MULTI-STEP INFORMATION-DIRECTED SENSOR QUERYING IN DISTRIBUTED SENSOR NETWORKS

Juan Liu
Palo Alto Research Center
3333 Coyote Hill Road
Palo Alto, CA 94304

Dragan Petrovic
Department of EECS
U. C. Berkeley

Feng Zhao
Palo Alto Research Center
3333 Coyote Hill Road
Palo Alto, CA 94304

ABSTRACT

Sensor tasking is essential to many sensing applications in resource-constrained wireless ad hoc sensor networks. In this paper, we present a multi-step lookahead algorithm for sensor selection and information routing. The algorithm is based on the information-driven sensor querying (IDSQ) that uses mutual information as a utility meature for potential information contribution of individual sensors, and extend it to prediction of information gain over a finite horizon while balancing cost such as the number of communication hops. Simulation results on target tracking problems have shown that the multi-step lookahead algorithm significantly improves the tracking performance compared to the original greedy algorithm, when "sensor holes" are present in a sensor network.

1. INTRODUCTION

The study of distributed sensor networks is an emerging, interdisciplinary research area, with contributions from fields such as signal processing, communication, distributed algorithms, and MEMS sensor technology. Recently, there has been an increasing interest in using distributed sensor networks for large-scale sensing applications such as environmental monitoring, security surveillance, and battlefield awareness [1]. The primary reason is the flexibility and scalability of ad hoc sensor networks. Unlike traditional sensor arrays, a distributed sensor network can be flexibly deployed in an area where there is no a priori sensing infrastructure. By distributing computation over the network and invoking sensor collaboration within a local region, a distributed sensor network is more scalable than traditional centralized sensor array processing systems.

The sensor nodes in a large-scale sensor network are often battery-operated and bandwidth-limited. A major challenge in sensor network design is to optimize performance under various resource constraints. It becomes critical to carefully select the embedded sensor nodes that participate in the sensor collaboration, balancing the potential contribution of each sensor against its resource consumption or potential utility of other user. This is especially important in dense networks, where many measurements may be highly redundant.

Several approaches to sensor collaboration have been proposed in the sensor network literature. For example, Brooks et al. described a prediction-based sensor collaboration that uses estimation of target velocities to activate regions of sensors in the direction of movement [2]. This heuristic method provides a rough estimate of the relevance of nodes for the tracking task in the immediate neighborhood. Zhao et al. [3, 4] introduced the more general concept of information-driven sensor querying (IDSQ). The idea is to quantify information gain with measures such as Mahanolobis distance, Fisher information, or estimation entropy, and use the measured information gain to guide sensor selection and information routing. The sensor with the highest information gain is activated to participate in the collaboration. A greedy version of the IDSQ method suffers from getting trapped at local maxima when a sensor network contains "sensor holes". In this paper, we present a multi-step lookahead algorithm for IDSQ that uses more global knowledge of information gain to guide the sensor selection and routing.

2. INFORMATION-DIRECTED SENSOR COLLABORATION

While the idea of information-directed sensor collaboration applies to general sensing problems using wireless sensor networks, here we illustrate the sensor collaboration in the context of target tracking.

Consider the problem of tracking a target in a 2D region. The goal is to estimate the target location $x^{(t)}$ based on sensor measurements up to time t, denoted as $\overline{z^{(t)}} = \{z^{(0)}, z^{(1)}, \cdots, z^{(t)}\}$, using a set of spatially distributed sensors. We follow the IDSQ framework [3, 4] in which sensors are activated based on their utility and cost. Zhao et al. discuss IDSQ in conjunction with a leader-based tracking scheme, where at any time only one sensor, called

This work is partially supported by the DARPA Sensor Information Technology program under Contract F30602-00-C-0139.
The email addresses of the authors are: juan.liu@parc.com, dragan@eecs.berkeley, and zhao@parc.com.

the leader node, is active, while the rest of the network is idle. The leader node combines its measurement with the previous estimation to produce a posterior belief of the target state $p(x^{(t)}|\overline{z^{(t)}})$, using a sequential Bayesian estimation algorithm. Based on this belief, the leader selects a sensor with the best predicted information from those that are one hop away, and send the current belief to that sensor. The original leader goes back to sleep, and the new leader which receives the belief repeats the process of sensing, estimation, and leader selection. The leader-based selective sensor activation scheme has several advantages. It prolongs network lifetime and lowers probability of detection by turning off unnecessary sensors. It allows a network to support multiple concurrent user operations by leaving unused sensors available for other sensing tasks.

We use mutual information as a metric for quantifying information contribution from individual sensors. Mutual information is calculated based on statistical observation models, and provides a common ground for comparing different type of sensors. Therefore, it is suitable for heterogeneous sensor network systems and can be easily combined with statistical information fusion methods. Furthermore, it is applicable to models with arbitrary probability density functions (pdf) and is not restricted to Gaussian models.

The mutual information between two random variables U and V with a joint probability density function $p(u,v)$ is defined as

$$\begin{aligned} I(U;V) &\triangleq E_{p(u,v)}\left[\log\frac{p(u,v)}{p(u)p(v)}\right] \\ &= D(p(u|v)||p(u)), \end{aligned}$$

where $D(\cdot||\cdot)$ is the relative entropy between two distributions, also known as the Kullback-Leibler divergence [5]. It indicates how much information V conveys about U. From a data compression perspective, it measures the savings in bits of encoding U if V is already known. In classification and estimation problems, mutual information can be used to establish performance bounds.

Using the information metric, a leader selects the most "informative" sensor according to

$$k_{IDSQ} = \arg\max_{k\in\mathcal{N}} I(X^{(t+1)}; Z_k^{(t+1)}|\overline{Z^{(t)}} = \overline{z^{(t)}}), \quad (1)$$

where $\mathcal{N}$ denotes the set of sensors the leader node can talk to, and $p(x^{(t)}|\overline{z^{(t)}})$ is the current belief state. Essentially, this criterion selects a sensor whose measurement $z_k^{(t+1)}$, combined with the current measurement history $\overline{z^{(t)}}$, would provide the greatest amount of information about the target location $x^{(t+1)}$. The mutual information can be interpreted as Kullback-Leibler divergence between $p(x^{(t+1)}|\overline{z_k^{(t+1)}})$ and $p(x^{(t+1)}|\overline{z^{(t)}})$, the belief after and before applying the new measurement $z_k^{(t+1)}$, respectively. Therefore, this criterion favors the sensor which on average gives the greatest change to the current belief.

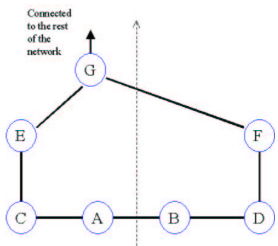

Fig. 1. An example of routing in the presence of a sensor hole.

3. PROBLEMS OF GREEDY IDSQ

The sensor selection criterion (1) is greedy: the leader selects among its neighbors the most "informative" sensor for the next step. While it is highly efficient and uses no global knowledge, the greedy algorithm may get stuck in local information maxima.

Consider the following sensor network example, as plotted in Fig. 1, where circles denote sensor nodes and edges one-hop communication links. Neighbor nodes are those that are one-hop away in the graph, e.g., nodes B and C are neighbors of node A. Since the problem with the greedy algorithm is independent of the choice of information measure, we use the simpler metric of inverse of Euclidean distance between a sensor and the target. The closer a sensor is from the target, the more informative it is in this metric. Consider the case that the target moves along the dashed line between A and B. At time $t = 0$, node A is the leader, and selects a sensor to query among its neighbors B and C. Sensor B has a higher information gain and therefore becomes the leader for time $t = 1$. B then selects A since it is its most informative neighbor. The sensor handoff keeps going back and forth between nodes A and B, while the target moves away. The leadership never gets to nodes E, F, or G, who might be closer to the target as it moves up. The culprit in this case is the "sensor hole" the target went through. In general, inhomogeneity in sensor layout can cause problems for the greedy algorithm due to its lack of knowledge beyond the immediate neighborhood.

4. MULTI-STEP LOOKAHEAD IDSQ

The above example demonstrates that locally optimal choice of sensors can be detrimental to overall optimality of tracking. To alleviate this problem, we propose a multi-step sensor selection algorithm for IDSQ that optimizes the sensor choice based on a prediction of information gain over a lookahead horizon. Instead of selecting a single sensor as in (1), we find a path in the network along which the sensors maximize the accumulated information gain. The path length, denoted by M, is the number of lookahead steps and a design parameter. In general, M should be large enough to guide the sensor selection around sensor holes, but not too large to make the computational cost prohibitive.

In general, information gain obtained by incorporating the readings of multiple sensors is upper-bounded by the sum of information gains of individually incorporating each sensor measurement, i.e.,

$$I(X; Z_1, Z_2, \cdots, Z_K) \leq \sum_{k=1,\cdots,K} I(X; Z_k). \quad (2)$$

The exact computation of $I(X; Z_1, Z_2, \cdots, Z_K)$ involves the high-dimensional pdf $p(x, z_1, z_2, \cdots, z_K)$ and is expensive. For computational efficiency, we use the upper bound $\sum_k I(X; Z_k)$ as the objective function to maximize.

With this objective function, the multi-step lookahead algorithm can be stated as follows. Given a M-hop neighborhood of a node, compute a path of length M or shorter with maximum accumulated gain $\sum_k I(X; Z_k)$ along the path. This problem is combinatorial, since the information gain one can collect at a node depends on whether the node has been visited earlier on the path. Revisiting a node does not bring in new measurement, hence the information gain is zero. The optimal solution to the path-finding problem has an exponential complexity. One needs to enumerate all paths up to length M and selects the best one.

We propose a near-optimal path-finding algorithm with a much lower complexity, called the *min-hop* algorithm. The algorithm first selects the destination node as the node within its M-hop neighborhood with the highest information gain. Let m be the minimum number of hops from the current leader to the destination. The algorithm then considers all paths from the source to the destination with m hops (hence the algorithms is called *min-hop*), and selects the path whose accumulated gain is the greatest. The reason for considering only the minimum hop paths is that the information gain is calculated based on the current belief, which becomes less accurate as time advances. The first node along the selected path becomes the leader for the next iteration, and the entire process of path-finding repeats.

Finding the minimum hop path to a destination with maximum accumulated information is simple. The following trick converts the maximum information gain problem into a standard shortest-path problem, which can be solved efficiently using Dijkstra's algorithm [6], with computational complexity $O(N \log N + E)$, where N is the total number of nodes in the local graph, and E is the number of edges. The conversion is designed as follows: for each node i, we assign to each edge going into the node the cost of $L - u_i$, where L is some large number and u_i is the information utility value at node i. Fig. 2 shows an example of local neighborhood before and after the conversion. The information gain of each node is marked next to the node in

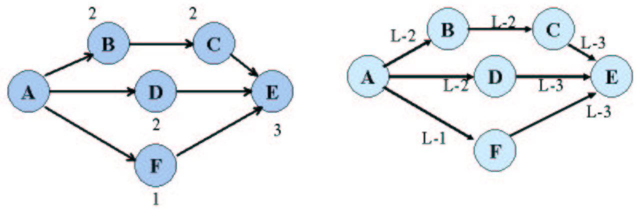

(a) (b)

Fig. 2. (a) local neighborhood of the current leader A; (b) the converted graph for shortest-path algorithm.

Fig. 2a. The cost of edges are marked in Fig. 2b. The shortest path in terms of accumulated edge cost in the converted graph is the same as the path with maximum accumulated information gain. This is easy to see: paths with minimum number of hops wins over longer paths since L is large, and all min-hops have the same number of hops, hence the one with the maximum gain wins. To ensure the correctness of this algorithm, we let $L > E \cdot u_{max}$, where u_{max} is the maximum information gain in the local neighborhood.

The min-hop algorithm reduces computation by considering only a subset of the paths. It finds the optimal information gain path among a family of paths with minimum hop length from the current leader to the node with maximum information value. The selected path provides a good tradeoff between information gain and communication cost. Strictly speaking, it is suboptimal. For example, in Fig. 2a, the maximum information gain path from A to E is $ABCE$, with accumulated gain of 7. The min-hop algorithm returns the path ADE, with accumulated gain of 5. Clearly, the algorithm trades optimality with computational complexity.

5. SIMULATION RESULTS

Simulations for a tracking scenario were carried out to validate and characterize the performance of the proposed multi-step lookahead algorithm. The sensor network used consists of two types of sensors, acoustic amplitude sensors and direction-of-arrival (DOA) sensors. The acoustic amplitude sensors output sound amplitude measured at each microphone, and estimate the distance to a target based on the physics of sound attenuation. The DOA sensors are small microphone arrays. Using beamforming techniques, they determine the direction where the sound comes from, i.e., the bearing of the target. The detailed description of these two types of sensors can be found in [7].

In the simulation, the domain where sensors are laid down is a 225×375 m^2 region. The sensor layout is generated as follows (see Fig. 3). First, a uniform grid of 15 rows and 6 columns is generated to evenly cover the region. A sensor hole is created by removing the grid points with rows 5-6 and columns 2-5, assuming lower-left corner is

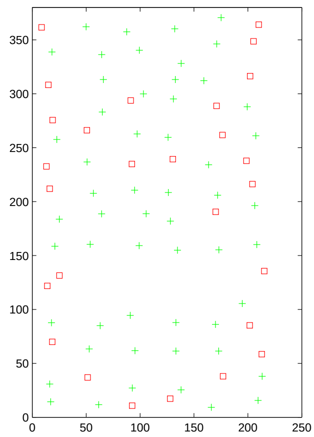

Fig. 3. Simulated sensor layout. The points marked with a "+" denote amplitude sensors, ant the points marked with a square denote the DOA sensors.

the origin. The remaining grid is perturbed with a Gaussian noise $N(0, 25)$. The sensor network consists of 70% amplitude sensors and 30% DOA sensors, randomly spread over the sensor region. Each sensor can directly communicate to eight nearest neighbors around itself, corresponding to the 3×3 subgrid centered at the sensor. The target travels from south to north along a straight line in the middle of the field, with a speed of $v = 15$ m/s. A sequential Bayesian estimation algorithm, whose details can be found in [7], is used to track the target location.

We compare the performance of a greedy algorithm and a 3-step min-hop lookahead algorithm. The performance is reported in Table 1. The results are averaged over 100 runs. Due to the presence of sensor hole, the greedy algorithm often loses track of the target. In our simulations, a track is considered lost if by the time the vehicle reaches the north end of the sensor field, the estimate of target location of the last five steps is on average more than 60 meters away from the true target location. By this standard, in 93% of the simulated runs, the greedy IDSQ algorithm loses track, and the 3-step min-hop lookahead loses track in only 19% of the runs. The improvement is significant. Also of interest is the mean-squared error (MSE) and the variance of the belief state $p(x^{(t)}|\overline{z^{(t)}})$ for the "good" runs in which the track is not lost. MSE measures the tracking accuracy, and the variance shows the compactness (hence the confidence) of the belief. These statistics are also shown in Table 1. The numbers are averaged over 100 runs and all time steps. Note that the MSE using the greedy IDSQ algorithm is twice higher than using the 3-step min-hop algorithm, indicating that even the greedy algorithm did recover from the sensor hole and keep track of the target, the recovery is slow, causing the average MSE to be high. From the results we can see that the 3-step min-hop algorithm is much more robust against the presence of sensor holes.

	Target loss prob.	statistics of good runs	
		MSE	variance
greedy IDSQ	93%	23.61	261.7
3-step min-hop	19%	11.95	281.9

Table 1. Simulation results: tracking performance averaged over all tracking steps and 100 runs.

6. DISCUSSIONS

This paper addresses the important problem of jointly optimizing for estimation quality and information routing cost in ad hoc sensor networks. A sensor network is typically designed for one or more sensing and information gathering tasks. Routing algorithms that minimize communication cost alone may find information poor paths.

We have presented an information-based optimization technique, the multi-step lookahead IDSQ algorithm, that maximizes information utility for a sensing task and at the same time balances with resource usage. Simulation results have shown that the technique can significantly improve estimation quality for a target tracking problem in the presence of sensor holes in the network.

In the formulation we presented, the objective function is optimized with respect to the upper bound of the information gain along a path. Hence, the resulting paths are not strictly optimal, but are close to the optimal ones when the upper bound is tight, i.e., the measurement at each sensor is as dependent of each other as possible. As mentioned earlier, solving the optimization problem with mutually dependent measurements is computationally prohibitive. Finding other efficient approximate solutions remains as a future research topic.

7. REFERENCES

[1] D. Estrin, L. Girod, G. Pottie, and M. Srivastava, "Instrumenting the world with wireless sensor networks," in *Proc. ICASSP 2001*, (Salt Lake City, Utah), May 2001.

[2] R. Brooks, C. Griffin, and D. Friedlander, "Self-organized distributed sensor network entity tracking," *International Journal of High-Performance Computing Applications*, vol. 16, no. 3, Fall 2002.

[3] F. Zhao, J. Shin, and J. Reich, "Information-driven dynamic sensor collaboration," *IEEE Signal Processing Magazine*, vol. 19, no. 2, pp. 61–72, March 2002.

[4] M. Chu, H. Haussecker, and F. Zhao, "Scalable information-driven sensor querying and routing for ad hoc heterogeneous sensor networks," *International Journal of High-Performance Computing Applications*, vol. 16, no. 3, Fall 2002.

[5] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. New York, NY: John Wiley and Sons, Inc., 1991.

[6] T. H. Cormen, C. E. Leiserson, and R. L. Rivest, *Introduction to Algorithms*. Cambridge, MA, and New York, NY: MIT Press and McGraw–Hill, 1989.

[7] J. Liu, J. E. Reich, and F. Zhao, "Collaborative in-network processing for target tracking," *EURASIP, Journal on Applied Signal Processing*, to appear in 2002.