

Distributed Adaptive Sampling Using Bounded-Errors

Kévin Huguenin
ENS Cachan / Lab. I3S, CNRS-UNSA
Sophia Antipolis, France
huguenin@i3s.unice.fr

M. João Rendas
Lab. I3S, CNRS-UNSA
Sophia Antipolis, France
rendas@i3s.unice.fr

Abstract—This paper presents a communication/coordination/processing architecture for distributed adaptive observation of a spatial field using a fleet of autonomous mobile sensors. One of the key difficulties in this context is to design scalable algorithms for incremental fusion of information across platforms robust to what is known as the “rumor problem”. Incremental fusion is in general based on a Bayesian approach, and algorithms (e.g. the Covariance Intersection, CI) which propagate consistent characterizations of the estimation error under this challenging situation have been proposed. In this paper, we propose to base inter-sensor fusion on a deterministic approach which considers that bounds to the observation errors are known, which is intrinsically robust to the rumor problem. We present the equations that enable the determination of the ellipsoidal domain of uncertainty that covers the intersection of the individual sets describing sensor’s uncertainty, and show that they solve some pathologies associated to CI. The results presented corroborate a previous claim of the robustness of our control strategy (the criterion used for adaptively choosing the nodes positions) with respect to the conservativeness of fusion methods able to handle rumor .

I. INTRODUCTION

Over the last few years, sensor networks have been progressively recognized as providing an efficient solution for real-world applications in the context of environmental monitoring, enabling rapid and focused observation of spatially extended fields such as temperature or humidity inside a given region of interest. Recent advances in the areas of wireless communications and robotic sensors have made the design of autonomous mobile distributed sensing systems possible. An example of such networks are fleets of underwater vehicles with embedded sensors, micro-processors and wireless communication devices. They enable more powerful and autonomous applications in the sense that, contrary to fixed sensors, this kind of networks do not rely on preliminary computation for sensor placement, allowing proper observation of dynamic environments, and adapting the network geometry to the characteristics of the measured field. Also, increasing computational power available at each node makes decentralized and distributed path control and data fusion possible. Finally, wireless communications allow collaborative behavior and dynamic data exchange making sampling more efficient.

So far, most of the research on sensor networks has focused on *sensor placement* and *data processing* for *static* sensor networks with one or more *data fusion centers* [1], [2]. Dealing

with autonomous *mobile* sensors and designing *distributed* systems brings new challenges. Due to the dynamic nature of the network, full-connectivity can no longer be guaranteed, making traditional data collection/dissemination schemes and centralized path control infeasible. To date there has been little work on adaptive collaborative sensing using mobile sensors. Rahimini *et al.* propose an algorithm for adaptive sampling using Networked Info-Mechanical System (NIMS) [3] but they restrict to the independent control of each sensor and do not take into account energy and communication constraints since NIMS are power supplied and use a wired communication medium. In the framework of underwater applications, the concept of adaptive sampling using a fleet of autonomous sensors has been advocated for some years by the members of the ASAP (Autonomous Sampling and Prediction) project, see for instance [4]. However, this project considers a centralized control of the underwater fleet, eliminating most of the problems and difficulties that we address in this paper.

The ultimate goal of an autonomous sensing network is to *gather data which collectively allows the estimation of a given spatial field over an area of interest \mathcal{A}* , this estimate being obtained once data gathering is finished, by globally fusing the data sets collected by the nodes. The key point of such an architecture is the collaboration between nodes to determine the sensing points where the measurements will jointly bring the most information, which requires that nodes share their knowledge about the field. In [5], [6], we presented an architecture integrating communication, control and information processing algorithms that implements a distributed and decentralized strategy for acquiring an estimate of a spatially extended field, based on a probabilistic Bayesian framework, under strong communication constraints. While iterative integration of additional sensor measurements can be easily done in the context of Bayesian estimation theory, the fusion of local estimates is prone to the *rumor* problem: there is no simple way to distinguish fresh information conveyed by a distinct node from re-circulation, through the network, of a node’s information back to itself. In these conditions a naive application of the Bayesian recursion would lead to overconfident error characterizations. This problem is overcome by accepting some degree of information loss, and using fusion methods that propagate upper bounds on uncertainty (conservative error estimates), like the Covariance Intersection

(CI) algorithm [7]. In [5] we presented a new criterion which prevents fresh measurements from being discarded by CI by choosing sensing positions that correspond to different observation structures, and we compared the performance of our distributed sensing architecture using CI and naive Bayes. Our analysis shows the superiority of CI and that our criterion is well adapted to CI.

The entire issue of “rumor propagation” is eliminated if a deterministic, error-bounded [8] approach to the fusion problem, which propagates upper bounds to the sets of parameters values compatible with the acquired observations, is used. Bounded-error methods have been popular in control and estimation problems, mostly for recursive state estimation problems, and update equations for the set dilation and intersection steps have been published. The problem of *distributed fusion* requires that two or more *uncertainty domains* be directly fused. In this paper we present equations that allow the fusion of two ellipsoids that are a special case of the general problem solved in [9]. Besides providing a better solution to the fusion problem, as we show in section VI, the use of a bounded-error based methodology for fusion, based on a completely geometric characterization of the uncertainty domain affecting the current estimates at each node, is a natural match to our adaptive sensing criterion, giving a stronger coherency to our overall architecture.

The paper is organized as follows. Section II outlines the problem of distributed collaborative sensing and presents an overview of the global architecture proposed. In Section V we briefly present and justify the adaptive sensing criterion used. Sections III and IV are dedicated to recursive estimation and fusion, focusing in the presentation of the update equations for the bounded-error methods, and pointing out their analogies to Gauss-based Bayesian methodologies. In Sections VI and VII we present simulation results and summarize our conclusions and some ideas for future work.

II. PROBLEM STATEMENT: GENERAL ARCHITECTURE

Our goal is to implement a *decentralized strategy* for controlling the *sensors positions* $X^t = \{X_{t_c}\}_{t_0 \leq t_c \leq t}$ such that the information collectively gathered by the network optimizes some aggregated statistic functional of the (error of the) reconstructed field, $\hat{f}(x; X^t)$

$$\mathcal{C}(X^t) = \mathbb{E} \left\{ \iint_{\mathcal{A}} g \left(f(x, y), \hat{f}(x, y; X^t) \right) dx dy \right\} \quad (1)$$

where \mathbb{E} denotes the (ensemble) expectation operator.

We consider a set of working assumptions, which reflect constraints typical of ocean-observation problems:

- the observed field $f(z)$ is scalar and assumed *static* during the time interval required to perform its observation inside a known region, $\mathcal{A} \subset \mathbb{R}^2$.
- the robots are able to localize themselves (GPS, network of acoustic beacons, distributed location algorithm, ...). Their speed is upper-bounded $v \leq v_{\max}$ and they can (and must, to perform a measurement) station-keep.

- error-free fixed radius communications: two nodes can communicate with each other without errors provided that the distance between them is lower than the communication range R_{com} .

We present below the major components of our distributed sensing architecture [5]. Our solution involves three main stages that are executed in a cyclic manner: (i) *sensing*, (ii) *communicating* and data *fusion*, and (iii) *collaborating* to decide on the next network deployment.

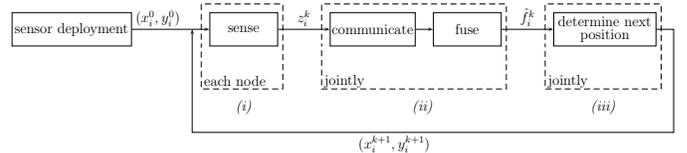


Figure 1. Coarse Block-based overview of execution cycle of the Distributed Algorithm for Collaborative Adaptive Sensing

While sensing is performed independently at each node, choosing the next sensing position can be drastically improved by sharing information across nodes. The second task is thus denoted by *fusion* insofar as it consists in fusing the estimates stored at each node to compute a more informative one while the third one is called *adaptive collaborative observation* in the sense that the nodes rely on the exchanged information to determine the set of future sensing points which jointly bring the most additional information about the observed field.

Network partitioning

Obviously, the larger the number of nodes that collaborate to select the future sensing positions, the more their individual decisions will reflect the global state of knowledge of the network. However, under low spatial node density and limited communication range, effective network connectivity may be very limited. In particular, the communication graph may be disconnected, with communication being impossible between some network nodes. Our architecture is based on the dynamic definition of *logical groups of collaborating nodes*, designed to allow easy information exchange amongst their members. Logical groups (*clusters*) are defined at the beginning of each cycle, by assigning a virtual group number to each robot and then restricting collaboration to nodes belonging to the same cluster.

Distributed clustering is a well-known problem and has been addressed under the single-hop structural constraint [10]. By imposing direct communication between cluster members, the nodes of our network are released from the communication problems and requirements characteristics of Mobile Ad-Hoc Networks (MANET) such as intermittent connectivity and delay tolerant routing. A “Distributed Clustering Algorithm” (DCA) which dynamically creates and destroys these clusters using only *gossip messages* has been proposed in [5]

Distributed data fusion

Data fusion consists in updating the estimate of the observed field at each node using either new measurements (we call this operation *update*) or estimates received from other nodes (that we designate by *fusion*). Thus, clustering must be immediately

followed by data sharing and fusion, that need to occur before the sensing positions are chosen in order to ensure consistency of this last operation.

In [6], [5] we compared two Bayesian methods for update and fusion (“naive Bayes” and “Covariance Intersection”) and demonstrated the superiority of CI. The fusion operation is the main new contribution of this paper, which proposes to assess it in a distinct uncertainty modeling framework (deterministic, bounded-error methods), see Section III.

Choice of future positions

Selection of the sensing points is the core problem of adaptive field observation. Its potential impact is evident, since without collaboration nearby nodes will have a similar local “view” of the observed field, and will most likely decide to sense at the same or nearby locations. The criterion presented in [6], [5] to choose the next set of sensing points tries to reach a balance between (i) maximizing the new information *jointly* brought by the measurements at the target sensing points; and (ii) imposing geographic dispersion of the nodes, to increase the probability they establish communication with nodes from distinct clusters.

Search of the best set of sensing points is performed within a bounded region around the nodes, called *action zone* using a *greedy distributed approach*: sensing points are determined one-by-one, taking into account the previously selected ones. Thus, the computational power required to determine the sensing points grows linearly with the number of nodes in the clusters and so does the available computational power. Once the set of the sensing positions has been chosen, each of them is assigned to one node such that the total amount of energy (this can be replaced by a different criterion) used by the robots to reach them is minimized, see [5] for details.

III. RECURSIVE ESTIMATION

As we stated in the Introduction, scalability constraints forbid storing the entire set of raw data acquired by each node. In this section we briefly present (A) the fundamental equations of Bayesian estimation theory that propagate a sufficient statistic summarizing all relevant information about the estimated field contained in the past observed data, as well as (B) a deterministic estimation method based on feasible-parameter set outer-bounding, both allowing recursive update.

We assume that the sensor observations are noisy versions of the field’s values, such that

$$z_k = f(x, y) + \epsilon_k, \quad (x, y) \in \mathcal{A},$$

A. Bayesian estimation

In the context of the probabilistic Bayesian estimation, we assume that a parametric model for the field is known,

$$f(x, y) = f(x, y; \mathbf{a}), \quad \mathbf{a} \sim p_{\mathbf{a}}: \text{unknown parameter},$$

where $p_{\mathbf{a}}$ is the prior distribution of \mathbf{a} , such that the task of estimating the field is formally equivalent to the estimation of \mathbf{a} . If the *noise components affecting distinct measures are statistically independent*, then the posterior density $p(\mathbf{a}|\mathbf{z}^k)$ of

a given observed data $\mathbf{z}^k = \{z_i\}_{i=1}^k$ is a sufficient statistic for the estimation of the field f , and it can be recursively updated [11] using the Bayes relation. For finite-dimension linear models,

$$f(x, y; \mathbf{a}) = \mathbf{a}^T \mathbf{f}(x, y), \quad \mathbf{f}(x, y) = [f_1(x, y), \dots, f_I(x, y)]^T,$$

where $\{f_i\}_{i=1}^I$ are a set of known functions, and Gaussian distributions (both prior and observation noise), the posterior densities are normal, $a|z^k \sim \mathcal{N}(\mu_k, \Sigma_k)$, and closed-form expressions can be obtained for their means and variances:

$$\Sigma_{k+1}^{-1} = \Sigma_k^{-1} + \frac{1}{\sigma_\epsilon^2} \mathbf{f}_{k+1} \mathbf{f}_{k+1}^T \quad (2)$$

$$\Sigma_{k+1}^{-1} \boldsymbol{\mu}_{k+1} = \Sigma_k^{-1} \boldsymbol{\mu}_k + \frac{1}{\sigma_\epsilon^2} z_{k+1} \mathbf{f}_{k+1} \quad (3)$$

In this equation σ_ϵ^2 is the variance of the observation noise. This version of the Bayes equations are known as the “information form” of the least-squares estimator, and allow asynchronous integration of new measures. The statistics $\eta_k = (\Sigma_k^{-1}, \Sigma_k^{-1} \boldsymbol{\mu}_k)$ maintained at each node all have the same form, and their size does not grow with the number of measurements. Since matrix addition is associative and commutative, the fusion equations presented above can be easily extended to several measures.

Note that under the linear-Gauss assumption, the covariance matrix of the posterior density does not depend on the measurement itself (z_k) but only on the local observation operator ($\mathbf{f}(x_k, y_k)$). Since the amount of information conveyed by \mathbf{z}^k about \mathbf{a} is a function of Σ_k , the impact of future measurements on the quality of the estimate can be predicted.

B. Bounded-error

The estimation method presented above relies on knowledge of the statistical properties of the noise, which is a rather strong assumption. Fogel and Huang proposed in their seminal paper [8] a deterministic approach to parameter estimation under a *bounded error* assumption. The key idea is that, assuming observation noise is bounded, each measurement defines a set of feasible parameters values (i.e., which are consistent with the measurement, the model and the bound). The actual value of the parameter lies in the intersection of the sets associated to the measures. Several representations for the so-called *feasible parameter set* have been proposed, including parallelotopes [12], polytopes [13] and ellipsoids [8]. Using polytopes gives the optimal solution but the state maintained at each node may potentially grow as the number of measurements, violating the scalability requirement. On the contrary, using ellipsoids leads to a scalable solution (the size of the estimate does not grow with k and is $\mathcal{O}(I^2)$) and is directly comparable to the statistical Bayesian approach with Gaussian posteriors. In the following we use ellipsoids to describe the feasible parameter sets:

$$\mathcal{E}_k = \left\{ \mathbf{a} : (\mathbf{a} - \mathbf{a}_k^c)^T P_k^{-1} (\mathbf{a} - \mathbf{a}_k^c) \leq 1 \right\} \text{ and } \hat{\mathbf{a}}_k = \mathbf{a}_k^c \quad (4)$$

Assume the measurement errors ϵ_k are bounded by r : $|\epsilon_k(x, y)| \leq r$. Then, each new measurement z_k defines two

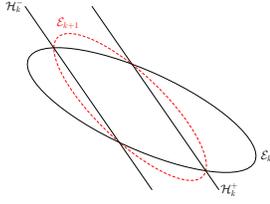


Figure 2. Minimum-determinant ellipse of the form $\mathcal{E}_{k+1}(q)$ (5)

hyperplanes that bound the set of feasible parameters:

$$\mathcal{H}_k^+ = \left\{ \mathbf{a} : \frac{1}{r} (z_k - \mathbf{a}^T \mathbf{f}_k) \leq 1 \right\}$$

and

$$\mathcal{H}_k^- = \left\{ \mathbf{a} : \frac{1}{r} (z_k - \mathbf{a}^T \mathbf{f}_k) \geq -1 \right\}$$

The updated feasible parameter set \mathcal{E}_{k+1} must be of the form (4) and enclose $\mathcal{E}_k \cap \mathcal{H}_{k+1}^+ \cap \mathcal{H}_{k+1}^-$. A simple one-parameter family of sets verifying these properties is [8]:

$$\mathcal{E}_{k+1}(q) = \left\{ \mathbf{a} : (\mathbf{a} - \mathbf{a}_k^c)^T P_k^{-1} (\mathbf{a} - \mathbf{a}_k^c) + q \frac{1}{r^2} (z_k - \mathbf{a}^T \mathbf{f}_k)^2 \leq 1 + q \right\}, \quad q \geq 0 \quad (5)$$

This ellipsoid can be rewritten in the form (4) using the following recursive updating rules:

$$P_{k+1} = \zeta_k Z_k \quad (6)$$

$$\mathbf{a}_{k+1}^c = Z_k (P_k^{-1} \mathbf{a}_k^c + qr^{-2} z_k \mathbf{f}_k) \quad (7)$$

$$Z_k = \left(P_k^{-1} + qr^{-2} \mathbf{f}_k \mathbf{f}_k^T \right)^{-1} \quad (8)$$

$$\zeta_k = 1 + q - \frac{qr^{-2}}{1 + qr^{-2} G_k} e_k^2 \quad (9)$$

where $e_k = z_k - \mathbf{a}_k^T \mathbf{f}_k$, $G_k = \mathbf{f}_k^T P_k \mathbf{f}_k$, and $q \geq 0$. The value of q is chosen to minimize some optimization criterion. In the traditional case of volume minimization (minimizing $\det P_{k+1}$), the optimal value q_k^* of q is the – positive if any, or zero otherwise – root of a second order polynomial and can be expressed in closed form. Finally, $\mathcal{E}_{k+1} = \mathcal{E}_{k+1}(q_k^*)$.

Contrary to the Bayesian approach, the bounded-error (BE) method enables optimization of the feasible parameter set (i.e, the ellipsoid) according to any criterion – such as the determinant, the trace or the radial spectrum of P_k – while integrating new measurements. On the other hand, due to its deterministic nature, a new measurement may not convey any additional information. This case occurs when the hyperplane \mathcal{H}_{k+1}^+ and \mathcal{H}_{k+1}^- do not intersect the ellipsoid \mathcal{E}_k . In the case where only one hyperplane intersects the ellipsoid, a correction has been proposed so that the one-parameter family of ellipsoids (5) encloses more tightly the intersection [8].

IV. DISTRIBUTED DATA FUSION

The fusion step is the weak link of collaborative sampling architectures. In this section we present Covariance Intersection (CI), the most frequently used fusion algorithm under

unknown correlation, and illustrate its performance on simple examples. Also, we highlight pathological behaviors of CI on some special cases. Based on purely geometric considerations, we extend the bounded-error estimation to a fusion method which does not suffer from pathologic cases aforementioned.

A. Covariance Intersection (CI)

For linear-Gaussian problems, fusion is a linear operation. Imposing unbiasedness of the estimator resulting of the fusion of unbiased estimates implies that the linear combination be convex, and finding the fusion operator consists in choosing an element in a I^2 -dimensional space. In their original paper [7], Julier and Uhlmann proposed CI which reduces the aforementioned problem to a one-dimensional optimization problem, using a convex linear combination:

$$\Sigma_C^{-1} = \omega \Sigma_A^{-1} + (1 - \omega) \Sigma_B^{-1} \quad (10)$$

$$\Sigma_C^{-1} \boldsymbol{\mu}_C = \omega \Sigma_A^{-1} \boldsymbol{\mu}_A + (1 - \omega) \Sigma_B^{-1} \boldsymbol{\mu}_B \quad (11)$$

CI has been proved to provide consistent estimates (i.e., $\Sigma_C \geq \mathbb{E}[(\boldsymbol{\mu} - \mathbf{a})(\boldsymbol{\mu} - \mathbf{a})^T]$) for any $\omega \in [0, 1]$ meaning that the free parameter ω can be chosen to optimize some criterion.

The consistency of the estimator resulting from CI can be interpreted as a conservativeness property since the fusion method never “creates” virtual information – as the Bayesian method may do – when fusing correlated estimates. However, CI may discard fresh information, considering them as redundant. This is revealed by the fact that the centers of the ellipsoids are not taken into account in the determination of the covariance Σ_C : suppose that two nodes have two estimates with the same covariance matrix but different centers as illustrated in Figure 3. Then, CI returns the same covariance matrix for any value of ω leaving the center (i.e., the estimate $\hat{\mathbf{a}}$ of \mathbf{a}) indetermined.

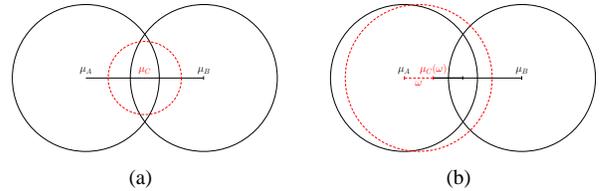


Figure 3. Fusion of two estimates with the same covariance using (a) Naïve Bayes (NB) (b) Covariance Intersection (CI)

B. Feasible Parameters Set Intersection

Although CI behaves strangely in some particular cases and ignores the precious information conveyed by the centers of the ellipsoids, its core idea of intersecting uncertainty ellipsoids can be adapted to a deterministic context: let \mathcal{E}_1 and \mathcal{E}_2 be two estimates of \mathbf{a} , then any set enclosing $\mathcal{E}_1 \cap \mathcal{E}_2$ contains the actual value of \mathbf{a} , and is thus a valid feasible parameter set. To be directly and fairly comparable to CI, we parametrize our search space with only one parameter. Inspired by both (5) and (11), we naturally define the search space as:

$$\mathcal{E}(\omega) = \left\{ \mathbf{a} : (1 - \omega) (\mathbf{a} - \mathbf{a}_1^c)^T P_1^{-1} (\mathbf{a} - \mathbf{a}_1^c) + \omega (\mathbf{a} - \mathbf{a}_2^c)^T P_2^{-1} (\mathbf{a} - \mathbf{a}_2^c) \leq 1 \right\}, \quad 0 \leq \omega \leq 1 \quad (12)$$

yielding a set of ellipsoids given by [9] (see [6] for a detailed proof):

$$P_0 = (1 - \zeta_0)Z_0 \quad (13)$$

$$\mathbf{a}_0^c = Z_0 \left((1 - \omega)P_1^{-1}\mathbf{a}_1^c + \omega P_2^{-1}\mathbf{a}_2^c \right) \quad (14)$$

$$Z_0 = \left((1 - \omega)P_1^{-1} + \omega P_2^{-1} \right)^{-1} \quad (15)$$

$$\zeta_0 = (1 - \omega) (\mathbf{a}_1^c - \mathbf{a}_0^c)^T P_1^{-1} (\mathbf{a}_1^c - \mathbf{a}_0^c) - \omega (\mathbf{a}_2^c - \mathbf{a}_0^c)^T P_2^{-1} (\mathbf{a}_2^c - \mathbf{a}_0^c) \quad (16)$$

Figure 4 shows an example of the minimum-determinant ellipsoid enclosing the intersection of two uncertainty sets (Figure 4(a)) and in the case where the uncertainty sets are the same (Figure 4(b)). One may notice that contrary to the CI method, the estimates $\hat{\mathbf{a}}$ are taken into account yielding a smaller volume of uncertainty and a unique new estimate.

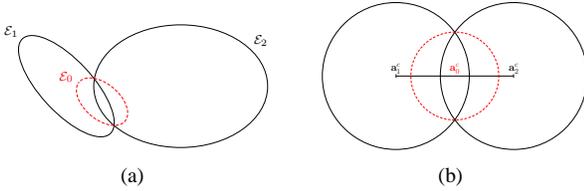


Figure 4. (a) Sample minimum-determinant ellipse of the form $\mathcal{E}(\omega)$ (12), (b) fusion of two estimates with the same covariance using ellipsoid intersection

V. COLLABORATIVE ADAPTIVE SENSING

As stated in [6], the key points of distributed estimation are (i) collaboration between nodes and (ii) adaptation to the sensed field. The first aims at maximizing the joint information conveyed by the measurements performed by the nodes while the latter uses a model for field so as to determine the points that will bring the most information, taking into account the current state of knowledge. In both the deterministic and probabilistic frameworks, the confidence associated to an estimate is described by an ellipsoid. For vector estimates (\mathbf{a} is not a scalar) the uncertainty set obtained by integration of additional data sampled does not allow the definition of a complete order relation over \mathcal{A} that can be used to select where the next measures should be made. The use of a scalar function of the matrix M describing the uncertainty set (the positive definite matrices P or Σ depending on whether CI or BE is used) – such as trace, determinant or spectral radius – can be used to define a total order relation on which the choice of the next measuring points can be based. These metrics have a direct geometric interpretation, making the analysis of sensing points determination method easier to apprehend.

The most common criterion is $\det(M)$ which is a increasing function of the uncertainty volume of the uncertainty ellipsoid. It can be shown that maximizing the volume is equivalent to maximizing $\mathbf{f}^T M \mathbf{f}$, which can be interpreted as choosing the sensing point where \mathbf{f} is closest to the direction along which $\hat{\mathbf{a}}_k$ is the most uncertain. In the deterministic case, choosing the observation direction collinear to the longest axis (eigenvector associated to the largest eigenvalue) of M is the optimal choice, since “cutting” the ellipsoid with hyperplane

orthogonal to the observation vector leads to the largest possible reduction of the uncertainty volume and deformation.

Based on a different line of reasoning, a *structural* algorithm was proposed for choosing the sensing positions. The algorithm uses a orthonormal basis \mathcal{B} , which is iteratively updated starting with $\mathcal{B} = \mathcal{B}_0$. At each stage the new sensing point is determined by maximizing the projection of \mathbf{f} on $\text{Span}(\mathcal{B})^\perp$ and \mathcal{B} is updated by a Gram-Schmidt iterative orthonormalization algorithm:

$$\{x^i, y^i\} = \underset{x, y \in \mathcal{G}}{\text{argmax}} \left\| (I - B_{i-1} B_{i-1}^T) \mathbf{f}(x, y) \right\|,$$

and

$$B_i = B_{i-1} \cup \frac{(I - B_{i-1} B_{i-1}^T) \mathbf{f}(x^i, y^i)}{\left\| (I - B_{i-1} B_{i-1}^T) \mathbf{f}(x^i, y^i) \right\|}.$$

The choice of \mathcal{B}_0 determines the places where the nodes should not go and is chosen with two goals:

- impose motion to a different position:

$$\mathcal{B}'_0 = \{ \mathbf{f}(x_k^i, y_k^i) \}_{i=1}^N$$

- avoiding the directions along which \mathbf{a} is best known:

$$\mathcal{B}''_0 = \{v_1, \dots, v_{I-2N}\}$$

where v_i is the eigenvector associated to the i -th smallest eigenvalue (shortest axis of the uncertainty ellipse) of Σ .

The initial basis is obtained by combining \mathcal{B}'_0 with \mathcal{B}''_0 .

VI. RESULTS

In this section, we study the performance of the estimation framework presented in this paper using computer simulations. The field to be estimated is composed of 20 Gaussian shaped basis functions with unit amplitude, and random scale and localization. The coefficients α_i are drawn from a uniform distribution. We consider a network of 6 robots, initially deployed at random inside a small launching zone in the region of interest \mathcal{A} (of size $10,000 \times 10,000$) – reflecting a realistic deployment scenario. The communication radius and the diameter of the action zone are respectively 2,000 and 2,500. We conducted a Monte-Carlo study, performing two sets of 25 independent simulations of the architecture keeping both the field and the launching zone fixed. One set uses Bayesian estimation together with CI for the fusion step, and the other uses a deterministic approach to both estimation and fusion.

We use two related criteria to evaluate the performance of our architecture: (i) the error on the field estimate \hat{f} and (ii) the error on the globally reconstructed parameter itself, $\hat{\mathbf{a}}$ (i.e., the fusion of the set of estimates over the whole network):

$$\hat{\mathbf{a}} = \text{fuse}(\hat{\mathbf{a}}_1, \dots, \hat{\mathbf{a}}_N) \quad \text{and} \quad \hat{f}(x, y) = \hat{\mathbf{a}}^T \mathbf{f}(x, y)$$

where the fusion of the estimates is performed iteratively.

The first is motivated by the original goal of providing a global picture of the field, as mentioned in the Introduction. We use distortion as a metric to evaluate the quality of the field estimates using equation (1) together with the \mathcal{L}^1 norm.

However, this criterion does not directly reflect the precision of the estimate along each component of \mathbf{a} , which may – in the underwater framework – be related to interesting natural phenomena. In fact, the strategy that drives the network (the determination of the sensing points) has been designed to reduce the deformation of the uncertainty ellipsoid associated to \mathbf{a} (trying to avoid that some of its components be much better known than other). We plot these two criteria as a function of time, observing the rates of convergence and the asymptotic behaviors of the two different estimation/fusion frameworks. The performance improvement of the collaborative architecture with respect to independent node operation has already been demonstrated in [5].

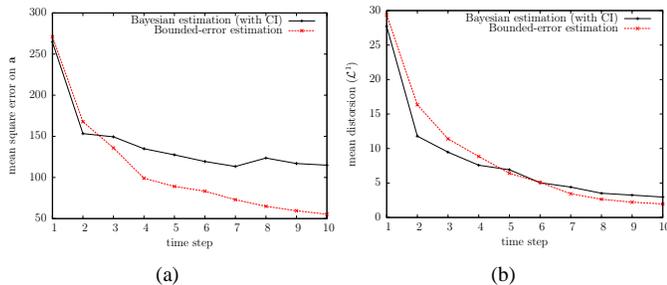


Figure 5. Metrics on fusion of the collection of network estimates over time: (a) $\|\mathbf{a} - \hat{\mathbf{a}}\|^2$; (b) \mathcal{L}^1 norm.

Figure 5 presents the mean square-error on the estimate of \mathbf{a} together with the distortion of the field globally reconstructed over time. These plots support our two major conclusions: (i) the BE fusion method is superior to CI leading to better estimates of \mathbf{a} (see 5). (ii) the network control strategy based on the geometric criterion is robust to poor error characterization of the fusion methods, as the comparable overall performance of the CI and BE demonstrates. In our architecture, what determines the network dynamics is the subspace generated by the directions of current smallest uncertainty, irrespective of the detailed amount of uncertainty along each one of them. This means that many different matrices will lead to the same orthonormal basis \mathcal{B} , and thus to the same set of chosen sampling points which explains the similar performances for the different fusion methods observed in Figure 5(b).

VII. CONCLUSIONS AND FUTURE WORK

In this paper we extended the overall distributed algorithm for collaborative adaptive sensing proposed in [5] proposing a deterministic approach to parameter estimation and fusion. Assuming bounded-errors, and bounding the uncertainty domains by ellipsoids, provided a comprehensive geometric view of the information processing aspects of the architecture (estimation, fusion and adaptive determination of the sensing points). All blocks of the architecture now rely on the same formalism, becoming matched to each other, and improving the architecture consistency and coherency. Besides this overall coherency, the deterministic setting provided an inherent robustness with respect to one of the difficult problems affecting distributed fusion (the rumor problem). It also improved estimation when

compared to the traditional Bayesian estimation with CI as fusion method, as the computer simulations demonstrated.

More generally, design of efficient distributed fusion methods, together with robust criteria for sampling site selection, are critical steps for being able to fully exploit the potential advantages of the flexible and powerful framework of mobile sensor networks, providing effective answers to the difficulties resulting from the stringent constraints of these systems.

Since our architecture is built at the application level, the communication aspects are idealized (fixed radius and error-free communications with unlimited bandwidth). For real-world implementation, these aspects must be taken into account, especially in underwater environment where the communication conditions are harsh.

Our work can be improved in several aspects. The final fusion is done *off-line* with no computational power limitation or time restrictions, and thus may be performed more efficiently than simply fusing iteratively (pairwise) the node estimates. The same problem affects fusion inside each cluster, but under stronger constraints in terms of memory and computational power. Designing a fusion method that directly assesses the problem of fusing multiple ellipsoids is a point for future study. Also, looking for the *subspace* that results in the largest observation rank instead of iteratively optimizing a geometric criterion should lead to a more efficient set of sampling points.

REFERENCES

- [1] T. Arici and Y. Altunbasak, "Adaptive Sensing for Environment Monitoring using Wireless Sensor Networks," *IEEE Wireless Communications and Networking Conference*, vol. 4, pp. 2347–2352, March 2004.
- [2] M. Rabbat and al., "Decentralized Compression and Predistribution via Randomized Gossiping," in *IPSN '06: Proc. of the Conference on Information Processing in Sensor Networks*, April 2006, pp. 51–59.
- [3] M. Rahimi and al., "Adaptive Sampling for Environmental Field Estimation using Robotic Sensors," in *IROS '05: Proc. of the Conference on Intelligent Robots and Systems*, August 2005, pp. 3692–3698.
- [4] N. E. Leonard and al., "Collective Motion, Sensor Networks and Ocean Sampling," *Proc. of the IEEE, Special Issue on Networked Control Systems*, vol. 95, pp. 48–74, 2007.
- [5] K. Huguenin, "An Architecture for Distributed Collaborative Adaptive Sensing," in *UUST'07: Proc. of the International Symposium on Unmanned Untethered Submersibles Technology*, 2007.
- [6] —, "An Architecture for Distributed Collaborative Adaptive Sensing," Master's thesis, ENS Cachan/UNSA, I3S Lab., Sophia Antipolis, France, June 2007.
- [7] S. Julier and J. Uhlmann, "A Non-divergent Estimation Algorithm in the Presence of Unknown Correlations," in *Proc. of the American Control Conference*, June 1997.
- [8] E. Fogel and Y. F. Huang, "On the Value of Information in System Estimation–Bounded Noise Case," *Automatica*, vol. 18, no. 2, pp. 229–238, 1982.
- [9] L. Ros, A. Sabater, and F. Thomas, "An ellipsoidal calculus based on propagation and fusion," *IEEE Trans. on Systems, Man, and Cybernetics, Part B*, vol. 32, no. 4, pp. 430–442, 2002.
- [10] S. Basagni, "Distributed Clustering for Ad Hoc Networks," in *ISPAN '99: Proc. of the International Symposium on Parallel Architectures, Algorithms and Networks*, 1999, p. 310.
- [11] H. L. V. Trees, *Detection, Estimation, and Modulation Theory, Part I*. New York: John Wiley and Sons, 1968.
- [12] A. Vicino and G. Zappa, "Sequential Approximation of Feasible Parameters for Identification with Set Membership Uncertainty," *IEEE Trans. on Automatic Control*, vol. 41, no. 9, pp. 774–786, June 1996.
- [13] E. Walter and H. Piet-Lahanier, "Exact Recursive Polyhedral Description of the Feasible Parameter Set for Bounded-Error Models," *IEEE Trans. on Automatic Control*, vol. 34, no. 8, pp. 911–915, August 1989.