Coverage-Aware Sensor Engagement in Dense Sensor Networks [⋆]

Jun Lu, Lichun Bao and Tatsuya Suda

Bren School of Information and Computer Sciences
University of California, Irvine, CA 92697
`{lujun,lbao,suda}@ics.uci.edu`

Abstract. The critical issue in sensor networks is the search of balance between the limited battery supply and the expected longevity of network operations. Similar goals exist in providing a certain degree of sensing coverage and maintaining a desirable number of sensors to communicate under the energy constraint. We propose a novel sensor network protocol, called Coverage-Aware Sensor Engagement (CASE) for coverage maintenance. Different from others, CASE schedules active/inactive sensing states of a sensor according to the sensor's contribution to the network sensing coverage. The contribution is quantitatively measured by a metric called coverage merit. By utilizing sensors with large coverage merit, CASE reduces the number of the active sensors required to maintain the level of coverage. Simulation results show that CASE considerably improves the energy efficiency and reduces the computation and communication costs to maintain the required coverage degree in a dense sensor network.

1 Introduction

Wireless sensor networks are networks of a large number of small wireless devices, which collaborate to monitor environments and report sensing data via wireless channels. Wireless sensor networks have emerged rapidly in a variety of applications. For instance, thermal sensors are being deployed to monitor temperature in a forest, and to report the temperature information back to data collection nodes for further analysis. In another instance, a large number of seismic sensors are employed to monitor animal activities in a wild field. The seismic sensors, when triggered by the vibrations caused by animal movements, can record the vibration signals and report them to data collection nodes. Information about animal activities, like their moving tracks and velocities, can be acquired through analyzing the collected signals.

Wireless sensors are very limited in their processing, computing and communication capabilities as well as the storage and power supply. The typical Crossbow MICA mote MPR300CB [XBOW1] has a low-speed 4MHz processor equipped with only 128KB

[⋆] This work was supported by the National Science Foundation through grants ANI-0083074 and ANI-9903427, by DARPA through Grant MDA972-99-1-0007, by AFOSR through Grant MURI F49620-00-1-0330, and by grants from the University of California MICRO Program, Hitachi, Hitachi America, Novell, Nippon Telegraph and Telephone Corporation (NTT), NTT Docomo, Fujitsu, and NS-Solutions.

flash, 4KB SRAM and 4KB EEPROM. It has a maximal data rate of 40kbps and a transmission range of about 100 feet, powered by two AA batteries. Therefore, a wireless sensor network is usually deployed with high density. Dense deployment not only helps to improve a sensor network's reliability, but also extends its longevity. In practice, large-scale wireless sensor networks are usually deployed randomly.

Given such a randomly and densely deployed wireless sensor network, it is desirable to have sensors autonomously schedule their duty cycles while satisfying the degree of sensing coverage required by the application. The problem is called coverage maintenance.

Coverage maintenance problem in sensor networks has drawn intense research attention recently. Tian *et al.* [TG1] presented a node-scheduling algorithm to turn off redundant sensors if their sensing areas are already covered by their neighbors. Randomized as well as coordinated sleep algorithms were proposed in [HL1] to maintain network coverage using low duty-cycle sensors. The randomized algorithm enables each sensor to independently sleep under a certain probability. The coordinated sleep algorithm allows each sensor to enter sleep state if its sensing area is fully contained by the union set of its neighbors. A K-coverage maintenance algorithm was proposed in [HT1] so that each location of the sensing area is covered by at least K sensors. A sensor decides whether it is redundant only by checking the coverage state of its sensing perimeter. In [GWL1], the redundancy of the sensing coverage of wireless sensor networks is analyzed, and the relation between the number of neighbors and the coverage redundancy is studied. Abrams *et al.* studied a variant of the NP-hard SET K-COVER problem in [AGP1], partitioning the sensors into K covers such that as many areas are monitored as frequently as possible. Yan *et al.* proposed an adaptable energy-efficient sensing coverage protocol, in which each sensor broadcasts a random time reference point, and decides its duty schedule based on neighbors' time reference points [YHS1].

We propose a new coverage maintenance scheme called Coverage-Aware Sensor Engagement (CASE). CASE is based on a probabilistic sensing model, which is more practical than the disk sensing model assumed by many others. Rather than fixing the sensing range of a sensor as the disk sensing model, the probabilistic sensing model defines the sensing ability as the probability to detect an event happening at a location. To the best of our knowledge, this work is the first to address the K-coverage problem under the probabilistic sensing model. In fact, the disk sensing model is a special case of the probabilistic sensing model, and CASE works for the disk sensing model as well. In CASE, each sensor is initially inactive in sensing, but checks to see whether it is necessary to turn on its sensing unit according to its contribution (which we called coverage merit) to meet the required degree of sensing coverage. Before actually turning on itself, each sensor waits for a back-off period decided by its coverage merit. Sensors with larger coverage merit have shorter back-off period. In this way, sensors turn on themselves (if necessary) in the decreasing order of their coverage merit. By utilizing sensors with large coverage merit, CASE can reduce the active sensor density needed to maintain the required coverage degree.

The rest of this paper is organized as follows. The differences of our work from the others are examined in Section 2. Section 3 describes the assumptions of CASE.

Section 4 specifies CASE in more details. Simulation results are presented in Section 5 for performance evaluations. Section 6 concludes the paper.

2 Prior Works

In the scheduling algorithm proposed by [TG1], every sensor is active at the beginning. A sensor is eligible to turn off if its sensing area is covered by the union of the sponsored sectors by its neighbors. A sensor eligible to turn off broadcasts a `TURNOFF` beacon to inform neighbors. Every other sensor receiving such a beacon re-evaluates its eligibility to turn off. If not eligible, the sensor cancels the timer and stays active. [YHS1] presented an elegant approach to dynamically schedule sensors in order to guarantee a certain degree of coverage. Each sensor generates a random reference time and exchanges the reference time with its neighbors. Each sensor can setup its working schedule by examining the reference time of its neighbors.

There are several major differences between the proposed algorithm and the algorithms proposed in [TG1] and [YHS1]. First, we differentiate sensors according to their coverage merit, which decreases the active sensor density to provide the coverage degree required by the application. Second, while both of the existing schemes assume the disk sensing model, CASE solves the K-coverage problem under the probabilistic sensing model. Treating the disk sensing model as a special case of the probabilistic model, CASE also works for the disk sensing model. Third, unlike [TG1], CASE sets sensor initial state as inactive. Following the scheduling procedure, each sensor tries to turn on to provide the required coverage degree. This feature is favorable for dense deployment in that the communication and computation overhead is reduced due to less sensor state changes.

The scheme proposed by [TG1] does not work for the probabilistic sensing model. In order to compare CASE with the scheme proposed in [TG1], we modified the eligibility rule of [TG1] to accommodate the probabilistic sensing model. We refer the scheme proposed in [TG1] as *the sponsored sector scheduling scheme* or *Tian-Sector* and the scheme with the modified eligibility rule as *the grid point scheduling scheme* or *Tian-Grid*. Section 5.1 explains in details about the differences between Tian-Grid and Tian-Sector. The validation of the working schedule setup algorithm in [YHS1] is tightly coupled with the assumption of the disk sensing model, which prevents it from being ported to the probabilistic sensing model. Therefore, in the simulation evaluation, we only compare CASE with Tian-Grid and Tian-Sector.

3 Assumptions

We assume that sensors are static, and that each sensor knows its location as well as its neighbors'. Such assumptions are conveniently taken by other works [HT1] [TG1] [YHS1] and are supported by the existing research [ACZ1] [BHE1] [BP1] [PCB1]. The location information can be absolute or relative to neighbors. We also assume that sensors can synchronize their timers [DH1] [EGE1].

We assume that the sensing ability model of sensors is available before deployment through calibration process. A sensor detects an event based on its measurement, and the event is detected if the measurement strength is above a preset threshold. Due to the signal attenuation and noise, a sensor's measurement is modeled by a probability density function (PDF), which varies with the type of signals and the propagation channel. In CASE, the sensing ability of a sensor is modeled as the probability of a successful detection of certain events of interests.

Apparently, a sensor's sensing ability is a function of the distance between the sensor and the event [MKQP1]. We use $S_j(P_i)$ to describe sensor j's sensing ability at location P_i. A sensor's sensing range, which is denoted by SR, is defined as the range, beyond which the sensor's sensing ability can be neglected. The disk sensing model is regarded as a special case of the probabilistic sensing model, where a sensor detects an event within the SR with the probability 1 and outside the SR with the probability 0.

We assume that sensors have the same SR, and that the sensor communication range is greater than or equal to $2 \cdot SR$. This is usually true in practice. For example, ultra-sonic sensors have a sensing range of approximately 0.2-6m [ROB1] while the transmission range of MICA motes is about 30 meters [XBOW1]. In the cases that the communication range is less than $2 \cdot SR$, our algorithms can work through multi-hop transmissions.

4 Coverage-Aware Sensor Engagement (CASE)

4.1 K-Coverage

The objective of CASE is to guarantee K-Coverage with the least number of active sensors. Under the disk sensing model, location P_i is K-covered if the location is monitored by at least K sensors. Under the probabilistic sensing model, however, we need to modify the definition of K-Coverage. We say a location P_i is K-covered if the *expectation* of the number of sensors that monitor an event at the location is at least K, or essentially the weighted sum of active sensors is no less than K, as shown in Eq. (1).

$$\sum_j S_j(P_i) \geq K \tag{1}$$

where $S_j(P_i)$ is the probability of sensor j to detect an event at location P_i. Note that coverage degree K can be a real number under the probabilistic sensing model. For example, an application may require the target area to be 1.5-covered, which means the expected number of sensors that detect an event at any location in the area needs to be at least 1.5.

4.2 Coverage Merit

Apparently, when a location P_i is already covered by a group of sensors A, the additional coverage needed to fulfill the K-coverage requirement is

$$C(P_i) = K - \sum_{m \in A} S_m(P_i). \tag{2}$$

If $C(P_i)$ is greater than 0, more sensors are required to provide additional coverage.

In order to see sensor j's coverage merit at location P_i, we compare sensor j's probability of detecting an event at location P_i with $C(P_i)$, the minimum of which is defined as sensor j's coverage merit at the location. That is, sensor j's coverage merit at location P_i is:

$$CM_j(P_i) = \begin{cases} \min(C(P_i), S_j(P_i)), & C(P_i) > 0 \\ 0, & C(P_i) \leq 0 \end{cases} \tag{3}$$

Note that when $C(P_i)$ is less than or equal to 0, P_i is already K-covered, therefore the sensor j's coverage merit at this location is 0. It is easy to see that $CM_j(P_i)$ is a continuous function over the sensing area of sensor j, and is dependent on the active states of its neighbors.

In order to evaluate sensor j's coverage contribution to the sensor network as a whole, the summation of $CM_j(P_i)$ over sensor j's sensing area is computed as:

$$CM_j = \int\int CM_j(P_i)dxdy$$

Since the existence of a sensor only affects the area covered by the sensor, its coverage merit can be calculated by only considering the area within its SR. For computation convenience, the above equation is converted into polar coordinates:

$$CM_j = \int_0^{2\pi}\int_0^{SR} CM_j(P_i)rd\theta dr \tag{4}$$

4.3 Coverage-Aware Sensor Engagement

To provide K-coverage with the minimum number of sensors, CASE applies a greedy strategy by gradually activating sensors in decreasing order of their coverage merit. In contrast, previous schemes schedule sensors regardless of their contribution to meet the required degree of network coverage (*e.g.*, in [TG1], redundant sensors have the same chance to power off based on a random back-off timer). More specifically, CASE runs in two phases as follows:

1. Wakeup phase: the first phase when sensors start in inactive sensing state, and gradually enter the active state according to their coverage merits.
2. Optimization phase: the second phase when sensors optimize the coverage by turning off redundant sensors to meet coverage requirements.

In the wakeup phase, each sensor is inactive in sensing, and computes an initial coverage merit. Note that the inactive/active states are logical states in CASE, *i.e.*, sensors are actually active to execute the CASE algorithm. Because no neighbor is active, the initial coverage merit of a sensor is maximum, and given by Eq. (5).

$$CM_{max} = \int_0^{2\pi}\int_0^{SR} \min(K, S_j(P_i))rd\theta dr \tag{5}$$

Afterward, each sensor sets a back-off timer T to announce its active state. The back-off timer T is determined according to its coverage merit using Eq. (6).

$$T = \xi \cdot (CM_{max} - CM_j) + \epsilon \tag{6}$$

where ξ is a configurable system parameter, and ϵ is a small positive random number. ξ determines the convergence latency of the wakeup phase in CASE. Small value of ξ means fast convergence but may increase the chance of collisions among neighboring sensors. The choice of an appropriate value for ξ is out of the scope of this paper, and will be part of our future work. ϵ is used to prevent the potential collision between two neighboring sensors with the same coverage merit.

According to Eq. (6), sensors with larger coverage merit have shorter back-off period. When a sensor times out, the sensor changes to active state, and broadcasts a `TURNON` message to its neighbors within its transmission range, which is approximated by $2 \cdot SR$. When a sensor receives a `TURNON` message before the timer expires, it recalculates its coverage merit, and adjusts its back-off timer accordingly.

According to Eq. (2), (3) and (4), the sensor's coverage merit is reduced when a new neighbor is turned into active state, thus the back-off timer is always delayed. Once all the locations within the sensing range of a sensor are K-covered, the sensor's coverage merit becomes 0. The sensor cancels its back-off timer and stays inactive.

The wakeup phase ends at $\xi \cdot CM_{max}$. After the wakeup phase, there may be redundant active sensors. This is because that the coverage of the sensors turning on later may overlap with the sensing area of the active sensors. In the optimization phase, we use a similar random back-off algorithm as [TG1] to turn off redundant sensors. Accordingly, each redundant sensor sets a random timer, and re-checks its eligibility to turn off when received `TURNOFF` messages from other sensors. If the sensor realizes that it is not eligible to turn off, it cancels its timer and stays active. Otherwise, it broadcasts a `TURNOFF` message and turns off upon timeout.

In order to simplify the computation, we cover the target area by a virtual square grid (Fig. 1) and sensors only consider the grid points within the SR when calculating coverage merit (this technique is also used by Yan *et. al* [YHS1]). The coverage merit is approximated by the summation of the coverage merit on the grid points within the SR, *i.e.*, P_i in Eq. (4) is a grid point.

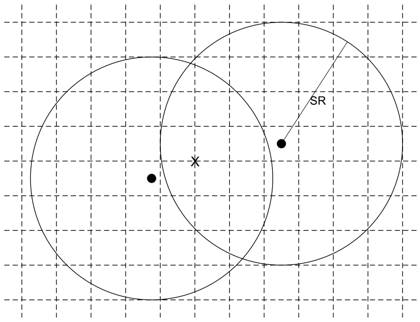

Fig. 1. Grid Point X in the Target Area

5 Simulation Evaluations

5.1 Experiment Setup

We carried out experiments under two sensing models — the probabilistic sensing model and the disk sensing model, both over a square deployment area of $100 \times 100 m^2$.

In Eq. (6), parameters are chosen as $\xi = 0.1$ and $\epsilon = 0.01$. If not otherwise specified, the deployment density is set to 0.08 sensors/m^2, and the network is designed to provide the coverage degree of 1.0.

The probabilistic sensing models depend on the sensor capabilities and environments. Although CASE shall work with any realistic sensing model, for simplicity, we assume a virtual probabilistic sensing model for the sensors, two examples of which are shown below,

$$S_j(P_i) = f(D_{ij}) = \frac{1}{1 + \alpha D_{ij} + \beta D_{ij}^2 + \cdots + \gamma D_{ij}^k}$$

$$S_j(P_i) = f(D_{ij}) = \frac{1}{\chi^{D_{ij}}}$$

where D_{ij} is the distance between sensor j and location P_i; α, β, γ and χ ($\chi > 1$) are system parameters reflecting the physical characteristics of sensor j and deployment environments.

Specifically, we assume the following virtual probabilistic sensing model in the simulations:

$$f(D_{ij}) = \frac{1}{(1 + \alpha D_{ij})^{\beta}} \tag{7}$$

where α is set to 0.1 and β is set to 3 or 4. Assuming that detection probability lower than 4% is negligible, two SRs, *i.e.*,15 and 20 meters, are simulated. For the disk sensing model, the SR is set to 15 meters.

As we explained in Section 2, under the probabilistic model, CASE is compared with the modified Tian-Sector based on virtual grids, which we call Tian-Grid. Like CASE, Tian-Grid checks the expected number of monitoring sensors at each grid point within its sensing range. A sensor is eligible to turn off if the expected number of monitoring sensors of each grid point within its sensing range is at least K. Also, different from Tian-Sector, which only examines the sectors sponsored by neighbors within SR, Tian-Grid considers all the neighbors within $2 \cdot SR$. In the disk sensing model, we compare CASE with both Tian-Grid and Tian-Sector because Tian-Grid natively works for the disk sensing model.

5.2 Result Analysis

The simulation results show the performance of CASE in terms of active sensor density, communication overhead and computation overhead. The communication overhead is computed as the number of beacons sent and received for the `TURNON` messages in the wakeup phase and the `TURNOFF` messages in the optimization phase in CASE. Because the eligibility checking is the most costly computation operation, the computation

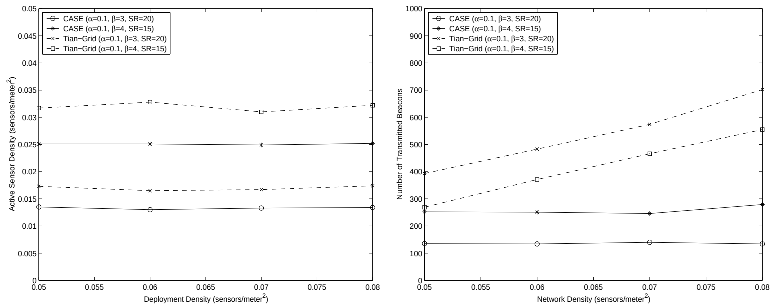

(a) Active Sensor Density (b) Transmitted Beacons

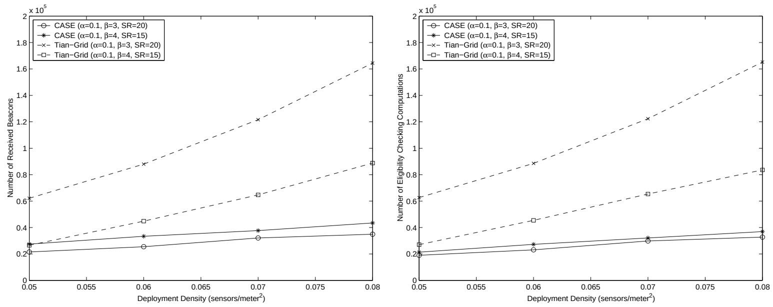

(c) Received Beacons (d) Computation Overhead

Fig. 2. Various Deployment Densities

overhead is calculated as the times of checking the eligibility of a sensor to be in active state, which is determined by coverage merit.

We analyze the results in the probabilistic sensing model and the disk sensing model, separately.

Probabilistic sensing model In section 4.3, we have proposed to compute the coverage merit based on virtual grids. For comparison purposes, we simulate the modified Tian-Sector protocol, which we refer as *Tian-Grid* in the figures, and collect corresponding statistics.

The results under various deployment densities are shown in Fig. 2. Results for different required coverage degrees are shown in Fig. 3.

Fig. 2(a) indicates that both CASE and Tian-Grid provide stable active sensor density. However, CASE results lower active sensor density than Tian-Grid under different deployment densities because CASE activates sensors with large coverage merit, therefore allowing less active sensors in order to provide the same degree of coverage. For instance, when the sensor network has the deployment density of 0.05 sensors/m^2 and sensors have the SR of $20meters$, CASE provides 1.0-coverage with the active sensor density of only 0.0137 sensors/m^2, whereas Tian-Grid requires 0.0175 sensors/m^2.

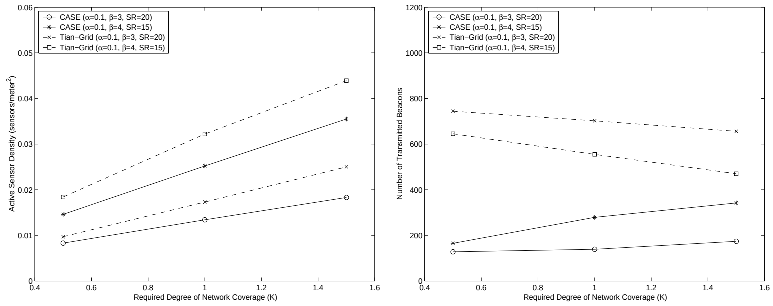

(a) Active Sensor Density (b) Transmitted Beacons

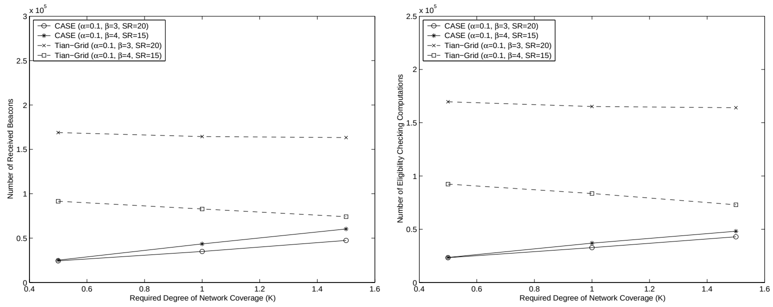

(c) Received Beacons (d) Computation Overhead

Fig. 3. Various Required Coverage Degrees (K)

Fig. 2(b) shows that CASE uses less beacons than Tian-Grid. This is due to the fact that sensors are gradually switched on from inactive state to active state in CASE, whereas Tian-Grid has all sensors initially in active state and turn off redundant sensors, which translates into different amount of beacons transmitted in order to inform state changes. If the network deployment is dense enough, the number of redundant sensors is much larger than the number of active sensors needed to provide the required coverage degree. Thus CASE involves less state changes than Tian-Grid.

Furthermore, we observe that the number of transmitted beacons in CASE changes little along with the increase of deployment density. In contrast, Tian-Grid suffers when the deployment density increases in Fig. 2(b). This is because the active sensor density is almost stable along with the deployment densities in CASE, whereas in Tian-Grid, most beacons are the `TURNOFF` messages sent by redundant sensors. When the deployment density increases, more redundant sensors need to turn off with more beacons.

Similar to Fig. 2(b), Fig. 2(c) shows that CASE has less received beacons than Tian-Grid, and that the number of beacons received in both schemes increases with the deployment density because of the broadcast nature of the wireless channel. However, the increasing rate of received beacons in CASE is less than that in Tian-Grid because the increase of the received beacons in CASE is mainly caused by the increase of sensor density. In Tian-Grid, however, the increase is caused by the increase of both the

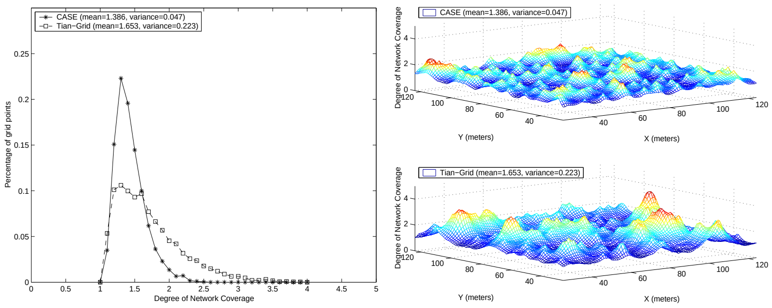

(a) (b)

Fig. 4. Coverage Distribution; $\alpha = 0.1, \beta = 3, SR = 20$

number of the transmitted beacons and sensor density. Because the eligibility checking computations are often triggered by the received beacons, we have similar observation for computation overhead as shown in Fig. 2(d).

In Fig. 3, we show the results based on various coverage degrees requirements. Again, CASE performs better than Tian-Grid under various coverage degree requirements. However, the difference between the two protocols in Fig. 3(b), 3(c) and 3(d) diminishes with the increase of the required coverage degree. This is because the higher the coverage degree, the more sensors are needed active. Because the sensors initially assume inactive in CASE, higher coverage degree means more sensors need to turn on. While in Tian-Grid where sensors are initially active, higher coverage degree means less sensors need to turn off.

To further investigate the performance improvement of CASE, we show a normalized histogram of the number of grid points under different coverage degrees in Fig. 4(a). As we can see, majority of the grid points are covered by a degree from 1.0 to 2.0 in CASE, while the grid-point coverage under Tian-Grid varies from 1.0 to 3.0. From a different observation angle, we plotted the coverage degree of different points in the sensor network as shown in Fig. 4(b), which indicates that CASE provides more even coverage than Tian-Grid does.

Disk sensing model We compare CASE with Tian-Grid and Tian-Sector under the disk sensing model in Fig. 5. In [YZLZ1], a theoretical lower bound on the active sensor density to achieve 1-coverage is provided as $2/\sqrt{27}SR^2$, and is again calculated here in Fig. 5(a) as a baseline for the comparison purposes. Note that although we only present the results of 1-coverage, similar results are observed for other K values (*e.g.*, K = 2).

Fig. 5(a) shows that Tian-Grid achieves the same required coverage degree with less than half of the active sensor density required by Tian-Sector. This is because that Tian-Sector is conservative about the sensor redundancy by only considering the neighbors within SR, and ignoring the coverage provided by sensors in range from SR to $2 \cdot SR$. Thus Tian-Sector results in relative high density of active sensors. Again, CASE performs better than Tian-Grid by reducing 20% of the active sensor density.

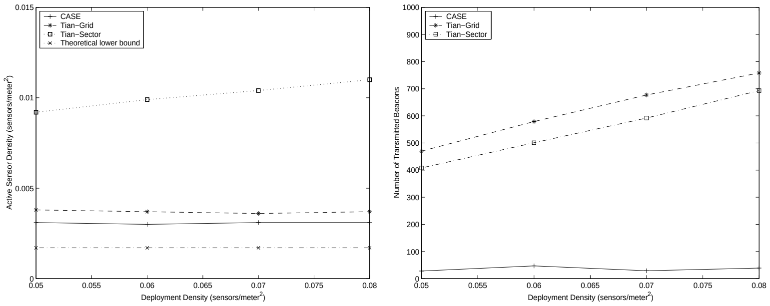

(a) Active Sensor Density (b) Transmitted Beacons

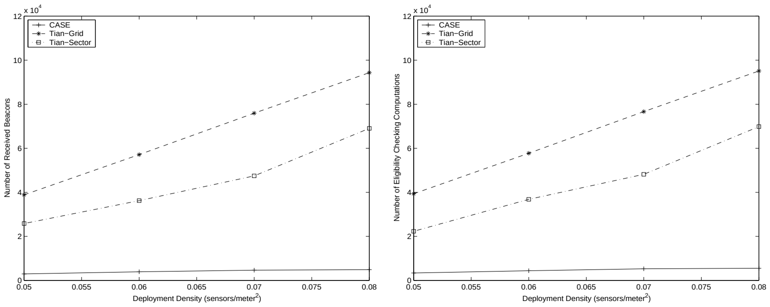

(c) Received Beacons (d) Computation Overhead

Fig. 5. Disk Sensing Model

A larger discrepancy between CASE and the other two protocols are shown in terms of the communication and computation overheads in Fig. 5(b), 5(c) and 5(d).

6 Conclusions

We have proposed a novel coverage maintenance scheme called Coverage-Aware Sensor Engagement (CASE). CASE conserves energy while providing the required coverage degree by allowing sensors to autonomously decide their active/inactive states. Unlike prior works, CASE considers local coverage information of sensors, *i.e.* coverage merit, when scheduling sensors' active/inactive states. Simulation results show that CASE provides the required coverage degree for a dense sensor network with lower active sensor density and less communication and computation costs than existing solutions. Furthermore, CASE is highly scalable to sensor network deployment density due to the low increasing rate of communication and computation costs relative to the increase of deployment density.

References

[ACZ1] Albowicz, J., Chen, A., Zhang, L.: Recursive Position Estimation in Sensor Networks. Proceedings the IEEE International Conference on Network Protocols (ICNP), 2001.

[AGP1] Abrams, Z., Goel, A., Plotkin, S.: Set K-Cover Algorithms for Energy Efficient Monitoring in Wireless Sensor Networks. Proceedings IPSN, 2004.

[BHE1] Bulusu, N., Heidemann, J., Estrin, D.: GPS-less Low Cost Outdoor Localization For Very Small Devices. IEEE Personal Communications, 2000.

[BP1] Bahl, P., Padmanabhan, V.: RADAR: An In-Building RF-based User Location and Tracking System. Proceedings INFOCOM, 2000.

[DH1] Dai, H., Han, R.: TSync: A Lightweight Bidirectional Time Synchronization Service for Wireless Sensor Networks. Mobile Computing and Communications Review, 2004.

[EGE1] Elson, J., Girod, L., Estrin, D.: Fine-Grained Network Time Synchronization using Reference Broadcasts. Proceedings OSDI, 2002.

[GWL1] Gao, Y., Wu, K., Li, F.: Analysis on the Redundancy of Wireless Sensor Networks. Proceedings WSNA, 2003.

[HL1] Hsin, C., Liu, M.: Network Coverage Using Low Duty-Cycled Sensors: Random and Coordinated Sleep Algorithms. Proceedings IPSN, 2004.

[HT1] Huang, C., Tseng, Y.: The Coverage Problem in a Wireless Sensor Network. Proceedings WSNA, 2003.

[MKQP1] Meguerdichian, S., Koushanfar, F., Qu, G., Potkonjak, M.: Exposure In Wireless Ad-Hoc Sensor Networks. Proceedings MOBICOM, 2001.

[PCB1] Priyantha, N.B., Chakraborty, A., Balakrishnan, H.: The Cricket Location-Support System. Proceedings MOBICOM, 2000.

[ROB1] Robosoft Advanced Robotics Solutions: http://www.robosoft.fr/SHEET/02Local/1001LAUN/LAUN.html. Last visited on 09/07/2005.

[TG1] Tian, D., Georganas, N.D.: A Coverage-Preserving Node Scheduling Scheme for Large Wireless Sensor Networks. Proceedings WSNA, 2002.

[XBOW1] XBOW Inc.: http://www.xbow.com/Support/Support_pdf_files/MPR-MIB_Series_Users_Manual.pdf. MPR/MIB User's Manual, Last visited on 06/30/2005.

[YHS1] Yan, T., He, T., Stankovic, J.A.: Differentiated Surveillance for Sensor Networks. Proceedings SenSys, 2003.

[YZLZ1] Ye, F., Zhong, G., Lu, S., Zhang, L.: Energy Efficient Robust Sensing Coverage in Large Sensor Networks. UCLA Technical Report, 2002.