Energy-Quality Tradeoffs for Target Tracking in Wireless Sensor Networks

Sundeep Pattem[1], Sameera Poduri[2], and Bhaskar Krishnamachari[1]

[1] Department of Electrical Engineering,
[2] Department of Computer Science,
University of Southern California,
Los Angeles, CA 90036, USA
{pattem, spoduri, bkrishna}@usc.edu, http://ceng.usc.edu/~bkrishna/

Abstract. We study the tradeoffs involved in the energy-efficient localization and tracking of mobile targets by a wireless sensor network. Our work focuses on building a framework for evaluating the fundamental performance of tracking strategies in which only a small portion of the network is activated at any point in time. We first compare naive network operation with random activation and selective activation. In these strategies the gains in energy-savings come at the expense of increased uncertainty in the location of the target, resulting in reduced quality of tracking. We show that selective activation with a good prediction algorithm is a dominating strategy that can yield orders-of-magnitude energy savings with negligible difference in tracking quality. We then consider duty-cycled activation and show that it offers a flexible and dynamic tradeoff between energy expenditure and tracking error when used in conjunction with selective activation.

1 Introduction

There is an emerging trend towards the use of sophisticated wireless networks of unattended sensor devices for intelligence gathering and environmental monitoring [1] -[6]. One canonical application of sensor networks that has received considerable attention in the literature is the tracking of a mobile target (point source) by the network.

In a tracking scenario, information obtained from nodes far away from the region of activity is of little or no use. For a typical sensor network with a large number of nodes, a major portion of these falls in the above category. In addition, if the nodes are densely deployed, information obtained from some sensors close to the region of activity might be redundant. An obvious way to save energy is to switch on only a subset of the sensor nodes. We discuss in this paper various possible activation strategies: (1) naive activation, (2) randomized activation (3) selective activation based on trajectory prediction and (4) duty-cycled activation.

In these sensor activation strategies, energy savings come at the expense of a reduction in the quality of tracking. In other words, relying on the information provided by a small subset of the sensor nodes results in an increased uncertainty

F. Zhao and L. Guibas (Eds.): IPSN 2003, LNCS 2634, pp. 32–46, 2003.

in the sensed location of the mobile. In this paper we study the energy-quality tradeoffs involved by building a model to quantify both the energy expenditure and the quality of tracking. Also for a particular strategy, we study the impact of the following: a) deployed/activated density of sensors b) their sensing range c) capabilities of activated and un-activated nodes d) the target's mobility model.

Our efforts are not directed *per se* at proposing new techniques for mobile tracking. Rather the focus is on the evaluation and analysis of general strategies which may be incorporated into a real system. We start with a simple model for tracking and substantiate the intuition that it is possible to obtain orders of magnitude savings in energy while keeping the uncertainty within acceptable limits. We also discuss the extensions of the model to relate closely with real life scenarios. The results in this work are a first step in our attempt to understand the fundamental bounds on the the tracking quality that can be obtained under various energy constraints and sensor models.

The rest of the paper is organized as follows. In section 2, we discuss related work from the existing literature, presenting the context for our work. We describe our basic model, assumptions and evaluation metrics for target tracking in section 3. The general tracking strategies that we investigate are detailed in section 4. Section 5 contains the description of our experiments to evaluate the performance of these strategies, and an analysis of the results presented. Finally, we present concluding comments in section 6.

2 Related Work

Target tracking is considered a canonical application for wireless sensor networks, and work in this area has been motivated in large part by DARPA programs such as SensIT [18].

Zhao *et al.* present the information driven sensor querying (IDSQ) mechanism in [8], [7]. IDSQ is a sensor-to-sensor leader handoff based scheme in which at any given time there is a leader sensor node which makes the decisions about which sensors should be selectively turned on in order to obtain the best information about the target. A combined cost function which gives weight to both energy expenditure and information gain is considered. The generic selective activation strategy which we describe in this paper is closest in spirit to IDSQ. As our focus in this paper is to evaluate general strategies, our findings regarding selective activation are applicable to the performance of intelligent tracking strategies such as IDSQ. Liu *et al.* develop a dual-space approach to tracking targets which also enables selective activation of sensors based on which nodes the target is likely to approach next.

Along these lines, Ramanathan, Brooks, *et al.* advocate a location-centric approach to performing collaborative sensing and target tracking in [13], [14]. The idea is to develop programming abstractions that provide addressing and communication between localized geographic regions within the network rather than individual nodes. This makes localized selective-activation strategies simpler to implement.

Brooks *et al.* present self-organized distributed target tracking techniques with prediction based on Pheromones, Bayesian, and Extended Kalman Filter techniques [21], [22]. The implementation and testing of a real distributed sensor network collaborative tracking algorithm in a military context is described in [23].

A number of recent papers have focused on the question of deploying sensors to ensure adequate coverage of moving targets. Megerian, Meguerdichian, Potkonjak, et al. [20], [19], investigate the question of the minimum exposure path that a target can take in a given sensor field - which is a worst-case metric to evaluate the tracking quality that can be obtained for a given deployment. Clouquer *et al.* [16] use a related metric to evaluate sensor deployment strategies that enhance the worst-case probability of target detection, taking into account factors such as equipment and deployment costs. Chakrabarty, Iyengar *et al.* discuss the problem of tolerating faults while ensuring sensor coverage of an area to ensure that the target moving through the area can be tracked at all times [10]. Jung and Sukhatme examine target tracking by a mobile robotic sensor network in [12].

The problem of multiple targets has also attracted some attention. Bejar, Krishnamachari, *et al.* formulate a sensor tracking problem as that of distributed constraint satisfaction. They show that there is a critical combination of sensing and communication needed to ensure that multiple targets can be tracked satisfactorily by a sensor network. In [15], Li, Wong *et al.* tackle the problem of distinguishing between multiple targets, describing and developing several target classification mechanisms. Fang, Zhao and Guibas describe a distributed mechanism for counting the number of targets in a given field in [9].

In the context of these related works, we should emphasize that our attention is primarily focused on single-target tracking. Our interest is in analyzing and evaluating the fundamental energy-quality tradeoffs involved in tracking with different generic tracking strategies, rather than designing/advocating yet another tracking protocol.

3 Model and Metrics

We now describe the models, assumptions and metrics used in our work.

3.1 Basic Model

We consider a sensor network consisting of N nodes deployed in some operational area, operating for a total time duration T. There is a single target moving through the area. We assume that all sensors in the network are binary detectors with a fixed sensing range S. In other words, at each instant, each sensor returns a '1' if the target is present within a distance S of that sensor, and a '0' otherwise. Given this simple sensor model, we take the centroid of the *locations* of all detecting sensors as an estimate of the target's location at any given time t. Say there are k sensors at locations $X_i = (x_i, y_i),\ i = 1 \ldots k$, detecting the target at time t. Then the estimated location of target $X_s(t) = (x_s(t), y_s(t))$, where

$$x_s(t) = \frac{\Sigma^k x_i}{k} \tag{1}$$

$$y_s(t) = \frac{\Sigma^k y_i}{k} \tag{2}$$

We assume two different modes of operation for each node:

1) A high power tracking mode : Nodes in this mode use a higher power H, which depends on their sensing capabilities. A node in this mode is capable of both sensing a target and also communicating with neighbor nodes.

2) A low power communication mode : Nodes in this mode use a lower power L, which is an indicator of the farthest distance they can communicate. A node in this mode can only communicate with neighbor nodes.

3.2 Quality Metric: Tracking Error

The two performance measures of interest to us in evaluating different tracking strategies are the average total energy expenditure P (averaged over a period of time T), and some measure of the tracking quality, which reflects the uncertainty in the target's location. We use the Euclidean distance between the estimated and actual locations of the target to measure the tracking error. If $X_a(t) = (x_a(t), y_a(t))$ is the actual position of the target at time t, we denote the *instantaneous* tracking error metric as $q(t)$:

$$q(t) = d(X_s(t), X_a(t)) = \sqrt{(x_s(t) - x_a(t))^2 + (y_s(t) - y_a(t))^2} \tag{3}$$

For the time T spent by a target in the area of interest, the time average error, which we denote as Q is given as

$$Q = \frac{1}{T} \int_0^T q(t)dt \tag{4}$$

We note that one drawback of the tracking error metric Q is that it is dependent on the target's specific trajectory[1] $X_a(t),\ t = 0 : T$. An alternative trajectory-independent metric can be obtained by assuming that the target's movement is an *Ergodic* random process, and that its location probability distribution is independent of time. (A random process is ergodic if the time average of any instantiation of the process converges to the mathematical expectation.) Then we can use an alternative tracking error metric Q', the expected distance between the estimated and actual positions of the target:

$$Q' = E[q(t)] = E[\sqrt{(x_s(t) - x_a(t))^2 + (y_s(t) - y_a(t))^2}] \tag{5}$$

[1] Note that in our model, once the location of all N nodes in the network is fixed, and assuming the nodes that are sensing at each time is known, the estimated trajectory $X_s(t)$ can be determined from the actual trajectory $X_a(t)$.

Note that this tracking error metric Q' depends not on a time-dependent trajectory, but rather the probability distribution of the target's location in the operational area.

3.3 Energy Metric: Tracking Energy

For a given tracking strategy, let n_s denote the number of nodes that are in tracking/sensing mode and $n_c = N - n_s$ the number of nodes that are in communication mode. The average energy expenditure for a network of N nodes is then

$$P = (n_s H + n_c L) = P = (n_s H + (N - n_s)L) \tag{6}$$

To simplify our analysis, we assume that the cost of communication is comparable across the different tracking strategies[2]. We therefore compare strategies primarily on the basis of their respective energy expenditure for tracking. To the first order, one can consider the sensing power expenditure as being a power law function of the sensing range S of the nodes: $H(S) = H_0 S^\alpha$, where α could be considered the decay exponent for the sensed signal and would depend upon the sensor modality and deployment factors such as terrain characteristics. Normalizing $H_0 = 1$, we get the following energy metric useful for evaluating a tracking strategy:

$$P_t = n_s H = n_s S^\alpha \tag{7}$$

4 Tracking Strategies

We now describe some general tracking strategies:

- **Naive activation (NA):** In naive activation, all nodes in the network are in tracking mode all the time. While clearly this strategy offers the worst energy efficiency, it is a useful baseline for comparison because it provides the best possible quality of tracking. For this strategy, we have that

$$n_{s,NA} = N \tag{8}$$
$$P_{t,NA} = NS^\alpha \tag{9}$$

- **Randomized activation (RA):** In this strategy, each node is on with a probability p. On average a fraction p of all the nodes will be on and in tracking mode. In this case,

$$n_{s,RA} = pN \tag{10}$$
$$P_{t,RA} = pNS^\alpha \tag{11}$$

[2] This is a reasonably valid assumption particularly when one takes into account recent studies suggesting that reception power for wireless sensor nodes can be comparable to the transmission power.

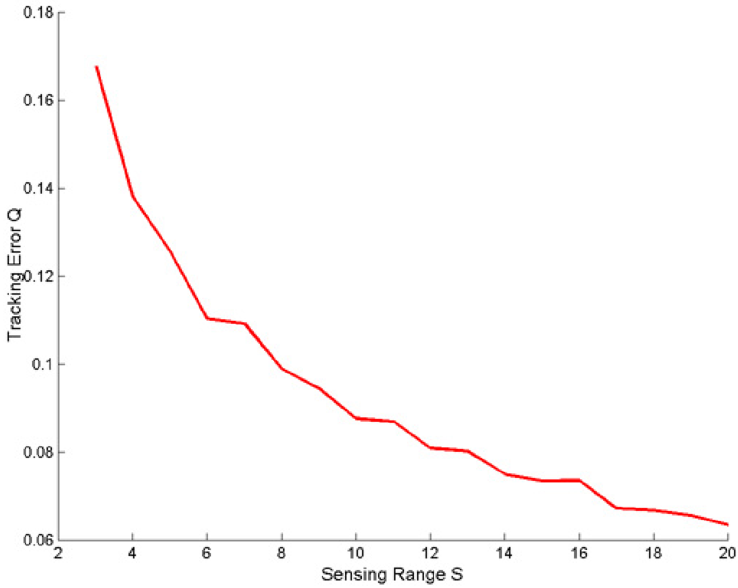

Fig. 1. Tracking Error versus Sensing Range for Naive Activation

– **Selective activation based on prediction (SA):** In this strategy, only a small subset of all the nodes are in tracking mode at any given point of time. They also predict the "next" position of the target and hand over tracking to nodes best placed to track the target in the "next" position. The rest of the nodes are in communication mode and can switch to tracking mode on being alerted by signals from tracking nodes.

  Let X_a be the actual position of the target, and $X_b = X_s$ the belief position of target as before; define X_p as the predicted target position. The idea of selective activation is to use prior history of X_b to determine X_p for the next step. (For example, if we discretize time, knowing sensors could use a simple linear predictor to predict the next location of the target $X_p(t+1)$, using the two latest previous belief positions to estimate the target velocity and assuming that it will continue to move in a straight line). All the sensors within a circle of radius S_p around $X_p(t+1)$ are then alerted to start sensing. Only the sensors within the sensing range S of the actual position $X_a(t+1)$ can possibly sense the target. Hence, the sensors lying in the overlap of the two circles sense the target and the new belief location $X_b(t+1)$ is obtained by finding the centroid of the positions of these sensors. This is illustrated in figure 3. With selective activation based on prediction, only the sensors within a radius S_p around X_p are in tracking mode at any point of time. If ρ is the density of deployment, we get

$$n_{s,SA} = \pi S_p^2 \rho \tag{12}$$

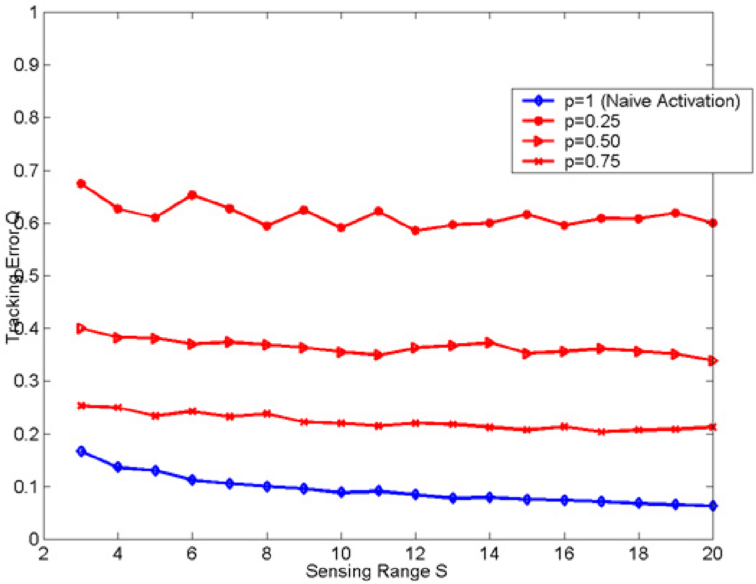

Fig. 2. Tracking Error versus Sensing Range for Random Activation

$$P_{t,SA} = \pi {S_p}^2 \rho S^\alpha \tag{13}$$

– **Duty-cycled Activation (DA):** In duty-cycled activation, the entire sensor network periodically turns off and on with a regular duty cycle. One key feature of duty-cycled activation is that it can actually be used in conjunction with any other activation strategy for target tracking (including NA, RA and SA). Let T_D be the period of the cycle, t_{ON} the on-time, and $n_{s,U}$ be the average number of tracking sensors in the underlying activation strategy U. Then

$$n_{s,DA} = \frac{n_{s,U} t_{ON}}{T_D} \tag{14}$$

$$P_{t,DA} = \frac{P_{t,U} t_{ON}}{T_D} = \frac{n_{s,U} S^\alpha t_{ON}}{T_D} \tag{15}$$

5 Experiments and Results

In the previous two sections we have developed useful common metrics for energy and tracking quality based on our sensor network model, described the tracking strategies we will consider, and their energy expenditure model. In order to compare these strategies, we now turn to simulation experiments.

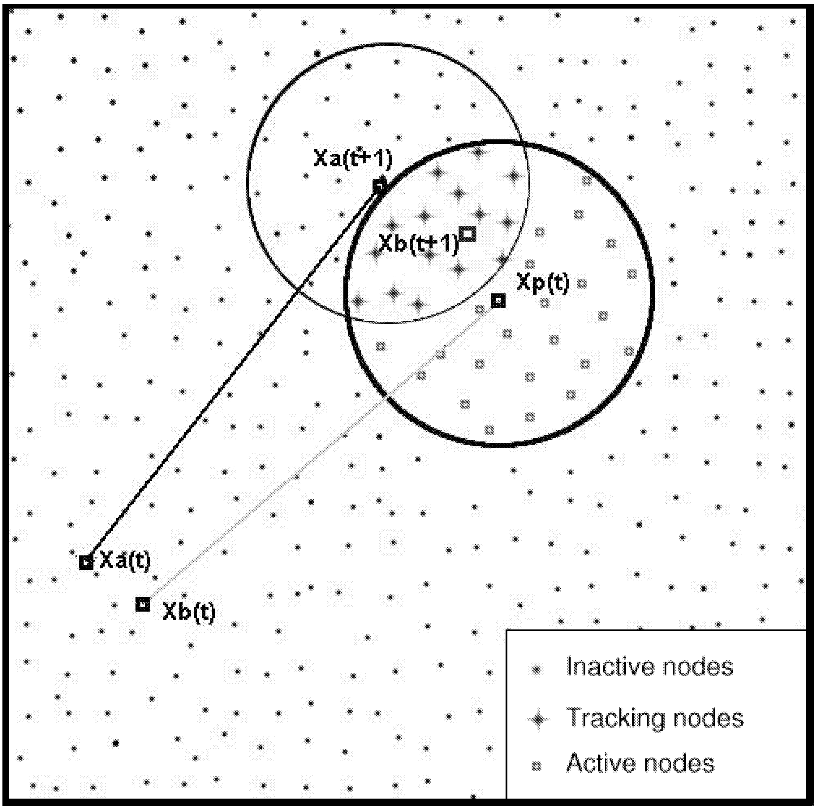

Fig. 3. Illustration of Selective Activation (note: $X_b = X_s$, the believed position)

We simulated a virtual large scale sensor network on a 200 unit x 200 unit area with random placement of sensors and density of deployment ρ= 1 sensor/unit area (i.e. a total of 40000 nodes). Linear, sinusoidal and other reasonable trajectories for the target motion were considered. To avoid edge effects in estimating uncertainty, our calculations are for trajectories in which the target stays away from the boundaries of the region. In the results presented, the target is assumed to follow a representative trajectory of the form $y(t) = Ax^B(t) + CsinDx(t) + E$.

5.1 Performance of Naive Activation, Random Activation, and Selective Activation

Since we are using the centroid of the sensors tracking at any point of time as the sensed position, this estimate can be improved by considering a larger number of sensors. One way of achieving this is to increase the sensing range S. Figure 1 shows how tracking error decreases with S for naive activation. Similarly, figure 2 shows the performance of randomized activation for different

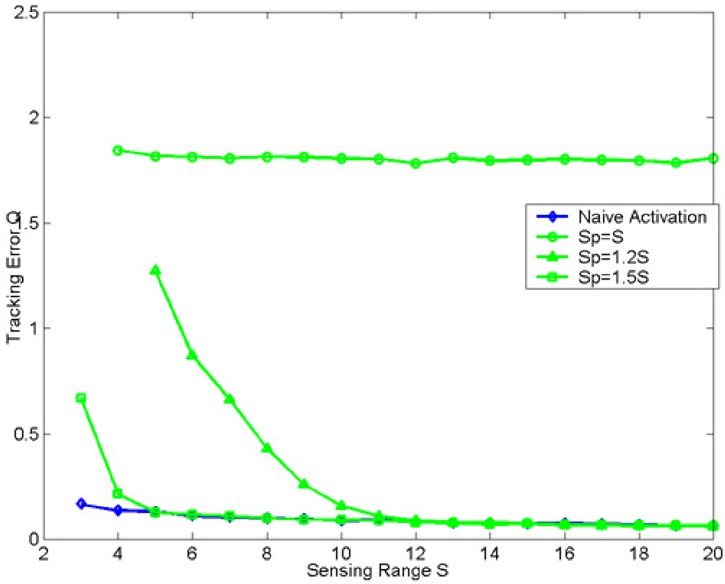

Fig. 4. Tracking Error versus Sensing Range for Selective Activation

values of p. It can be seen that the tracking quality of network-wide randomized activation deteriorates significantly as p is decreased. We also observe that while increasing S does result in a decrease in the tracking error, the decrease is not very substantial and diminishes with increasing S. This evidence of diminishing returns on quality leads us to conclude that it is best not to set the sensing range within the network too high.

Figure 4 compares the performance of selective activation with different settings of S_p. Naive activation is also plotted in the same figure as a baseline. It can be seen that the tracking error is quite high when $S_p = S$. In predictive selective-activation, as the intersection area of the two circles (the circle of radius S around the actual position and the circle of radius S_p around the belief position) becomes larger, sensors closer to the target's actual position are activated. This can be achieved by increasing S_p. For the particular trajectory considered, we find that selective activation with $S_p = 1.5S$ performs nearly as well as a naive network.

Figure 5 shows the energy-quality tradeoff between the NA, RA and SA strategies. It is a plot of the tracking error vs $log(P_t)$ for these strategies, with respect to the energy metric in log scale (as defined in section 3). In this figure, data points to the bottom left represent dominating, Pareto-optimal strategies, since they represent low tracking error (hence high tracking quality) as well as low energy expenditure. It is clear from the figure that selective activation with

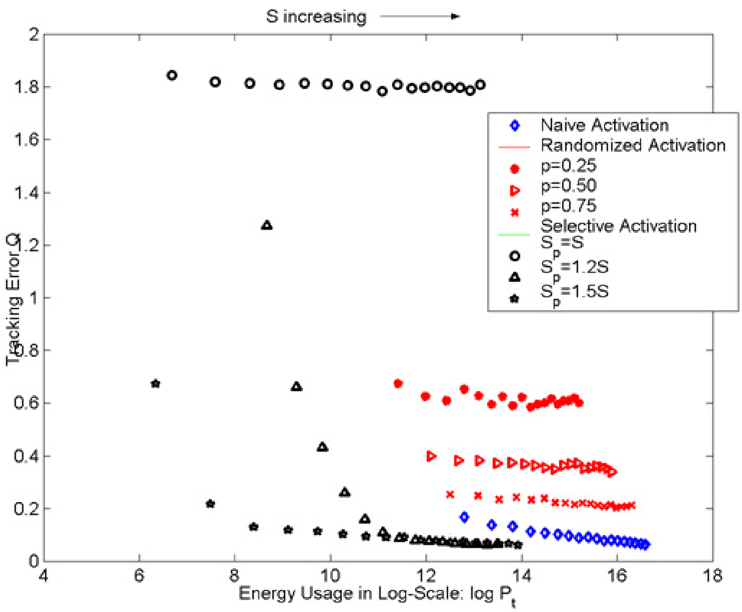

Fig. 5. Energy-Quality Tradeoff for Basic Activation Strategies: NA, RA, SA

reasonably high S_p is a dominating strategy. It provides overall significantly reduced traffic error for low energy expenditure.

Clearly, selective activation can provide a dominating design in terms of the energy-quality tradeoffs considered. Figure 5 shows that selective activation with optimal settings can offer 4 orders of magnitude savings in energy (corresponding to the size of the network) compared to NA or RA, for essentially the same tracking quality. Also, the sensing range should be chosen carefully and kept to a minimum based on the desired quality in order to effect the best tradeoff. For selective activation, the results suggest using the lowest feasible value of S and corresponding S_p. In general, the feasible values of S and S_p would depend on the mobility model of the target. The average speed of the target can provide a good indicator for determining these parameters. We found that the results do not vary much with trajectory for comparable values of target speed.

5.2 Performance of Duty-Cycled Activation

Let us now turn to the final strategy: duty-cycled activation. Let us understand the functioning of this scheme. If we consider a particular time period T_D, the instantaneous tracking error during time t_{ON} would be the same as for the network without duty-cycling (let's call this $q_U(t)$). However, once the network is shut down, the tracking error increase with time until the next time period starts - this is due to the drift between the estimated target location and the

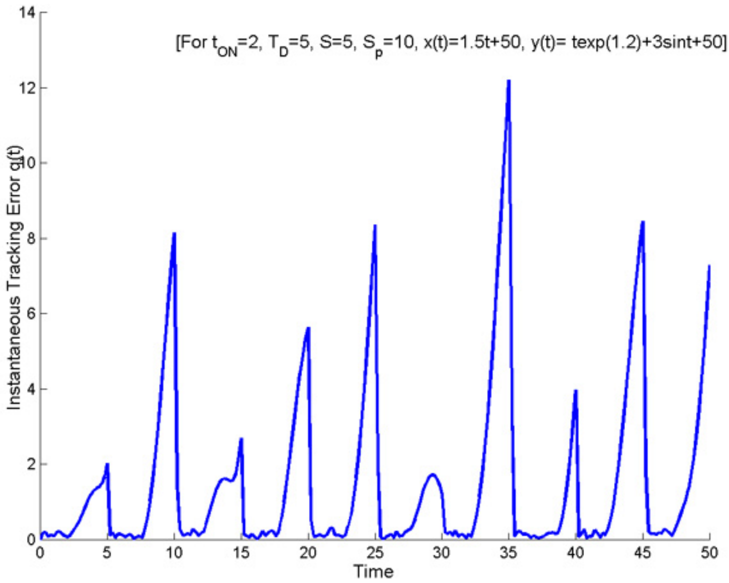

Fig. 6. Instantaneous Tracking Error versus Time for Selective Activation with Duty Cycling

actual target movement during the off-time. For the time period T_D, if v is the mean target speed, the tracking error at time t is

$$q(t) \simeq \begin{cases} q_U(t) & ,0 < t < t_O N \\ q_U(t) + v(t - t_{ON}) & ,t_{ON} < t < T_D \end{cases} \tag{16}$$

Hence the average tracking error for duty-cycled activation Q_{DA} can be approximated as

$$Q_{DA} \simeq Q_U + 0.5v(1 - \frac{t_{ON}}{T_D})^2 T_D \tag{17}$$

As we noted before, DA can be used in conjunction with other underlying strategies. Since our previous results have shown that selective activation is a dominating strategy, we focus on this combination: duty-cycled selective activation. Figure 6 shows a sample run illustrating how instantaneous tracking error varies with time for selective activation with duty-cycling. Figure 7 shows (as suggested by equation (17)) that for the same ratio t_{ON}/T_D, the average tracking error Q increases with the period T_D. Given an acceptable value for the tracking error and the mobility model of the target (v), the above approximation can help us arrive at the feasible values of T_D (t_{ON} should be kept to the minimum possible value, which might depend on the time-constants associated with device start-up and shut-down).

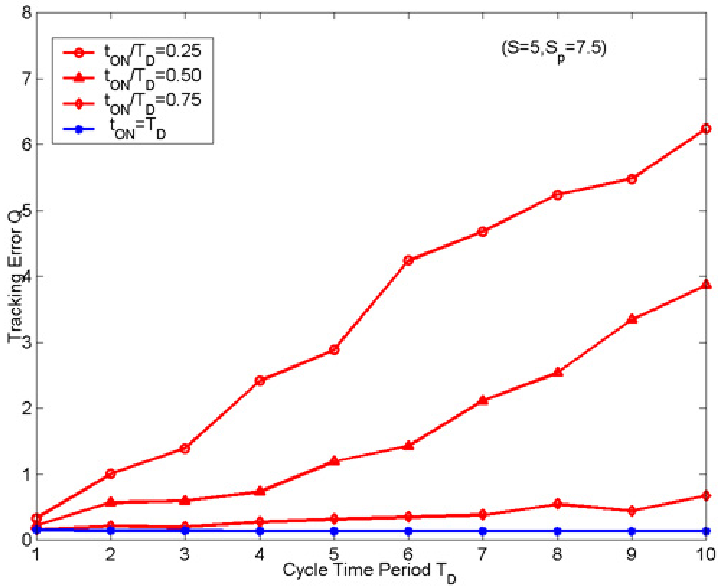

Fig. 7. Tracking Error versus Cycle Time for Selective Activation with Duty Cycling

Figure 8 shows the tracking error varies with energy usage when choosing different values of T_D and t_{ON}. The figure shows that duty-cycled activation is a flexible and efficient mechanism for tuning the energy-quality tradeoff of tracking.

6 Conclusions

The following is a summary of the main contributions of this paper:

- We identified four generic sensor activation strategies for target tracking that can be used to provide different energy-quality tradeoffs: naive activation, random activation, selective activation with prediction and duty-cycled activation.
- We developed simple metrics to evaluate the performance of these strategies with respect to energy usage and tracking quality.
- We examined how tracking performance for the basic strategies (NA, RA, SA) varies with sensor range, showing that there are diminishing returns in terms of tracking quality. This suggests that sensor range settings should be carefully chosen and kept to a minimum with these strategies.
- We showed that with the right parameters selective activation can provide orders of magnitude improvements in energy usage with near-optimal track-

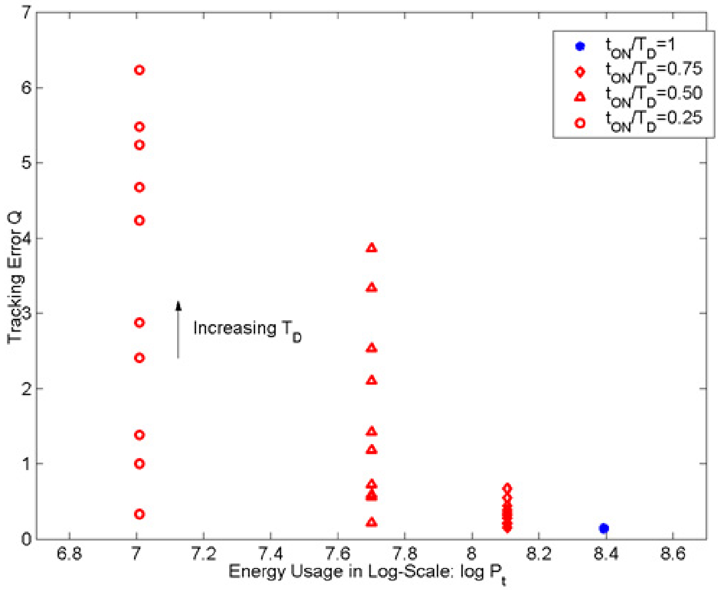

Fig. 8. Energy-Quality Tradeoff for Selective Activation with Duty Cycling

ing quality. With respect to random and naive activation, SA is a dominating strategy with Pareto-optimal points on an energy-quality plot.

– We then examined duty-cycled activation. Our analysis showed that for best energy performance the ratio t_{ON}/T_D should be kept as small as possible, while minimizing T_D improves the tracking quality. This allows us to use t_{ON} and T_D as tuning knobs to effect a flexible tradeoff between energy and tracking quality in conjunction with other base strategies such as selective activation.

Although we have taken a significant step in this direction, as future work, we would like to extend the mathematical treatment of the energy-quality tradeoffs involved in tracking. This will require the use of more tractable assumptions about the target mobility model. We would also like to consider richer sensor models and energy cost models to validate the generality of our findings.

References

1. A. Cerpa *et al.*, "Habitat monitoring: Application driver for wireless communications technology," *2001 ACM SIGCOMM Workshop on Data Communications in Latin America and the Caribbean*, Costa Rica, April 2001.
2. G.J. Pottie, W.J. Kaiser, "Wireless Integrated Network Sensors," *Communications of the ACM*, vol. 43, no. 5, pp. 551–8, May 2000.

3. J. Warrior, "Smart Sensor Networks of the Future," *Sensors Magazine*, March 1997.
4. D. Estrin, R. Govindan, J. Heidemann and S. Kumar, "Next Century Challenges: Scalable Coordination in Sensor Networks," *ACM/IEEE International Conference on Mobile Computing and Networks (MobiCom '99)*, Seattle, Washington, August 1999.
5. D. Estrin *et al. Embedded, Everywhere: A Research Agenda for Networked Systems of Embedded Computers*, National Research Council Report, 2001.
6. I. Akyildiz, , W. Su, Y. Sankarasubramaniam, and E. Cayirci, "A Survey on Sensor Networks," *IEEE Communications Magazine*, Vol. 40, No. 8, pp. 102–114, August 2002.
7. M. Chu, H. Haussecker, F. Zhao, "Scalable information-driven sensor querying and routing for ad hoc heterogeneous sensor networks." *International Journal on High Performance Computing Applications*, vol. 16, no. 3, Fall 2002.
8. F. Zhao, J. Shin, J. Reich, "Information-Driven Dynamic Sensor Collaboration for Tracking Applications." *IEEE Signal Processing Magazine*, March 2002.
9. Q. Fang, F. Zhao, L. Guibas, "Counting Targets: Building and Managing Aggregates in Wireless Sensor Networks." Palo Alto Research Center Technical Report pp. 2002–10298, June 2002.
10. K. Chakrabarty, S. S. Iyengar, H. Qi, E.C. Cho, "Grid Coverage of Surveillance and Target location in Distributed Sensor Networks" *To appear in IEEE Transaction on Computers*, May 2002.
11. R. Bejar, B. Krishnamachari, C. Gomes, and B. Selman, "Distributed constraint satisfaction in a wireless sensor tracking system," *Workshop on Distributed Constraint Reasoning, International Joint Conference on Artificial Intelligence*, Seattle, Washington, August 2001.
12. Jung, B. and Sukhatme, G.S. "Tracking Targets using Multiple Robots: The Effect of Environment Occlusion", *Autonomous Robots*, 2002.
13. P. Ramanathan, "Location-centric Approach for Collaborative Target Detection, Classification, and Tracking," *IEEE CAS Workshop*, 2002.
14. R. R. Brooks, P. Ramanathan, and A. Sayeed, "Distributed Target Tracking and Classification in Sensor Networks," *Proceedings of the IEEE*, Invited Paper, Submitted for review, September 2002.
15. D. Li, K. Wong, Y. Hu and A. Sayeed, "Detection, Classification, Tracking of Targets in Micro-sensor Networks," *IEEE Signal Processing Magazine*, pp. 17–29, March 2002
16. T. Clouqueur, V. Phipatanasuphorn, P. Ramanathan and K. K. Saluja, "Sensor Deployment Strategy for Target Detection," *The First ACM International Workshop on Wireless Sensor Networks and Applications (WSNA'02)*, Sep. 2002.
17. J. Liu, P. Cheung, L. Guibas, and F. Zhao, "A Dual-Space Approach to Tracking and Sensor Management in Wireless Sensor Networks," *The First ACM International Workshop on Wireless Sensor Networks and Applications (WSNA'02)*, Sep. 2002.
18. J. Reich, *SensIT Collaborative Signal Processing Scenario Discussions*, (http://www2.parc.com/spl/projects/cosense/pub/tracking_benchmarks.pdf).
19. S. Megerian, F. Koushanfar, G. Qu, G. Veltri, M. Potkonjak, "Exposure In Wireless Sensor Networks: Theory And Practical Solutions", *Journal of Wireless Networks, Vol. 8, No. 5, ACM Kluwer Academic Publishers*, pp. 443–454, September 2002.
20. S. Meguerdichian, F. Koushanfar, G. Qu, M. Potkonjak, "Exposure In Wireless Ad Hoc Sensor Networks", *International Conference on Mobile Computing and Networking (MobiCom '01)*, pp. 139–150, Rome, Italy, July 2001.

21. R. Brooks, and C. Griffin, “Traffic model evaluation of ad hoc target tracking algorithms,” *Journal of High Performance Computer Applications*, Accepted, 2002.
22. R. Brooks, C. Griffin, and D. S. Friedlander, “Self-Organized distributed sensor network entity tracking,” *International Journal of High Performance Computer Applications, special issue on Sensor Networks*, vol. 16, no. 3, Fall 2002
23. J. Moore, T. Keiser, R. R. Brooks, S. Phoha, D. Friedlander, J. Koch, A. Reggio, and N. Jacobson, “Tracking Targets with Self-Organizing Distributed Ground Sensors,” *2003 IEEE Aerospace Conference*, Invited Paper, Accepted for publication, November 2002.