile capabilities, since these are amenable to use in embedded sensor networks.

2.1 Delay Trade-off

An important point to note when sensors move to increase coverage is that the time spent in moving introduces a delay in sensing. This is not a great concern when motility is used to gradually approach the desired network configuration in response to static obstacles and a static phenomenon distribution, at deployment, with only infrequent updates. However, this delay becomes important when motion is used for adapting the network to phenomenon and medium dynamics in real time. The extent to which the actuation delay is acceptable, depends on the Nyquist frequency of the process. For a continuously sampled field, the Nyquist frequency is well defined and can be determined from the Fourier transform of the signal. Thus, as long as the motion delay is within the Nyquist interval of the sensed phenomenon, the delay is acceptable. For a sensor which is monitoring events, the event detection time and the motion delay should be less than the minimum persistence time of the event when it occurs, so that the event is captured before it expires. For changing medium anisotropies, the delay should be lower than the time-scale at which the medium changes. This delay will be an important consideration in determining the advantage in coverage due to motion.

3. SINGLE SENSOR MOTILITY

3.1 Actuating Node Appendages

Actuating node appendages are a form of mobility where a small appendage on the sensor node is re-oriented without physically relocating the node itself. It is amenable to sensor networks because:

(1) Energy requirements are low, since only the sensor transducer has to be moved while the bulkier parts such as the motors, the battery, and the processor board, can remain stationary.
(2) Navigation support required is minimal. Since the actuation does not depend on unreliable or arbitrary terrain characteristics, no extra sensors are needed to obtain terrain feedback.
(3) The motion is self contained and infrastructural support, such as localization beacons or trajectory markings, is not required.
(4) Since the node itself does not move, such actuation is also feasible in tethered sensor nodes, such as power intensive or high bandwidth sensors.

Despite their simplicity, actuated node appendages can significantly extend the sensing coverage. For instance, the area covered by a camera can be increased by panning and tilting the camera head. As a representative example, we analyze the increase in coverage for a commercially available pan-zoom-tilt camera (Sony SNC-RZ30N, http://bssc.sel.sony.com). Its actuation capabilities are summarized in Table I.

Actuation	**Range**	**Speed**
Pan	+170 to -170 degrees	170°/second
Tilt	+90 to -25 degrees	76°/second
Zoom	25X (Horizontal field of view angle changes from 45° to 2°)	3 seconds/(1X to 25X)

Table I. Example actuation specifications

Figure 2-(a) shows a simple model for the volume covered by the camera at its widest zoom, and without any pan and tilt. This volume, V_s, is approximated as a pyramid with a rectangular base of area $Ll \times Hl$ and height R. A small region near the vertex of the pyramid is not actually covered as the camera has a minimum focal distance and objects closer than this distance cannot be captured. For the above camera, at the widest zoom, this distance is 0.03m, which turns out to be negligible compared to the height of the pyramid, R, calculated below and this effect is thus ignored. Based on the minimum resolution desired for imaging, we can determine the range R which is satisfactorily covered by the camera. Let the image size be $L \times H$ pixels, taking L to be the horizontal dimension. Suppose the minimum spatial resolution needed for the application is such that a length l in space must cover at least one pixel. The image area of $L \times H$ pixels will then represent LHl^2 area in space. The horizontal field of view for the camera, θ_{fov}, can then be used to calculate R for the required l value:

$$tan(\theta_{fov}/2) = \frac{lL/2}{R} \tag{1}$$

The value of l depends on phenomenon size and the minimum image resolution needed by detection algorithms. Suppose $l = 1$cm is used. For the sample camera, $\theta_{fov} = 45°$, $L = 640$ pixels and $H = 480$ pixels, which yields $R = 7.27$m. This value is within the focal range of the camera. The focal range may limit R and hence the usable range of l. We also assume that the illumination is sufficient for the zoom range available on the camera. Thus, V_s is given by:

$$V_s = \frac{LHRl^2}{3} \tag{2}$$

Now consider the volume covered with motility. When motorized zoom with a zoom factor of z is used, the range covered becomes zR. The additional volume covered is a cuboid of depth $(z-1)R$ and area $Ll \times Hl$. The total volume covered with zoom, V_z, becomes:

$$V_z = V_s + LH(z-1)Rl^2 \tag{3}$$

and corresponding gain factor due to zoom motility, g_z, is $g_z = V_z/V_s = 3z - 2$.

Let the tilt angles with respect to the plane of pan rotation be denoted θ^-_{tilt} in the downward direction and θ^+_{tilt} upwards, with their total being θ_{tilt}. The volume covered when the tilt capability is used, V_t, can be viewed as the volume of a sector

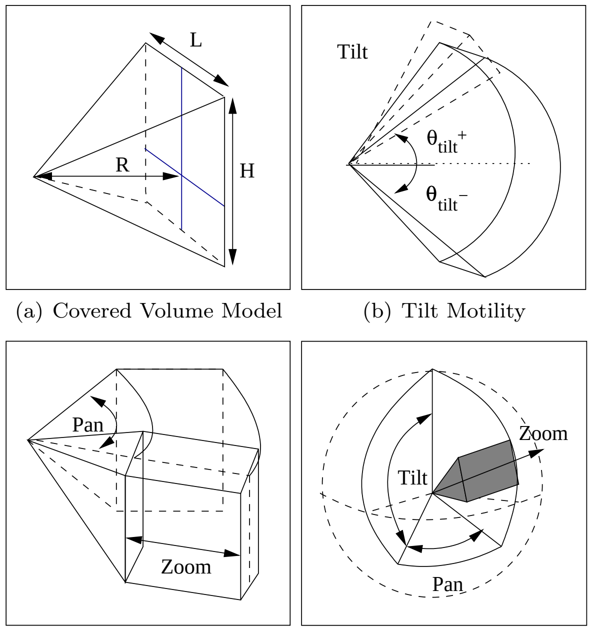

(a) Covered Volume Model (b) Tilt Motility

(c) Pan and Zoom Motility (d) PZT Motility

Fig. 2. Actuation in a PZT camera (a) volume covered without actuation can be modeled as pyramid, (b) tilt capability increases the effective volume covered, (c) combined pan and zoom capabilities further increase the volume covered, and (d) volume covered when pan, tilt, and zoom are combined: this can be viewed as a portion of a sphere swept out using pan and tilt, where the thickness of the sphere depends on the zoom range.

of a disk of radius R, thickness Ll, and sector angle θ_{tilt} minus the volume of two cone sectors of radius R, height $Ll/2$ and sector angle θ_{tilt}. This volume is shown in Figure 2-(b). A half of the pyramid extends beyond the tilt range on each side, and these additional edge volume segments extending beyond the sector total up to V_s. The total volume covered with tilt motility, V_t, is thus:

$$V_t = V_s + \frac{\theta_{tilt}}{2\pi}\left[\pi R^2 Ll - 2\left(\frac{1}{3}\pi R^2 \frac{Ll}{2}\right)\right] \tag{4}$$

The calculation of the volume covered using pan motility is very similar, and is omitted for brevity. Figure 2-(c),(d) show the volumes covered when two or more actuation modes are combined. Evaluating these volumes yields the overall gain due to pan, zoom and tilt combined, g_{pzt}, to be:

$$\begin{aligned} g_{pzt} &= (3z-2) + \frac{\theta_{tilt} R}{2Hl}(3z^2+1) + \\ &\quad \frac{\theta_{pan} z^3 R^2 (1-\cos\theta_{tilt})}{LHl^2} + \left[\frac{cos\theta^{-}_{tilt} + cos\theta^{+}_{tilt}}{2L}\right]\frac{\theta_{pan} R}{2l}(3z^2+1) \end{aligned} \tag{5}$$

This formula can also be used to calculate the gain due to motility primitives separately; for instance, substituting $z = 1$ and $\theta_{pan} = 0$ yields the formula for gain due to tilt alone. Note that g_{pzt} is independent of l as the quantity R/l in each term can substituted from equation (1) in terms of camera parameters θ_{fov} and L.

Substituting the values of L, H and other parameters from Table I in (5), we get the coverage improvements listed in Table II.

Actuation	**Gain**
Pan only	7.74
Tilt only	4.04
Zoom only	73
Pan and Tilt	27.71
Pan and Zoom	6361
Tilt and Zoom	2908
Pan, Tilt, and Zoom	226940

Table II. Coverage Improvement with PZT actuation

Clearly, actuation yields a significant advantage in coverage, especially when two or more simple forms of actuation are combined together. For example, the last row in the table shows that the pan, zoom and tilt actuation can reduce the number of cameras required by over five orders of magnitude, in unobstructed environments.

Let us also consider the actuation delay, T, for the sensor specified in table I. For event detection, suppose the camera is given a command to actuate to a position (p, t, z) where the three variables represent pan, tilt and zoom coordinates respectively, as soon as the event is detected (such as using PIR). The worst case delay is the maximum of the times to move the entire pan range ($2s$), the entire tilt range ($1.5s$), and change the zoom from one extreme to other ($3s$). All these may happen in parallel, yielding $T = 3s$. Thus the coverage gains listed in Table II are available only when this delay is acceptable. As an example consider a possible event detection scenario: individuals passing through an entry foyer are to be photographed each time a motion event is triggered and the coordinates of the detected motion location are delivered to the actuated camera. A pedestrian walking at $1m/s$ takes $6.4s$ to cross the horizontal field of view of the example camera, at resolution corresponding to $l = 1cm$. This time is sufficient for the camera to actuate and capture an image for our example. On the other hand, this delay may not be acceptable for capturing traffic signal violating drivers at a road intersection, and a multiple static cameras may be required to cover the view in each direction at the intersection.

3.2 Linear, Constrained Motion

Linear constrained motion refers to small relocation of a sensor node along a single dimension. Intuitively, such actuation can enable a sensor, whose coverage is blocked by an obstacle, to move to the edge of the obstacle and hence look around it. This form of actuation is expected to provide significant improvement in the presence of small obstacles such as trees in an outdoor environment and everyday objects or human beings in indoor environments. It is not expected to overcome the limitations in coverage due to large boundaries such as walls. This type of mobility is much easier to achieve than long range free terrain mobility because:

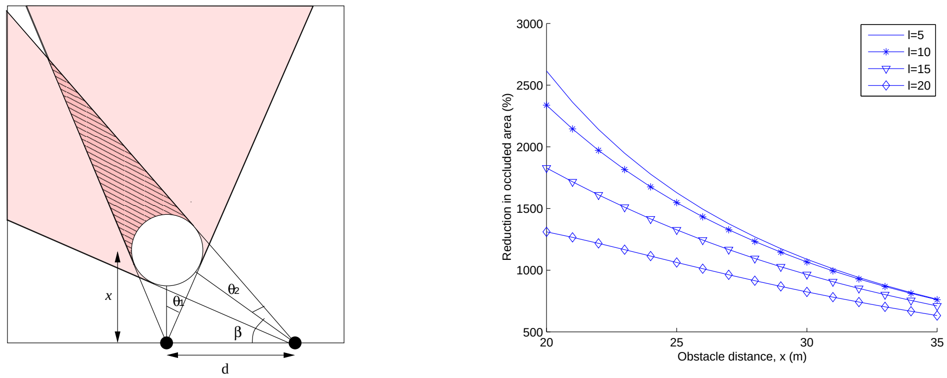

(a) Coverage obstacle (b) Coverage gain

Fig. 3. Coverage gain due to small linear motion (a) the shaded areas show the occluded regions from two static positions of the sensor while the hashed area shows the occluded region when the sensor can move between those two positions, and (b) improvement in coverage due to one-dimensional motion

(1) The motion need not occur on complex physical terrain but may use a pre-designed path, such as a small track attached to the sensor node or on a cable suspended to provide a low energy pathway for the node [Kaiser et al. 2003]. Energy required to move the node can be further reduced when the gravitational work done in moving is counterbalanced using spring-action or counterweights.
(2) Navigational support required is minimal since position along the track can be easily controlled.

A detailed simulation and experimental study of the advantages of such actuation was presented in [Kansal et al. 2004]. We only summarize the results here. Again assume a line of sight sensor. For simplicity of exposition, we restrict the discussion to coverage in a two dimensional plane. Assume that the field of view is 180°. Figure 3-(a) shows a sensor along the bottom edge of a square cell and an obstacle of diameter l at a distance x from the sensor. The sensor can move a distance of 0 to d. The advantage due to actuation is defined as the ratio of area occluded when the sensor is at its initial position to that of the area occluded when the sensor can move to any position along its allowed one dimensional trajectory. Figure 3-(b) shows the gain factor for some values of x and l (with d fixed at $2l$). Significant improvement is apparent.

4. NETWORKS OF MOTILE SENSORS

4.1 Collaborative Motility Advantage

While for a single node the aggregate coverage using the multiple perspectives of an actuated appendage increases, for a network of multiple such devices even the instantaneous coverage using one appropriately chosen sensor perspective at each node, is improved. This occurs due to two factors.

First, a random deployment may cause two sensors to be covering overlapping areas while other areas remain uncovered. Actuation allows efficient orientation of the sensors to cover non-overlapping areas.

Second, regions blocked by obstacles for one sensor may be viewable by another sensor, if that sensor chooses the appropriate position or orientation. Among the multiple possible network configurations, the positions of the nodes can be collaboratively adjusted to minimize occlusions. Coordination is required to ensure that among the multiple possible poses for each sensor, the sensors choose those that cover regions which cannot be efficiently covered by other sensors.

As an illustration, figure 4-(a) shows a random placement of 20 sensors in a $25m \times 25m$ area, each oriented randomly, and in the presence of obstacles. The circles represent obstacles in the sensing medium and the sensor locations are marked with small diamonds. The sectors of circles indicate covered regions, modeling a 45° field

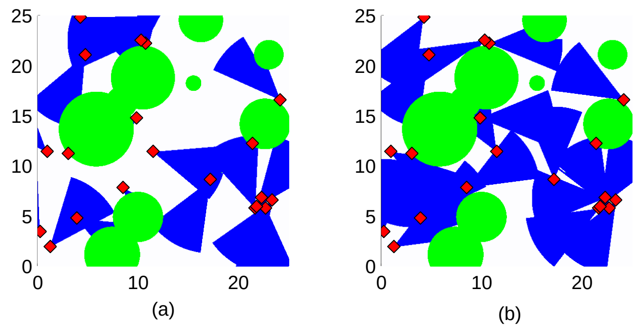

Fig. 4. Instantaneous coverage increase: (a) a sample topology showing randomly oriented sensors, and (b) sensor orientations changed using pan capability.

of view and a range, $R = 7.3m$, as calculated in section 3.1 for the sample sensor. If these sensors have pan capability, their orientations can be changed. Figure 4-(b) shows a changed network, where the number and location of nodes has remained the same, only the orientations have been updated using pan motility. Figure 5 shows this gain with varying node density. A distributed algorithm to update the sensor orientations is considered in Section 4.2.

Similarly, for linearly actuated sensors, an experimental evaluation was carried out in [Kansal et al. 2004] using real sensors (cameras) and model cylindrical obstacles to study the motility advantage in a prototype system. The effect of linear motion on detection probability was studied using image processing based detection algorithms operating on sensor data. The obstacle placement modeled tree locations of a sample forest [Windriver Forest 2003]. Coverage was measured by placing a target (a small cylinder) at evenly spaced locations in the intended coverage region. The target could be occluded by the obstacles at some of these positions depending on obstacle placement. The image processing algorithm relied on a minimum number of target pixels being in view. The test-bed allows experimenting under practical constraints such as sensor noise, constrained angle of view and image processing limitations. Figure 6-(a) shows a sample image from our lab test-bed with the target marked out. Figure 6-(b) shows the coverage quality for mobile

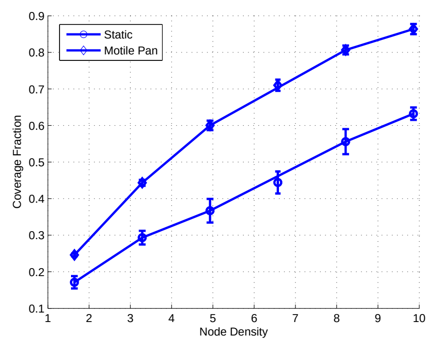

Fig. 5. Coverage advantage when pan actuation is used for initial reconfiguration only, at multiple densities. Error bars show standard deviation across 10 random topologies. (Node density is measured as the number of sensors in the area covered by a single sensor.)

and static cases as the number of cameras is varied. Note that the graph uses a logarithmic scale and the uncovered region, or the probability of mis-detection, is significantly lower with actuated sensors.

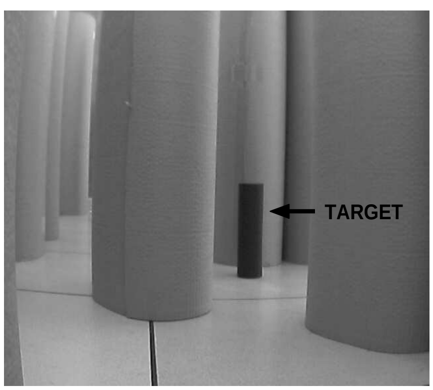

(a) Laboratory testbed

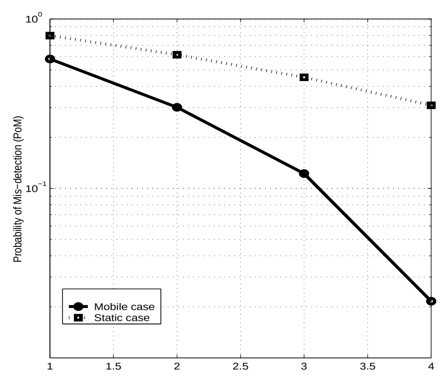

(b) Coverage gain

Fig. 6. Experiments with real sensors: (a) sample image showing field of view of the image sensor, with one target location marked out and (b) probability of mis-detection when the sensors are static and when allowed to move a small distance equal to twice the obstacle diameter

The above results show that limited motion can also be useful in helping reduce sensing uncertainty. This is useful for sensor networks because such simple motion is easier to incorporate in a wide variety of sensor network applications, compared to other alternatives.

4.2 Node Density and Delay Trade-offs

There are two obvious alternatives to the use of motility. One is to use a much higher density of static nodes and the other is to use an even lower density of mobile nodes with long range motion capabilities.

To decouple the delay trade-off initially, let us first assume that the persistence time or the Nyquist interval of the sensed phenomenon is $T = 3s$, which allows the pan and linear motion motility to be used fully. For nodes with unconstrained mobility, while the range of motion is potentially infinite, only a finite motion is possible within the tolerable delay. We assume a velocity $v = 100cm/s$ for such mobile nodes, as navigation and traction support for such speeds has already been demonstrated in experimental platforms. We also allow the mobile nodes to be able to turn with maximum agility at the same speed, which implies that an unconstrained mobile node may move and orient to any position within the distance vT. The exact relation between the numbers of static and mobile sensors required is dependent on the medium and the nature of deployment, and hence we generate random topologies with multiple obstacle and sensor configurations. Assuming a random initial deployment over a 20×20 area and simulating ten random topologies, we evaluate the coverage for each situation- static nodes, nodes with pan only motility, nodes with one dimensional limited linear motility and nodes with full mobility. Note that only a single calculation is performed for the fully mobile nodes, as the initial random configuration is not relevant for this case; these nodes may re-position themselves to the best initial configuration from which the allowable mobility range of vT then allows them to cover the maximum area.

Figure 7-(a) shows these results. Clearly, lower densities are required with actuation for the same coverage quality. For instance, at 90% coverage, the static density is 7 times higher than that with pan actuation.

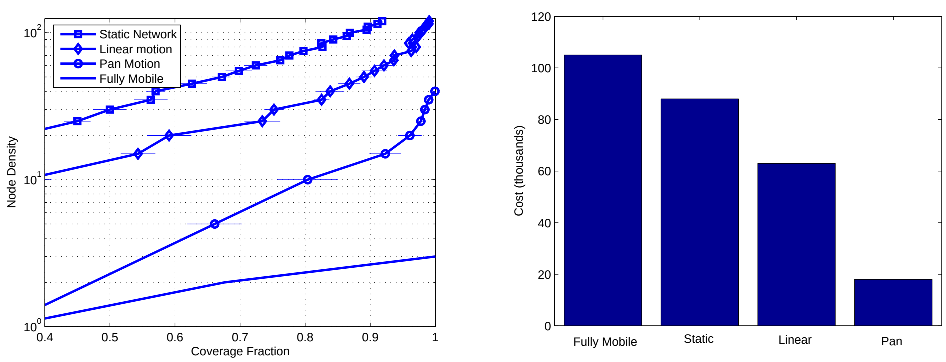

(a) Density Comparison

(b) Cost Comparison

Fig. 7. Alternative deployments: (a) Coverage fraction with and without actuation at different deployment densities. Node density is the number of nodes in a $20m \times 20m$ square region, where each sensor covers a sector of $45°$ with radius $Rs = 7.3m$. The error bars show the standard deviation across 10 random topologies simulated. (b) Converting the node density to total network cost, at 90% coverage point.

Since motile nodes, mobile nodes and static nodes have different costs, it may be more relevant to consider the overall system cost rather than density. The cost needs be evaluated for the particular sensors of interest for a system, and the costs may change over time. However, the factors responsible for the costs can be considered here. The cost of a static node includes the node production cost, c_n the installation cost (providing it power and network connectivity), c_i, maintenance cost, c_m, and the back-end data processing cost, c_p, to process the partially processed data received from that node. The total system cost is $N(c_n + c_i + c_m + c_p)$. It is reasonable to assume that only c_n varies across the three types of nodes, while the other costs are similar. Let us make three further assumptions which make the comparison harsher on the motile nodes. First, we do not account for $(c_i + c_m + c_p)$. Installation costs as high as 75% of the system cost have been reported [Farrar 2001]. Only c_n is considered, which means that the savings in $N(c_i+c_m+c_p)$ due to a lower N for the motile nodes compared to static ones are not being considered. Second, the off-the-shelf fully mobile nodes do not have built-in reliable navigation and localization subsystems, rather they have a traction platform with appropriate interfaces for adding deployment specific navigational support. We ignore the costs of these additions, thus loosing the significant advantage that motile nodes have in that they do not need such extra support. Third, the cost of motile node considered is with three motile capabilities- pan, tilt, and zoom even though only the pan capability is considered in the above density evaluation. Obviously, the cost of a pan-only camera will be lower than one with all three motility primitives. Figure 7-(b) shows the cost trade-off, and even with the harsher comparison, the motile nodes offer a significant advantage. The costs considered are USD 800 for a static network camera [SNC-CS3N 2004], 1300 for a pan-tilt-zoom network camera (same as chosen in table I), and 35000 for a fully mobile node [Packbot 2004]. For linear actuation, we assumed the cost to be 1200, assuming the static node cost plus an additional track, motor and motor controller.

Let us now consider the effect of actuation delay. Suppose the the tolerable delay is lower than the time taken to realize the full range of motility. An important observation here is that even in this case the motility range can help adapt the network configuration in two ways. First, the the full range can still be utilized for orienting away from static obstacles. Second, the full range should be exploited to choose an initial configuration for the network which allows maximum coverage within the motility range achievable for the tolerable actuation delay. For instance, for a pan sensor, the initial static position can be chosen at any point within the pan range. The position should however be chosen such that the regions accessible from it within the pan-delay are those that yield maximum coverage in view of the positions and constraints on other sensors.

While the focus of this paper is not to provide the distributed motion control algorithms (as shown in block III in Figure 1), we do mention a distributed algorithm below, which could be used to determine the initial sensor orientations. This algorithm is provided only to show feasibility and is used for calculating the achieved coverage within varying tolerable delays. Assume that at each sensor s, the set of coverage neighbors,$N_C(s)$, defined as the set of those sensors which are located within $2R_s$ of s, is known, and the locations (but not necessarily the pan orientations) of these nodes are known.

Distributed Coordination Algorithm

(1) **Contention:** Wait for a random back-off period, and send Request To Actuate (RTA) to each node in $N_C(s)$. Wait for $T_{timeout}$. If no CONFLICT message is received, begin actuating as per step 2. Each node which receives an RTA packet, refrains from sending a RTA until it receives an Actuation Finished (AF) packet from each node which had sent it an RTA. This ensures that among nodes with potentially overlapping coverage areas, only one node actuates at any instant.

(2) **Actuation:** If an RTA is received anytime during this phase, send a CONFLICT message to the sender of the RTA. That sender will then retry the RTA after random back-off (Step 1).

 (a) **Neighbor Orientation Information:** Look up the list of AF packets received up to now. AF packets contain node identity and chosen orientation after actuation. Thus the orientation of all nodes within $N_C(s)$ which have actuated up to now is known, and hence the area covered by them can be calculated.

 (b) **Medium Information:**Look up the medium mapping service (implemented either via a range sensor or other means on the node itself or on other dedicated nodes) to obtain information about obstacles within $2R_s$ from own location.

 (c) **Pose Selection:**

  i. Calculate potential coverage angle, θ, as per

$$\theta = \theta_{fov} + 2\omega T \tag{6}$$

  where θ_{fov} is the field of view without actuation, ω is the angular velocity of pan, and T is the time allowed for pan. A sector of angle θ and radius R_s can potentially be covered, since if the pan position is chosen at its center, any point within this sector can be accessed within T, by moving ωT in either direction.

  ii. At any pan position, the coverage area, C_A, is defined as the amount of area covered within the potential sector which is not occluded by an obstacle, and is not already covered by another sensor which has already actuated. Choose a pan position which maximizes C_A.

 (d) Transmit packet AF={identity, chosen orientation}. Other nodes may now contend to actuate. After a node receives an AF packet from every node from which it had received an RTA, it goes to Step 1 unless the termination condition (Step 3) is met.

(3) **Termination:** After a node has got a chance to actuate N_{iter} times, it stops executing this algorithm. N_{iter} determines the quality of optimization, as the node configuration is improved each time a node actuates; however, a larger N_{iter} requires more time to execute.

The algorithm above is a greedy search for the optimal orientations for the sensors. It may not always reach the global optima but is amenable to distributed implementation. Also, it is computationally less expensive than a global search since each sensor only searches for a locally optimal orientation.

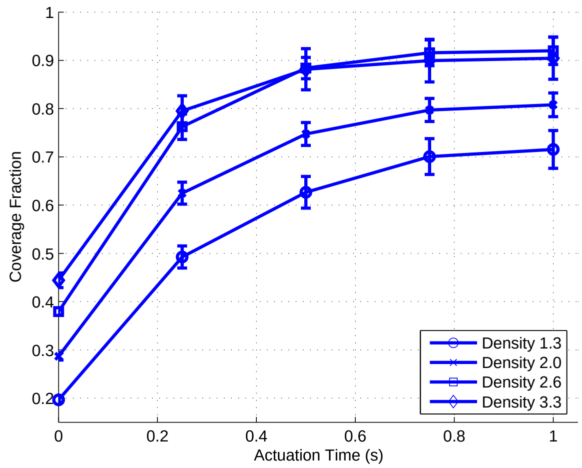

Fig. 8. Coverage achieved with T varying from 0 to 1 second. Different curves show different deployment densities. Density is measured as the number of nodes present in the coverage sector of one static sensor. Error bars show standard deviation across 10 random topologies.

Assuming that the network configuration is optimized using the above algorithm, we evaluate the coverage possible for pan motility, within a tolerable sensing delay T, and vary T from zero to the maximum time needed to realize the full pan range. The results appear in Figure 8. The parameter N_{iter} in step 3 was set to 1 in our simulations.

As expected, the coverage advantage improves with longer actuation delay and saturates as the delay reaches the maximum time required to actuate. However, even at smaller actuation times, there is a significant coverage advantage. This is addition to the advantage obtained due to choosing a better initial configuration with respect to the medium obstacles and overlap with other sensors.

Some other motility trade-offs may also be noted. In case a wireless channel is used, a lower density allows higher data-rates per node. Also, a lower density installation has lower interference with the existing environment, which may be an important consideration in pre-existing structures or in the observation of sensitive biological environments. On the other hand, certain scenarios may not allow even limited motility, such as when the sensor is embedded within concrete. Thus, motility is expected to be beneficial in several environments, though not all.

5. CONCLUSIONS

Sensor networks have traditionally been designed with the assumption of high density static nodes [Intanagonwiwat et al. 2002; Min and Chandrakasan 2002]. Alternatively, the use of robotic sensor nodes has been considered for certain problem domains [Sibley et al. 2002; Poduri and Sukhatme 2004; Wang et al. 2003; Butler and Rus 2003]. Our investigations reveal that these two design paradigms are only two ends of the spectrum and using motion in a limited way can in fact help exploit the advantages of both these paradigms. The objective is to obtain high fidelity data with minimal sensing resources. Motility can help achieved this objective by aligning the available sensing resources with the sensing requirements in the deployment scenario. We studied the benefits of several forms of such motility, and observed significant advantages for coverage, reaching multiple orders of magnitude in certain examples. An example algorithm to coordinate the use of motility capabilities across nodes in a distributed network was also presented.

Motility also has the advantage that it can be realized with a low actuation delay compared to long range mobility. The system costs were seen to be significantly lower for motile nodes than static or mobile nodes. Combined with the savings in installation and maintenance overheads, we expect motility to be applicable in many situations, especially for high bandwidth and high power sensors.

We also noted that the use of motility can be made more effective when support for event detection is available. This allows moving the sensor only when an event is detected rather than having it patrol its coverage region continuously. Further, information about obstacles can be used to improve the network configuration when the nodes are motile. Such information, if available in a static network will only help determine which nodes are affected by obstacles, and while that does allow ignoring the data from affected nodes, the sensing resources at those nodes are wasted. We are working on developing prototype systems with support for these services. Future work includes developing more sophisticated coordinated actuation algorithms, and determining the minimum required degree of actuation for meeting expected performance requirements in a given deployment scenario.

REFERENCES

Butler, Z. and Rus, D. 2003. Event-based motion control for mobile-sensor networks. *IEEE Pervasive Computing 2,* 4 (October-December), 34–42.

Capkun, S., Hubaux, J.-P., and Buttyn, L. 2003. Mobility helps security in ad hoc networks. In *ACM Mobihoc*. ACM Press, 46–56.

Farrar, C. 2001. *Lecture notes on Structural Health Monitoring using Statistical Pattern Recognition*. Los Alamos Dynamics, Los Alamos, NM, Chapter Historical overview of structural health monitoring.

Goldenberg, D., Lin, J., Morse, A. S., Rosen, B., and Yang, Y. R. 2004. Towards mobility as a network control primitive. In *ACM MobiHoc*. Tokyo, Japan.

Grossglauser, M. and Tse, D. N. C. 2002. Mobility increases the capacity of ad hoc wireless networks. *IEEE/ACM Transactions on Networking (TON) 10,* 4, 477–486.

Harle, R. K. and Hopper, A. 2003. Building world models by ray tracing within ceiling mounted positioning systems. In *UbiComp 2003*. Seattle, WA, USA, 1–17.

Intanagonwiwat, C., Estrin, D., Govindan, R., and Heidemann, J. 2002. Impact of network density on data aggregation in wireless sensor networks. In *Proceedings of the 22nd International Conference on Distributed Computing Systems*. IEEE, Vienna, Austria, to appear. See UCLA CSD TR-01-750 for an expanded version of this paper.

Kaiser, W., Pottie, G., Srivastava, M., Sukhatme, G., Villasenor, J., and Estrin, D. 2003. Networked infomechanical systems (NIMS) for ambient intelligence. Tech. Rep. 31, UCLA-NSF Center for Embedded Networked Sensing. December.

Kansal, A., Yuen, E., Kaiser, W. J., Pottie, G. J., and Srivastava, M. B. 2004. Sensing uncertainty reduction using low complexity actuation. In *ACM Symposium on Information Processing in Sensor Networks*. ACM Press, 388–395.

LaMarca, A., Brunette, W., Koizumi, D., Lease, M., Sigurdsson, S. B., Sikorski, K., Fox, D., and Borriello, G. 2002. Plantcare: An investigation in practical ubiquitous systems. In *Ubicomp*.

MERRILL, W., GIROD, L., ELSON, J., SOHRABI, K., NEWBERG, F., AND KAISER, W. 2002. Autonomous position location in distributed, embedded, wireless systems. In *IEEE CAS Workshop on Wireless Communications and Networking* (6). Pasadena, CA.

MIN, R. AND CHANDRAKASAN, A. 2002. A framework for energy-scalable communication in high-density wireless networks. In *ISLPED '02: Proceedings of the 2002 international symposium on Low power electronics and design*. ACM Press, 36–41.

Packbot 2004. Packbot, the next step in unmanned tactical mobile robots. http://www.packbot.com.

PODURI, S. AND SUKHATME, G. S. 2004. Constrained coverage for mobile sensor networks. In *IEEE International Conference on Robotics and Automation*. New Orleans, LA, 165–172.

RAHIMI, M., SHAH, H., SUKHATME, G. S., HEIDEMANN, J., AND ESTRIN, D. 2003. Studying the feasibility of energy harvesting in a mobile sensor network. In *IEEE Int'l Conference on Robotics and Automation*.

SCOTT, J. AND HAZAS, M. 2003. User-friendly surveying techniques for location-aware systems. In *Ubicomp*. Seattle, WA, USA, 1–17.

SHAH, R. C., ROY, S., JAIN, S., AND BRUNETTE, W. 2003. Datamules: Modelling a three tiered architecture for sparse sensor networks. In *First IEEE International Workshop on Sensor Network Protocols and Applications (SNPA)*.

SIBLEY, G. T., RAHIMI, M. H., AND SUKHATME, G. S. 2002. Robomote: A tiny mobile robot platform for large-scale ad-hoc sensor networks. In *IEEE International Conference on Robotics and Automation*. Washington, DC, USA.

SNC-CS3N 2004. Sony SNCCS3N. http://bssc.sel.sony.com/Professional/. CS Mount Fixed Network Color Camera.

SOMASUNDARA, A., KANSAL, A., JEA, D., SRIVASTAVA, M., AND ESTRIN, D. 2004. Intelligent fluid infrastructure for embedded networks. In *ACM Mobisys*. Boston, MA.

WANG, G., CAO, G., AND PORTA, T. L. 2003. A bidding protocol for deploying mobile sensors. In *Proceedings of the IEEE International Conference on Network Protocols (ICNP)*. Atlanta, GA, 315–324.

Windriver Forest 2003. Wind river canopy crane research facility. http://depts.washington.edu/wrccrf/12haplot/. 12-ha Permanent Plot Tree Data.

Received September, 2004.