*IEEE Signal Processing Magazine, March 2002, to appear*



# Information-Driven Dynamic Sensor Collaboration for Tracking Applications

Feng Zhao[1] and Jaewon Shin and James Reich, *Members, IEEE*
Xerox Palo Alto Research Center, 3333 Coyote Hill Road, Palo Alto, California, CA 94304
{zhao,jshin,jreich}@parc.Xerox.com

***Abstract*--This article overviews the information-driven approach to sensor collaboration in ad hoc sensor networks. The main idea is for a network to determine participants in a "sensor collaboration" by dynamically optimizing the information utility of data for a given cost of communication and computation. A definition of information utility is introduced, and several approximate measures of the information utility are developed for reasons of computational tractability. We illustrate the use of this approach using examples drawn from tracking applications.**

## I. Introduction

The technology of wirelessly networked micro-sensors promises to revolutionize the way we live, work, and interact with the physical environment. For example, tiny, inexpensive sensors can be "sprayed" onto roads, walls, or machines to monitor and detect a variety of interesting events such as highway traffic, wildlife habitat condition, forest fire, manufacturing job flow, and military battlefield situation.

Because of its spatial coverage and multiplicity in sensing aspect and modality, a sensor network is ideally suited for tracking moving phenomena (e.g. moving vehicles or people) traversing the range of many sensors in a large area, monitoring a large number of objects or events simultaneously (e.g. forest fires or large animal herds), or detecting low-observable events (e.g. stealthy, low signal-noise ratio sources, subject to loud distractors or other countermeasures). The detection, classification, and tracking of moving, non-local, low-observable events requires non-local collaboration among sensors. Aggregation of a multitude of sensor data can improve accuracy. Informed selective collaboration of sensors, in contrast to flooding data requests to all sensors, can reduce latency. Moreover, sensor collaboration can minimize bandwidth consumption (translating into energy savings) and mitigate the risk of network node/link failures. The longevity of a network depends on the rate the power is consumed performing computation and communication tasks at these un-tethered, battery powered sensors.

Therefore, one of the central issues for Collaborative Signal and Information Processing (CSIP) to address is *energy-constrained dynamic sensor collaboration* — how to dynamically determine who should sense, what needs to be sensed, whom the information must be passed on to. For non-local spatio-temporal events due to motion or spatial multiplicity of targets, sensor collaboration can dynamically invoke regions of a network informed by motion prediction as in tracking, or activate sensors around which there has been a significant change in physical measurement as in large-scale event monitoring (e.g. waking up sensors that are on the boundary of a forest fire). For low-observable events, sensor collaboration can selectively aggregate multiple sources of information to improve detection accuracy, or to actively probe certain nodes in order to disambiguate multiple interpretations of an event.

Although the ideas introduced in this paper apply to a wide range of distributed detection, classification, and monitoring problems, we will focus on the tracking problem and use it as a running example to introduce the techniques of information-driven dynamic sensor collaboration. We assume that each sensor has a local sensing and communication range; a physical phenomenon of interest has a local or non-local spatial extent. In addition, we assume that each node in the network can locally estimate the cost of sensing, processing, and communicating data to another node, and can monitor its power usage. The benefits of sensor collaboration can be measured as improvement in one or more of the following capabilities:

1. Detection quality: Detection resolution, sensitivity, and dynamic range; misses and false alarms; response latency.
2. Track quality: Tracking errors, track length, robustness against sensing gaps.
3. Scalability: Size of network, number of events, number of active queries.
4. Survivability: Robustness against node/link failures.
5. Resource usage: Power/bandwidth consumption.

The main idea of the information-driven approach is to base the decision for sensor collaboration on information constraints as well as constraints on cost and resource consumption. Using measures of information utility, the sensors in a network can exploit the information content of data already received to optimize the utility of future sensing actions, thereby efficiently managing the scarce communication and processing resources. The information-driven approach to sensor querying and data routing builds on the work of *directed diffusion* routing [Intanagonwiwat et al. 2000; Estrin et al. 1999] that has been successfully deployed for ad hoc sensor networks; it can enhance the diffusion

[1] Corresponding author.

mechanism in directed diffusion to further reduce latency and resource usage.

## II. A TRACKING SCENARIO

To illustrate the main idea of the information-driven approach, we consider a task of tracking a moving vehicle through a two-dimensional sensor field (see Fig. 1). A user initiates the following query: “report the position of the vehicle every 5 seconds”. A few interesting features of the problem are worth noting. There is no road constraint, and therefore no prior knowledge of possible vehicle trajectories can be exploited. Second, the vehicle can accelerate or decelerate, in between the nearest sensors. Both of these render traditional closest-point-of-approach (CPA) based trackers difficult to apply. Third, many sensors can potentially make simultaneous observations and flood the network with the information. This requires the network to make intelligent decisions about who should sense and who should communicate and at what time. For the sake of simplicity, we focus on the sensor collaboration during the tracking phase, ignoring the detection phase and glossing over the details of routing the query into regions of interest. We further assume there is one leader node active at any moment, and its task is to select and route tracking information to the next leader. A multiple, simultaneous leader protocol can be analogously developed, and is beyond the scope of this paper.

1. A user query enters the sensor network at node *Q*.
2. Meta knowledge guides the query towards a region of potential events.
3. Node *a* computes an initial estimate of vehicle state $\hat{\mathbf{x}}_a$, determines the next best sensor $Next(\hat{\mathbf{x}}_a, \lambda_i) = b,\ i \in neighbors(a)$, and hands off the state information to *b*. $\lambda_i$ is sensor characteristics for node *i*.
4. Node *b* computes a new estimate by combining its measurement $\mathbf{z}_b$ with the previous estimate $\hat{\mathbf{x}}_a$ using, say a Bayesian filter: $\hat{\mathbf{x}}_b = \hat{\mathbf{x}}_a \oplus \mathbf{z}_b$; $Next = c$
5. Node *c* computes: $\hat{\mathbf{x}}_c = \hat{\mathbf{x}}_b \oplus \mathbf{z}_c$; $Next = d$
6. Node *d* computes: $\hat{\mathbf{x}}_d = \hat{\mathbf{x}}_c \oplus \mathbf{z}_d$; $Next = e$
7. Node *d* sends current estimate back to the querying node *Q*.
8. Node *e* computes: $\hat{\mathbf{x}}_e = \hat{\mathbf{x}}_d \oplus \mathbf{z}_e$; $Next = f$
9. Node *f* computes: $\hat{\mathbf{x}}_f = \hat{\mathbf{x}}_e \oplus \mathbf{z}_f$; $Next = ...$
10. Node *f* sends current estimate back to querying node. …

As the above tracking scenario illustrates, sensor selection is a local decision. The decision must be based on a measure of information utility and cost, which can be locally evaluated and updated. The following section will overview the information-driven sensor querying (IDSQ) approach.

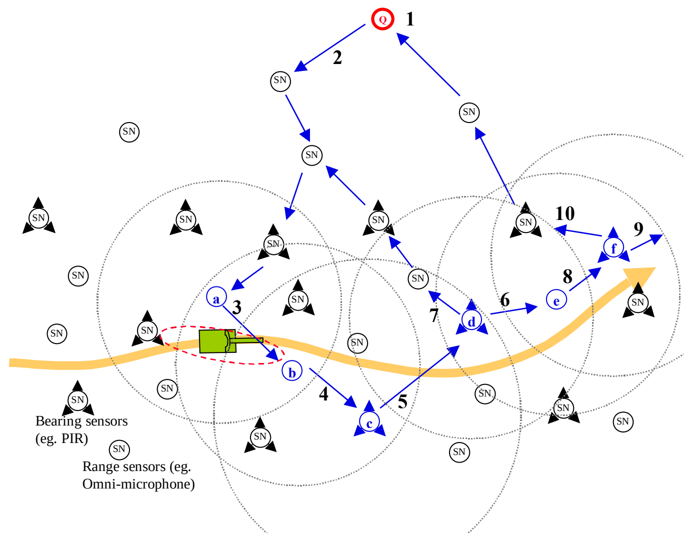


Figure 1. A tracking scenario illustrating how the decision of sensor collaboration is accomplished using a measure of information utility as well as a measure of cost. Here, a vehicle moves through the sensor field from left to right. A user query is initially routed from node *Q* to node *a* which performs an initial estimate of the vehicle position. The node *a* then selects the next sensor *b* which it believes will provide the best measurement for the next estimation at a reasonable cost, and hands the current estimate to *b*. … This process of sensor-to-sensor hand-off continues as the vehicle moves through the field. Periodically, the state estimation is sent back to the user using a shortest path routing algorithm such as directed diffusion routing.

## III. IDSQ: INFORMATION DRIVEN SENSOR QUERYING AND DATA ROUTING

We formulate the problem of distributed tracking as a sequential Bayesian estimation problem. Assuming that the state of a target we wish to estimate is $\mathbf{x}$.[2] Each new sensor measurement $\mathbf{z}_j$ is combined with the current estimate $p(\mathbf{x} \mid \mathbf{z}_1, \cdots, \mathbf{z}_{j-1})$, hereafter called belief state, to form a new belief state about the target being tracked $p(\mathbf{x} \mid \mathbf{z}_1, \cdots, \mathbf{z}_{j-1}, \mathbf{z}_j)$. The problem of selecting a sensor *j* that is likely to provide greatest improvement to the estimation at the lowest cost becomes an optimization problem. The objective function for this optimization problem can be defined as a mixture of both information gain and cost:

$$M(p(\mathbf{x} \mid \mathbf{z}_1, \cdots, \mathbf{z}_j)) = \alpha \cdot \varphi_{Utility}(p(\mathbf{x} \mid \mathbf{z}_1, \cdots, \mathbf{z}_j)) - (1-\alpha)\ \varphi_{Cost}(\mathbf{z}_j)$$

Where $\phi_{Utility}$ is the information utility measure, $\phi_{Cost}$ is the cost of communication and other resources, and α is the relative weighting of the utility and cost. We will refer to sensor *l*, which holds the current belief, as the leader node. This node might act as a relay station to the user, in which case the belief resides at this node for an extended time interval, and all information has to travel to this leader. In another scenario (such as in Fig.1), the belief itself travels through the network,

[2] We will use a superscript to denote the time stamp of the belief $\mathbf{x}^{(t)}$. This is not to be confused with the use of a subscript to index sensors; for example, $\mathbf{x}_j$ is the location of sensor *j*.

and leadership is transferred from node to node through the network. Depending on the network architecture and the tracking task, either of these cases or a combination thereof can be implemented. The detailed mathematical derivation, along with algorithms for implementing the information-driven sensor querying and data routing, can be found in [Chu et al. 2002]. In this article we will attempt to provide a more comprehensive overview of the approach and situate it in the context of the tracking applications.

The first term in the objective function $M(\cdot)$ characterizes the usefulness of the data provided by the sensor $j$. For example, when the sensor (e.g., a microphone measuring acoustic amplitude) provides a range constraint, the usefulness of the sensor data can be measured by how close the sensor is to the mean of the belief state under the Mahalanobis metric. We will return to this in some detail when we describe different criteria for sensor selection later. The second term measures the cost of obtaining the information, characterized by link bandwidth, transmission latency, node battery power reserve, etc. In the case of a moving leader node, this is the cost of handing the current belief state off to sensor $j$, acquiring data at sensor $j$, and combining the data with the current belief. In the case of a stationary leader node, this is the cost of requesting data from sensor $j$, acquiring the data and returning it to the leader to be incorporated into the belief state. In this case, the communication cost may be a function of the distance between sensor $l$ and sensor $j$, as a crude measure of the amount of energy required to transmit the data from sensor $j$ to sensor $l$. For example, with Mahalanobis distance as an information utility measure and Euclidean distance as an energy cost measure, the objective function becomes:

$$M(\mathbf{x}_j) = -\alpha \left(\mathbf{x}_j - \hat{\mathbf{x}}_T\right)^T \hat{\mathbf{\Sigma}}^{-1} \left(\mathbf{x}_j - \hat{\mathbf{x}}_T\right) - (1-\alpha) \left(\mathbf{x}_j - \mathbf{x}_l\right)^T \left(\mathbf{x}_j - \mathbf{x}_l\right)$$

where $\hat{\mathbf{x}}_T$, $\hat{\mathbf{\Sigma}}$, $\mathbf{x}_j$, $\mathbf{x}_l$ are the mean of the target position, its covariance, the position of queried sensor, and the position of querying sensor, respectively.

An example of using this objective function to query sensors and route data for a localization problem is illustrated in Fig. 2. The task here is to determine which sensors have the most useful information and ship the information back to a fixed querying node, denoted by '?' in the figure. It is important to note that incremental belief update during the routing dynamically changes both the shape and the offset of the objective function according to the updated values of $\hat{\mathbf{x}}_T$ and $\hat{\mathbf{\Sigma}}$ at every node along the routing path. As the updated values of $\hat{\mathbf{x}}_T$ and $\hat{\mathbf{\Sigma}}$ are passed on to the next node, all routing decisions are still made locally. The plotted objective function in the figure represents a snapshot of the objective function that an active routing node locally evaluates at a given time step.

## IV. Sensor selection

Given the current belief state, we wish to incrementally update the belief by incorporating measurements of other nearby sensors. However, among all available sensors in the network, not all provide useful information that improves the estimate. Furthermore, some information might be useful, but redundant. The task is to select an optimal subset and to decide on an optimal order of how to incorporate these measurements into our belief update.

It has to be emphasized that, due to the distributed nature of the sensor network, this selection has to be done without explicit knowledge of the measurement residing at each individual sensor, in order to avoid communicating less useful information. Hence, the decision has to be made solely based upon the sensor characteristics such as the sensor position or sensing modality, and the predicted contribution of these sensors.

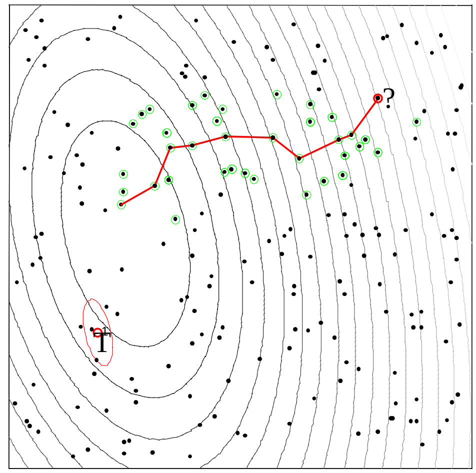


Figure 2. Sensor querying and data routing by optimizing an objective function of information gain and communication cost, whose iso-contours are shown as the set of concentric ellipses. This figure illustrates how a user query on the location of the target is being routed towards the maximum of the objective function — the center of the concentric ellipses — along the routing path. The circled dots are the sensors being queried for data along the path. 'T' represents the target position, with its covariance shown as a small red ellipse. '?' denotes the position of the query origin.

Fig. 3 illustrates the basic idea of optimal sensor selection. The illustration is based upon the assumption that estimation uncertainty can be effectively approximated by a Gaussian distribution, illustrated by uncertainty ellipsoids in the state space. In the figure, the solid ellipsoid indicates the belief state at time $t$, and the dashed ellipsoids are the incrementally updated belief after incorporating an additional measurement from a sensor, S1 or S2, at the next time step. Although in both cases, S1 and S2, the area of high uncertainty is reduced by the same amount, the residual uncertainty in the case of S2 maintains the longest principal axis of the distribution. If we were to decide between the two sensors, we might favor case S1 over case S2, based upon the underlying measurement task.

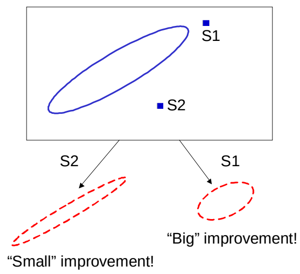


Figure 3. Sensor selection based on information gain of individual sensor contributions. The information gain is measured by the reduction in the error ellipsoid. In the figure, reduction along the longest axis of the error ellipsoid produces a larger improvement in reducing uncertainty. Sensor placement geometry and sensing modality can be used to determine the potential information gain to be provided by the two sensors S1 and S2.

Although details of the implementation depend on the network architecture, the fundamental principles introduced in this article hold for both, the selection of a remote sensor by a leader node (e.g., a cluster head), as well as the decision of an individual sensor to contribute its data and to respond to a query traveling through the network. The task is to select the sensor that provides the best information among all available sensors that have not been incorporated. As has been shown in [Chu et al. 2002], this provides a faster reduction in estimation uncertainty, and usually incurs a lower communication overhead for meeting a given estimation error requirement, compared to blind or nearest-neighbor sensor selection schemes.

## V. Information Utility

Information utility measures play a key role in the information-driven approach to sensor selection. In this section, we first introduce an information-theoretic definition of the utility measure. We then describe several heuristic approximations to the measure that prove to be practically useful.

Our goal is to *predict* the information utility of a piece of non-local sensor data before obtaining the data. In practice, the prediction must be based on the currently available information: the current belief state, the characteristics of the sensor of interest which includes information such as the sensor position and sensing modality that can be established beforehand.

We assume there are $N$ sensors labeled from 1 to $N$ and the corresponding measurements of the sensors are $\mathbf{z}_1, \cdots, \mathbf{z}_N$. Let $U \subset \{1, \cdots, N\}$ be the set of sensors whose measurements have been incorporated into the belief. That is, the current belief is $p(\mathbf{x} \mid \{\mathbf{z}_i\}_{i \in U})$. The sensor selection task is to choose a sensor whose data has not been incorporated into the belief yet and which provides the most information[3]. To be precise, let us define an information utility function $\varphi_{Utility}$ that assigns a value to each probability distribution.

[3] For a moving target, each sensor can be reused without loss of generality.

For the moment, we ignore the cost term in the objective function. The best sensor, defined by the earlier objective function, is given by

$$\hat{j} = \text{argmax}_{j \in V} \varphi_{Utility}(p(\mathbf{x} \mid \{\mathbf{z}_i\}_{i \in U} \cup \{\mathbf{z}_j\}))$$

where $V$ is the set of sensors whose measurements are potentially useful.

The following are possible definitions of the information utility function.

### A. Information-theoretic measure: entropy

The information utility function $\varphi(\cdot)$ evaluates the compactness of the belief state distribution. The natural choice of $\varphi(\cdot)$ is the statistical entropy, which measures the randomness of a given random variable. Mathematically, the entropy is defined as

$$H_p(x) = -\sum_{x \in S} p(x) \log p(x)$$

for a discrete random variable *x*. The equivalent mathematical quantity for a continuous random variable is

$$H_p(x) = -\int_S p(x) \log p(x) dx$$

In both of the entropy definitions, $S$ denotes the support of the random variable. Generally speaking, the smaller the entropy is, the more certain we are about the value of the random variable. Therefore, the information utility measure based on the entropy can be defined as

$$\varphi(p(\mathbf{x} \mid \{\mathbf{z}_i\}_{i \in U} \cup \{\mathbf{z}_j\})) \triangleq -H_p(x).$$

The entropy-based definition, while mathematically precise, is difficult to compute in practice since we need to have the measurement before deciding how useful the measurement is. A more practical alternative is to estimate the usefulness of a measurement based only on characteristics of a sensor such as its location or sensing modality.

### B. Mahalanobis distance measure

The idea of using Mahalanobis distance as a measure can be illustrated using Fig. 3. In the figure, the solid squares labeled S1 and S2 are sensors whose measurements can potentially improve the current belief state. Suppose each sensor's measurement provides a range constraint. A new belief state, whose uncertainty is shown as dashed ellipse, is computed by combining the measurement with the current belief state, using Bayes rule. In the sensor configuration shown, S1 would provide better information than $S2$ because S1 lies close to the longer axis of the uncertainty ellipse and S1's range constraint would more perpendicularly intersect this longer axis. To favor the sensors along the longer axes of an uncertainty ellipsoid, we use Mahalanobis distance, a distance measure normalized by the uncertainty covariance. Therefore, when the current belief can be well approximated by a Gaussian distribution, the utility function is

$$\varphi(\mathbf{x}_j, \hat{\mathbf{x}}, \hat{\mathbf{\Sigma}}) = -(\mathbf{x}_j - \hat{\mathbf{x}})^T \hat{\mathbf{\Sigma}}^{-1} (\mathbf{x}_j - \hat{\mathbf{x}})$$

where $\mathbf{x}_j$ is the position of sensor $j$, and $\hat{\mathbf{x}}$ is the mean of the belief (target position estimate).

The Mahalanobis distance-based utility measure works well when the current belief can be well approximated by a Gaussian distribution or the distribution is very elongated, and the sensors are range sensors. However, a bearing sensor reduces the uncertainty along the direction perpendicular to the target bearing. For a general uncertainty distribution or bearing sensors, we must develop alternative information utility measures.

### C. *Measures on expected posterior distribution*

The idea of using expected posterior distribution is to predict what the new belief state (posterior distribution) would look like if a simulated measurement of a sensor from the current belief state is incorporated. The utility of each sensor can then be quantified by the entropy or other measures on the new distribution from the simulated measurement.

We use the tracking problem to derive an algorithm for evaluating the expected utility of a sensor. In the ideal case when a real new measurement is available, the new belief or posterior is evaluated using the familiar sequential Bayesian filtering:

$$p(\mathbf{x}^{(t+1)} \mid \overline{\mathbf{z}^{(t+1)}}) = C \cdot p(\mathbf{z}_j^{(t+1)} \mid \mathbf{x}^{(t+1)}) \cdot \int p(\mathbf{x}^{(t+1)} \mid \mathbf{x}^{(t)}) \cdot p(\mathbf{x}^{(t)} \mid \overline{\mathbf{z}^{(t)}}) d\mathbf{x}^{(t)}$$

where $p(\mathbf{x}^{(t)} \mid \overline{\mathbf{z}^{(t)}})$ is the current belief given a history of the measurement up to time $t$: $\overline{\mathbf{z}^{(t)}} = \{\mathbf{z}^{(0)}, \cdots, \mathbf{z}^{(t)}\}$, $p(\mathbf{x}^{(t+1)} \mid \mathbf{x}^{(t)})$ specifies the predefined dynamics model, and $p(\mathbf{z}_j^{(t+1)} \mid \mathbf{x}^{(t+1)})$ is the likelihood function from the measurement of sensor $j$.

How do we compute the expected value of $p(\mathbf{x}^{(t+1)} \mid \overline{\mathbf{z}^{(t+1)}})$ without having the data $\mathbf{z}_j^{(t+1)}$ in the first place? The idea is to *guess* the shape of likelihood function from the current belief and the sensor position.

Without loss of generality, the current belief is represented by a discrete set of samples on a grid of the state space. This non-parametric representation of the belief state allows to represent highly non-Gaussian distribution and nonlinear dynamics. Figure 4 shows an example of the grid-based state representation. The gray squares represent the likely position of the target as specified by the current belief. The brighter the square, the more likely the target is there. For a sensor $i$, given the observation model $\mathbf{z}_i^{(t+1)} = h(\mathbf{x}^{(t+1)}, \mathbf{w}_i^{(t)})$, where $\mathbf{w}_i^{(t)}$ is the sensor noise, we can estimate the measurement $\mathbf{z}_i^{(t+1)}$ from the predicted belief and compute the expected likelihood function, that is,

$$\hat{p}(\mathbf{z}_i^{(t+1)} \mid \mathbf{x}^{(t+1)}) = \sum_{v_k \in S(\mathbf{x}^{(t+1)})} L_{ki}(\mathbf{x}^{(t+1)}, v_k) \cdot \left[ p(\mathbf{x}^{(t+1)} \mid \overline{\mathbf{z}^{(t)}})\Big|_{\mathbf{x}^{(t+1)}=v_k} \right]$$

where the marginal likelihood is defined as

$$L_{ki}(\mathbf{x}^{(t+1)}, v_k) \triangleq \hat{p}(\mathbf{z}_i^{(t+1)}(\mathbf{x}^{(t+1)} = v_k) \mid \mathbf{x}^{(t+1)})$$

and the prediction as

$$p(\mathbf{x}^{(t+1)} \mid \overline{\mathbf{z}^{(t)}}) \triangleq \sum_{u_k \in S(\mathbf{x}^{(t)})} \left[ p(\mathbf{x}^{(t+1)} \mid \mathbf{x}^{(t)})\Big|_{\mathbf{x}^{(t)}=u_k} \right] \cdot \left[ p(\mathbf{x}^{(t)} \mid \overline{\mathbf{z}^{(t)}})\Big|_{\mathbf{x}^{(t)}=u_k} \right]$$

Using the estimated likelihood function $\hat{p}(\mathbf{z}_i^{(t+1)} \mid \mathbf{x}^{(t+1)})$ from sensor $i$, the expected posterior belief can be obtained as follows

$$\hat{p}(\mathbf{x}^{(t+1)} \mid \overline{\mathbf{z}^{(t+1)}}) = C \cdot \hat{p}(\mathbf{z}_i^{(t+1)} \mid \mathbf{x}^{(t+1)}) \cdot p(\mathbf{x}^{(t+1)} \mid \overline{\mathbf{z}^{(t)}})$$

We can then apply measures such as the entropy to the expected belief $\hat{p}(\mathbf{x}^{(t+1)} \mid \overline{\mathbf{z}^{(t+1)}})$, as an approximation to the true belief $p(\mathbf{x}^{(t+1)} \mid \overline{\mathbf{z}^{(t+1)}})$. This approach can apply to non-Gaussian belief since the discrete approximation of the belief state assumes a general form. However, to compute the expected belief, we have conditioned the expected likelihood function on the predicted belief state. We will examine some of the consequences of this bias in the discussion section.

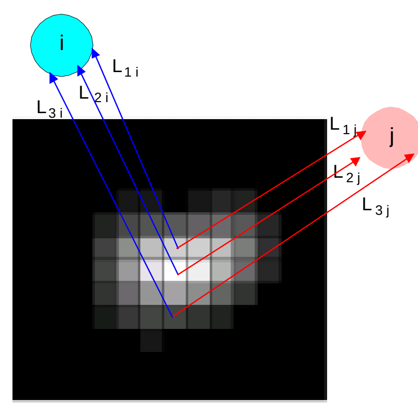


Fig. 4. The expected likelihood function for each sensor ($i$ or $j$) is a weighted sum of the marginal likelihood function conditioned at each grid point in the predicted belief distribution. The expected posterior can then be computed from this likelihood function.

## VI. Experimental Result

In this section, we present computational results from applying the information utility measures introduced in the last section to target localization and tracking problems.

*A. Localizing a stationary target*

We compare the information-driven sensor selection with blind nearest neighborhood selection in the context of localizing a stationary target. Figure 5 shows two snapshots of the tracking algorithm based on the nearest neighborhood criterion (NN). Fig. 5(a) shows the posterior distribution after combining the data from the sensor at the middle of the linear array with the data from its two nearest neighbors. Using the NN criterion, the next best sensor is one of their nearest neighbors in the linear array. The new posterior distribution remains as a bimodal distribution (Fig. 5(b)) till the sensor at the upper-left corner of the sensor field is selected.

In Fig. 6, the sensor selection is based on the Mahalanobis distance measure. Fig. 6(a) shows the posterior after combining the measurements from the same three sensors near the middle as in Fig. 5(a). The residual uncertainty, however, is elongated and thus the upper-left sensor is selected as the next sensor according to the Mahalanobis distance. The new measurement from that sensor effectively reduces the current uncertainty to a more compact region (Fig. 6(b); also compare Fig. 6(b) with Fig. 5(b)).

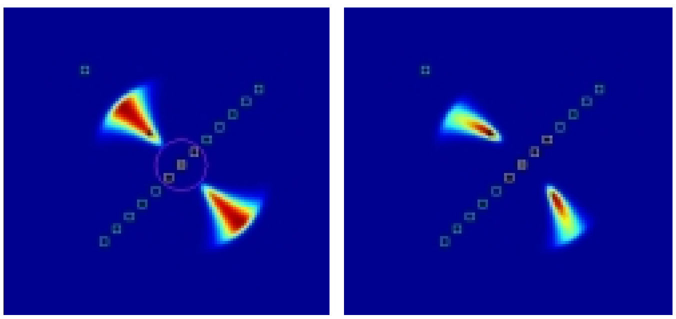

(a) (b)

Fig. 5. Sensor selection based on the nearest neighbor method. The estimation task here is to localize a stationary target labeled '∗'. Squares denote sensors. (a) Select the nearest sensor ; (b) Incorporate the new measurement from the selected sensor.

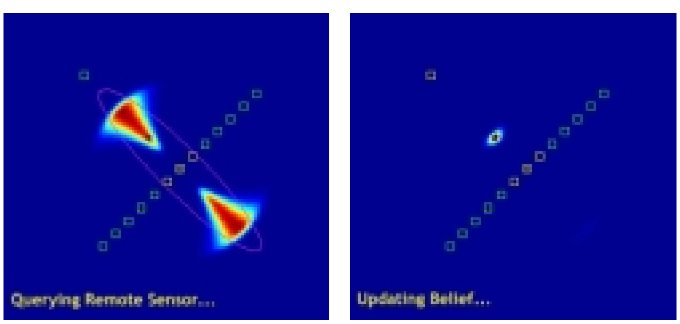


(a) (b)

Fig. 6. Sensor selection based on the Mahalanobis distance measure of information utility. The localization problem is the same as that in Fig. 5.

*B. Tracking a moving target with non-Gaussian distribution*

We now illustrate how the information utility measures can be applied to a tracking problem. In our study, we assume a leader node (the square-enclosed dot in Fig. 7) carries the current belief state. The leader chooses a sensor with good information in its neighborhood according to the information measure, and then hands off the current belief to the chosen sensor (the new leader).

As discussed earlier, the information-based approach to sensor querying and data routing selectively invokes sensors to minimize the number of sensing actions needed for a given accuracy and hence, latency and energy usage. It can optimize the use of multi-sensing-modality information (e.g., range and bearing sensing) to improve tracking accuracy. It can also handle non-constant target dynamics and is more general than the CPA-based method.

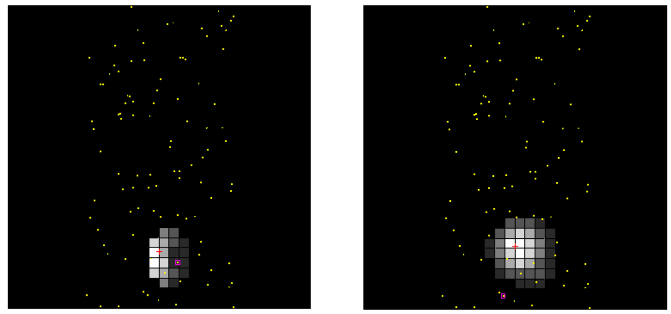

(a) (b)

Fig. 7. Tracking a moving target using an information-driven approach. A target is moving from the bottom of the field to the top. As the target moves through the field of sensors denoted by the dots, a subset of sensors are activated to carry the belief state. Each new sensor is selected according to an information utility measure on the expected posterior distribution of the new state. (a) Current belief distribution at time *t*. (b) New posterior distribution at time *t+1*, after incorporating a measurement from the selected sensor.

Fig. 8 shows the performance of various trackers based on four different sensor selection criteria: nearest neighborhood (a)-(b), minimizing Mahalanobis distance (c)-(d), minimizing entropy (e)-(f), and minimizing relative entropy (or maximizing Kullback-Liebler-distance (KL-distance) [Thomas&Cover 1991] between predicted belief and expected belief) (g)-(h). For each criterion, we also examine the effect of a simple heuristic that prevents each node from being selected more than *N* times. Since the target position estimation provided by a single range sensor is under-constrained, it is desirable that the same sensor not be selected over and over again in order to maintain a certain amount of spatial diversity. In practice, the repeated use of information from a single sensor may also lead to over confidence in the final estimation due to correlations between consecutive measurements from the same sensor.

The figures in the left column of Fig. 8 are the results without using the heuristic, and the ones in the right column with the heuristic. These empirical results indicate that two of the entropy-based utilities without the heuristic outperform the others by the huge margin of errors errors and confirm that the expected utility measure is a reasonable choice as information utility. On the other hand, it is observed that the Mahalanobis distance and the Nearest Neighborhood select the same set of sensors repeatedly and forces the target position estimation to be biased as a result. The heuristic, in this case, helps by preventing the same set of sensors from being selected repeatedly while it actually hurts in the case of the entropy-based measure. The larger error for the low sensor density in the Mahalanobis and the Nearest Neighborhood utility is because there are not enough sensors to work when the heuristic is used.

Figure 9 shows a snapshot of a simulation run using the relative entropy as the utility measure.

Table 1 summarizes the statistics from the simulation results shown in Fig. 8. A track is defined as lost when the final estimated position of the target is more than 15 meters away from the actual position. Fig. 10 graphs the estimated tracks for each choice of the information utilities, ordered in the same way as in Table 1, for sample runs with sensor density 60.

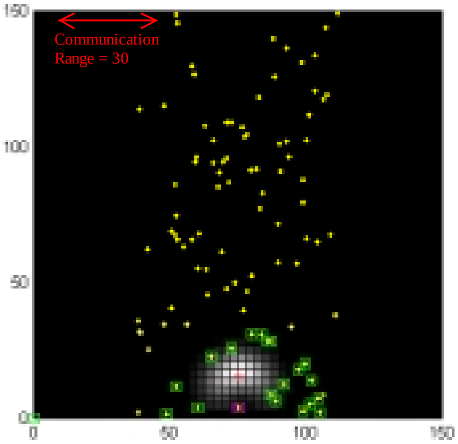


Figure 9. Snapshot of a typical simulation using the relative entropy measure to select sensors. The yellow dots are sensor nodes and the gray grids are current estimates on the target position. The pink-square-enclosed node is the current leader, and the green-square-enclosed nodes are its current neighbors. In this example, sensor density is 100, and the node communication range is set to30.

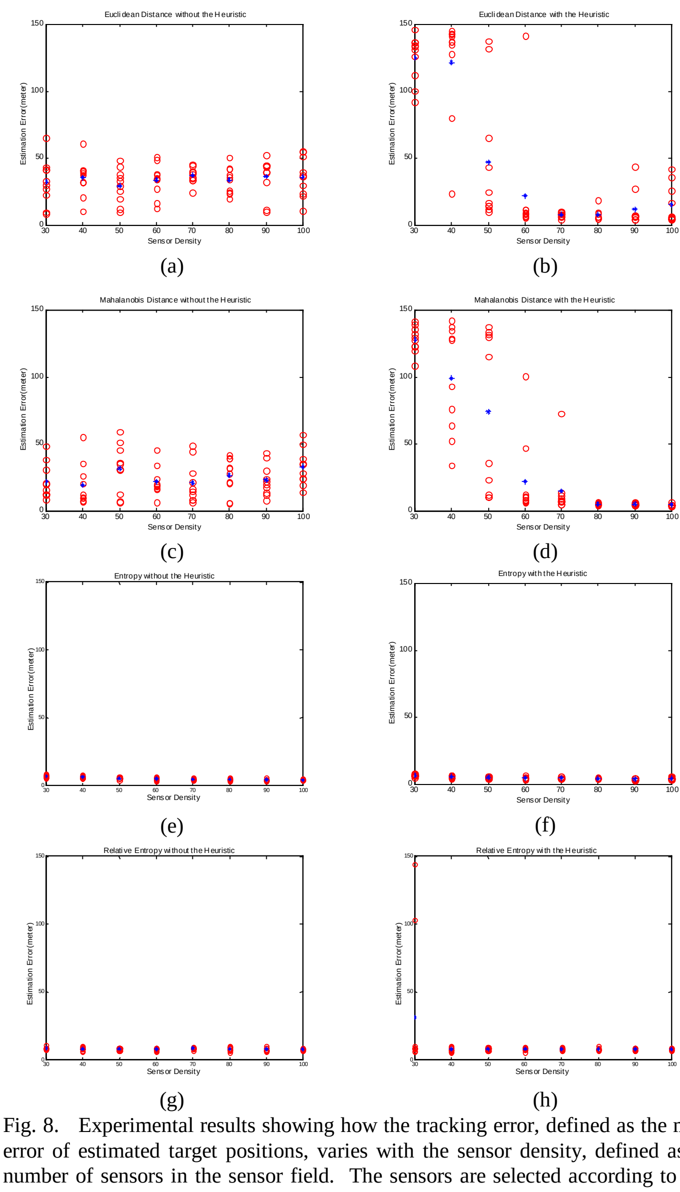


Fig. 8. Experimental results showing how the tracking error, defined as the mean error of estimated target positions, varies with the sensor density, defined as the number of sensors in the sensor field. The sensors are selected according to four different criteria, with or without using the max activation heuristic: (a) Nearest neighborhood. (b) Nearest neighborhood with the heuristic. (c) Mahalanobis distance. (d) Mahalanobis distance with the heuristic. (e) Entropy. (f) Entropy with the heuristic. (g) Relative entropy (KL-distance). (h) Relative entropy with the heuristic.

| | # Lost tracks/ total runs | Mean error |
|---|---|---|
| (a) Nearest neighbor | 75/80 | 34.39 |
| (b) Nearest neighbor with heuristic | 37/80 | 44.79 |
| (c) Mahalanobis distance | 70/80 | 24.86 |
| (d) Mahalanobis distance with heuristic | 32/80 | 44.20 |
| (e) Entropy | 0/80 | 5.13 |
| (f) Entropy with heuristic | 0/80 | 5.05 |
| (g) Relative entropy | 0/80 | 8.09 |
| (h) Relative entropy with heuristic | 2/80 | 10.79 |

Table 1. Statistics on tracker performance for different choices of the utility functions

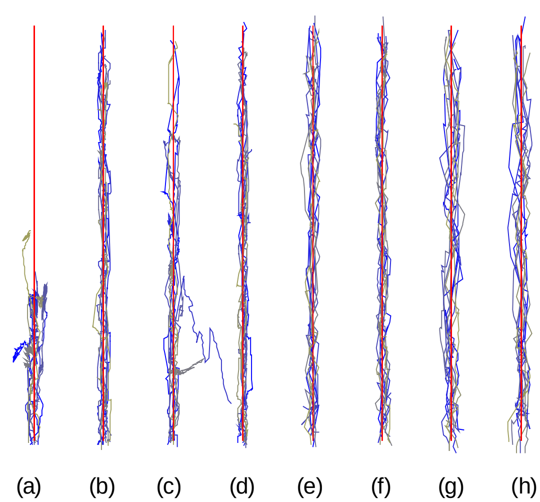


Figure 10. Tracking results for different choices of the utility functions. The red lines are the actual tracks, and the blue-gray curves are estimated tracks.

## VII. Discussions

### A. Representation of belief state

In our tracking example, the belief state is being passed from one leader to another. To efficiently implement the IDSQ tracking algorithm, it is important to design a compact representation of beliefs so as to minimize the communication requirement.

While the parametric representation is the most compact form, it is limited to specific classes of distributions that can be modeled this way. At the other extreme, one can approximate an arbitrary distribution with a set of discrete samples. This forms the basis of many successful Monte Carlo based algorithms such as sequential Monte Carlo or Markov Chain Monte Carlo methods [Doucet et al. 2000]. One drawback of the non-parametric approaches is the large numbers of samples that are needed to represent the state space, and addressing this problem is a very active topic of current research. Somewhere between the parametric and particle sample based approaches lies the grid-based representation of beliefs. In this representation, each grid approximates the value of the belief at the grid location. When the state space is sparse, meaning many parts of the space have negligible probability masses, the grid can be efficiently encoded in a sparse representation to minimize storage requirement. Fig. 11 provides a pictorial description of the three representations.

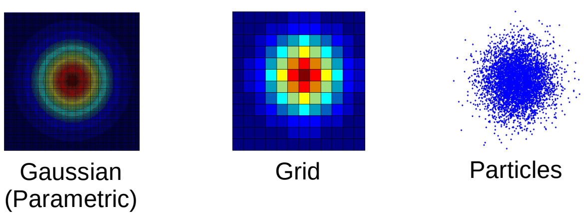


Figure 11. Representation of belief: parametric (e.g., Gaussian), non-parametric (e.g., grid samples, particle samples).

### B. Sequential vs. concurrent information exchange

In our discussion thus far, we have primarily focused on the case where a leader node selects the next leader to hand off the belief state information. While the idea is simple, this single node-to-node handoff may suffer information loss when the current leader/link incurs a failure. A more robust scheme uses a zone-to-zone hand-off where a group of sensors (in a zone) elect a leader and collectively maintain the belief state. The leader performs the normal node-to-node handoff to the next chosen leader of a cluster of sensors. But when the leader node fails, another sensor in the same cluster will step in as an acting leader, and the handoff continues. Pushing this idea to the extreme, every sensor may exchange information with another sensor in parallel, perhaps at the cost of consuming greater overall communication bandwidth. There are several possible types of information to exchange. Each sensor can send its own belief to a chosen next sensor, it can send its measurement to a neighbor sensor to whom the data is likely to be useful, or it can request other’s belief or measurement.

The selection of which style of information exchange to deploy for a sensor net depends on how the information will be extracted and used (e.g. query types) and the level of desired robustness to failure and tolerance for latency.

### C. Query types

In our straw-man tracking scenario (Fig. 1), we assumed that a user initiates a query about the location of a vehicle as a function of time. In other use cases, one might expect the network to possess some low-level awareness of the targets. When an event was detected and classified, the network can initiate a tracking task. As the network begins to activate sensors in a local region to maintain active belief states, the information about the current active regions and their event logs can help to guide user queries, which may enter the network at any point of the network, into the region of high information relevance.

Beyond the single target tracking scenario, one might be interested in tracking a group of targets or relations among the targets [Guibas 2002]. In these cases, a user initiated high-level query may, upon entering the network, break into several sub-queries, some of which are routed into regions where individual targets are present, and some of which coordinate the routing or interpretation of the first group of sub-queries. Complexities arise as targets merge, split, or cross-over. Sub-queries may have to reconcile with each other periodically to maintain consistency.

The taxonomy of query types, the nature of physical phenomenon being observed, and their corresponding in-network processing styles are future topics of investigation.

### D. Bias in sensor selection

The information utility measures we have introduced are approximate in nature. Mahalanobis distance is a heuristic for measuring the utility of range sensing data. In the presence of the finite-precision representation of probability distribution and possibly nonlinear utility functions, the sensor selection based on the expected posterior computed from predicted likelihood function may be strongly biased by the prior distribution. As the experimental results have shown, some of the cases where the track is lost are actually due to this bias. A poorly approximated prior may produce incorrect utility values, leading to the selection of less useful sensors eventually causing the tracking error to explode.

Similar to the sequential Monte Carlo method (also known as particle filters) [Doucet et al. 2000], one way to reduce the effect of the poorly approximated priors on sensor selection is to design proposal distributions that draw on multiple information sources.

### E. Tracking robustness

The quality of tracking is a complex function of several parameters: sensor placement density, sensing range, communication range, spatial extent of the physical phenomenon being observed, target dynamics, S/N ratio. This paper has skimmed over the important issue of quantifying track quality. Another critical issue for robustness is handshake during information handoff from node to node. We expect future research will seek to understand the effect of these parameters on the behaviors of trackers and design robust communication protocols.

### F. Related work

The idea of using information utility to manage sensing resources has been investigated in computer vision and robotics, e.g. active vision or active testing [Geman96]. Most of the approaches assume the selection is done centrally. [Manyika & Drrant-Whyte 1994] introduces expected utility measures for decentralized sensing systems based on local decision at each node; but the communication cost is not explicitly considered during the optimization. Likewise, [Byers 2000] uses a simple step or sigmoid function to describe utilities of each node, without explicit modeling of network spatial configuration. IDSQ generalizes these approaches to consider both network spatial configuration and communication cost in sensor selection.

The issues of data storage, retrieval, and naming in sensor networks have been studied by the wireless networking, decentralized database, and distributed tracking communities. Directed data diffusion (e.g., [Intanagonwiwat et al. 2000]) uses a publish-subscribe mechanism to name data and pair data sources with data sinks, exploiting the network topologies. [Brooks et al. 2002] also uses data to guide sensor communication; in this case, prediction from tracking history is used to invoke sensors for future processing. IDSQ builds on directed diffusion and allows local nodes to make routing decisions based on information gain and resource cost.

## VIII. Conclusion

We have formulated the problem of distributed tracking using wirelessly connected sensors as an information optimization problem and introduced several practically feasible measures of information utility. The information-driven approach to sensor querying and data routing balances the information gain provided by each sensor with the cost associated with acquiring the information. This way, the sensor network can seek to make an informed decision about sensing and communication in an energy constrained environment.

## IX. Acknowledgement

The authors thank Patrick Cheung, Julia Liu, and Leo Guibas for helpful discussions on trackers. This work is supported in part by the Defense Advanced Research Projects Agency (DARPA) under contract number F30602-00-C-0139 through the Sensor Information Technology Program. The views and conclusions contained herein are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the US Government.

## X. References

[Brooks et al. 2002] R.R. Brooks, C. Griffin, and D.S. Friedlander, "Self-organized distributed sensor network entity tracking", *Int'l J. of High-Performance Computing Applications*, to appear, 2002.

[Byers&Nasser 2000] J. Byers and G. Nasser, "Utility-Based Decision-Making in Wireless Sensor Networks (Extended Abstract)", Proceedings of IEEE MobiHOC 2000, Boston, MA, August 2000.

[Chu et al. 2002] M. Chu, H. Haussecker, and F. Zhao. "Scalable Information-Driven Sensor Querying and Routing for ad hoc Heterogeneous Sensor Netoworks". *Int'l J. of High-Performance Computing Applications*, to appear, 2002.

[Doucet et al. 2000] A. Doucet, J.F.G. de Freitas, and N.J. Gordon. *Sequential Monte-Carlo Methods in Practice*. Springer-Verlag, 2000.

[Estrin et al. 1999] D. Estrin, R. Govindan, J. Heidemann, S. Kumar, "Next Century Challenges: scalable coordination in sensor networks". In Proceedings of the Fifth Annual International Conference on Mobile Computing and Networks (MobiCOM '99), Seattle, Washington, August 1999.

[Geman&Jedynak 1996] D. Geman and B. Jedynak, "An active testing model for tracking roads from satellite images", *IEEE Trans. Pattern Anal. Mach. Intell.*, 18, 1-14, 1996.

[Guibas 2002] L. Guibas, "Sensing, Tracking, and Reasoning with Relations". This issue.

[Intanagonwiwat et al. 2000] C. Intanagonwiwat, R. Govindan, D. Estrin, "Directed diffusion: a scalable and robust communication paradigm for sensor networks". In Proceedings of the Sixth Annual International Conference on Mobile Computing and Networks (MobiCOM 2000), Boston, Massachusetts, August 2000.

[Manyika&Durrant-Whyte 1994] J. Manyika and H. Durrant-Whyte, *Data fusion and sensor management: a decentralized information-theoretic approach*. Ellis Horwood, 1994.

[Thomas&Cover 1991] T. M. Cover and J. A. Thomas, "*Elements of Information Theory*", Wiley, 1991