SENSOR NETWORK SOURCE LOCALIZATION VIA PROJECTION ONTO CONVEX SETS (POCS)

Alfred O. Hero, III and Doron Blatt

Department of EECS, University of Michigan, Ann Arbor, MI
{dblatt, hero}@eecs.umich.edu

ABSTRACT

This paper addresses the problem of locating an acoustic source using a sensor network in a distributed manner, i.e., without transmitting the full data set to a central point for processing. This problem has been traditionally addressed through the nonlinear least squares or maximum likelihood framework. These methods, even though asymptotically optimal under certain conditions, pose a difficult global optimization problem. It is shown that the associated objective function may have multiple local optima and saddle points and hence any local search method might stagnate at a sub-optimal solution. In this paper, we formulate the problem as a convex feasibility problem and apply a distributed version of the projection onto convex sets (POCS) method. We give a closed form expression for the projection phase, which usually constitutes the heaviest computational aspect of POCS. Conditions are given under which, when the number of samples increases to infinity or in the absence of measurement noise, the convex feasibility problem has a unique solution at the true source location. In general, the method converges to a limit point or a limit cycle in the neighborhood of the true location. Simulation results show convergence to the global optimum with extremely fast convergence rates compared to the previous methods.

1. INTRODUCTION

The problem of locating a source that emits acoustic waves using a wireless network of acoustic sensors has been addressed by several authors (see e.g. [1, 2, 3, 4] and references therein). This problem has been traditionally solved though nonlinear least squares estimation, which is equivalent to the maximum likelihood estimator when the observation noise is modeled as a white Gaussian process. Rabbat and Nowak [2, 5] proposed a distributed implementation of the incremental gradient algorithm to solve this problem in a distributed manner, i.e., without the need to transmit the data to a central point for processing. A drawback of their method, or any other local search method, is that it is sensitive to local optima and saddle points. As will be shown below, the objective function associated with this problem is indeed multi-modal and may have a number of local optima and saddle points.

To overcome this shortcoming, in this paper the problem is formulated as a convex feasibility problem instead of nonlinear least squares. Necessary and sufficient conditions are given under which, when the number of samples increases to infinity or in the absence of measurement noise, the convex feasibility problem has a unique solution at the true source location.

This research was partially supported by DARPA-MURI grant ARO DAAD 19-02-1-0262.

To solve the convex feasibility problem we propose the projection onto convex sets (POCS) method [6] (see also [7] Ch. 5). We show that this method can be implemented in a distributed manner, i.e., each sensor performs the bulk of its computations based on its own data and it is not required that the full data set be sent to a central point for processing. As in Nowak's distributed EM algorithm [8], a number of communication cycles across the network is sufficient for the implementation of the estimator. A closed form expression is given for the usually computationally demanding projection phase of POCS, which leads to a computationally efficient implementation. For a finite number of samples it is shown that convergence to a point or a limit cycle in the vicinity of the true source position occurs. Simulation results show global convergence of the proposed method in contrast to a local search method, with extremely fast convergence rates.

2. PROBLEM FORMULATION

Consider a sensor network composed of L sensors distributed at known spatial locations, denoted r_l, $l = 1, \ldots, L$, where $r_l \in \mathbb{R}^2$. Generalization to $\mathbb{R}^3$ is straightforward but is not explored here. An acoustic source is located at an unknown location $\theta^* \in \mathbb{R}^2$. Each sensor collects n noisy measurements of the acoustic signal transmitted by the source, denoted $\{y_t^l\}$, $l = 1 \ldots, L$, $t = 1, \ldots, n$. Following [1, 2, 3], the reading of the source's signal strength at sensor l is modeled by

$$y_t^l = \frac{A}{||r_l - \theta||^\beta} + w_t^l, \quad t = 1, \ldots, n$$

where A and β are the source signal strength and isotropic attenuation coefficient, respectively, and w_t^l is a zero-mean measurement noise with unknown variance σ^2. We assume that β and A are known. The first assumption is reasonable if a characterization of the terrain in which the network is deployed is available. The second one is valid when an additional sensor is added to an already deployed network and the new sensor transmits an acoustic signal with known power to enable the network to estimate its location.

Formulated as a nonlinear least squares problem[1], the source's location can be estimated by

$$\widehat{\theta}_{NLS} = \arg\min_\theta \sum_{l=1}^{L} \sum_{t=1}^{n} \left[y_t^l - \frac{A}{||r_l - \theta||^\beta} \right]^2. \qquad (1)$$

The fact that the objective function is a sum of L components was exploited in the implementation of the distributed incremental gradient method in [2, 5]. However, since the objective function has

[1] Equivalent to maximum likelihood when the noise is assumed to be a white Gaussian process.

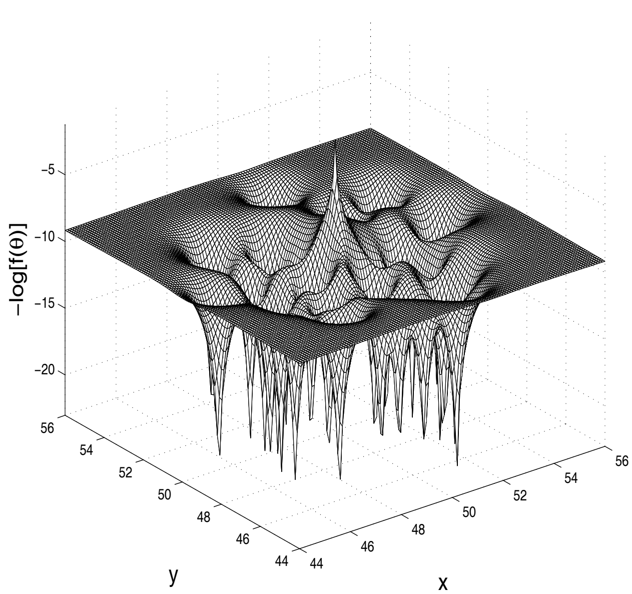

Fig. 1. The negative log of the nonlinear least squares objective function.

multiple local optima and saddle points, the incremental gradient method may stagnate at one of these sub-optimal solutions instead of converging to the optimal one. A realization of the negative log of the objective function in (1) is presented in Fig. 1. The details of the simulation that generated this figure are given in Sec. 4. It can be seen that the objective function has many local optima and saddle points.

An alternative problem formulation of the problem of estimating the source's location is the following. Consider the l summands in the objective function (1). It is easily seen that the function

$$f_l(\theta) = \sum_{t=1}^{n} \left[y_t^l - \frac{A}{||r_l - \theta||^\beta} \right]^2$$

obtains its minimum on the circle

$$C_l = \left\{ \theta \in \mathbb{R}^2 : ||\theta - r_l|| = \left[\frac{A}{\overline{y}^l} \right]^{1/\beta} \right\}$$

where $\overline{y}^l = n^{-1} \sum_{t=1}^{n} y_t^l$. Let D_l be the disk defined by

$$D_l = \left\{ \theta \in \mathbb{R}^2 : ||\theta - r_l|| \leq \left[\frac{A}{\overline{y}^l} \right]^{1/\beta} \right\}. \qquad (2)$$

Then our estimation problem is solved by finding a point in the intersection of the sets D_l, $l = 1, \ldots, L$, that is,

$$\widehat{\theta} \in D = \bigcap_{l=1}^{L} D_l \subset \mathbb{R}^2. \qquad (3)$$

Note that due to observation noise the intersection D might be empty. In this case, our estimator is any point that minimizes the sum of distances to the sets D_l, $l = 1, \ldots, L$, that is,

$$\widehat{\theta} = \arg\min_{\theta \in \mathbb{R}^2} \sum_{l=1}^{L} ||x - \mathcal{P}_{D_l}(x)|| \qquad (4)$$

where for a set $S \subseteq \mathbb{R}^2$ and a point $x \in \mathbb{R}^2$, $\mathcal{P}_S(x)$ is the orthogonal projection of x onto S, that is,

$$\mathcal{P}_S(x) = \arg\min_{y \in S} ||x - y|| \qquad (5)$$

where $||\cdot||$ is the Euclidean norm. Observe that (4) includes (3) as a special case when a minimum value of zero is attainable. Since the sets D_l are convex, both the consistent and inconsistent convex feasibility problems, (3) and (4), respectively, can be solved via the POCS method to be described below.

Before describing POCS, we give necessary and sufficient conditions for the consistency of the estimator (4). Denote by $\mathcal{H}$ the convex hull of the sensors' spatial locations, i.e.,

$$\mathcal{H} = \left\{ x \in \mathbb{R}^2 : x = \sum_{l=1}^{L} \alpha_l r_l,\ \alpha_l \geq 0,\ \sum_{l=1}^{L} \alpha_l = 1 \right\}.$$

It is possible to show geometrically (see Fig. 2) that when the number of samples increases to infinity, or in the absence of measurement noise, the convex feasibility problem (4) has a unique solution at the true source's location, denoted by θ^*, if and only if θ^* lies in $\mathcal{H}$, that is,

$$\bigcap_{l=1}^{L} \left\{ \theta \in \mathbb{R}^2 : ||r_l - \theta|| \leq ||r_l - \theta^*|| \right\} = \{\theta^*\}$$
$$\text{if and only if} \quad \theta^* \in \mathcal{H}$$

where $L \geq 2$.

In the general case of finite number of samples and finite signal to noise ratio, one of two cases can occur: (a) $D \neq \emptyset$, and (b) $D = \emptyset$. In the former, the POCS method is guaranteed to converge to a point in D. In the latter, the POCS method converges to a limit cycle in the vicinity of the point that minimizes the sum of distances to the sets D_l (2), or, when a certain sequence of relaxation parameters are used, the method converges to the optimal solution.

3. DISTRIBUTED IMPLEMENTATION OF POCS

The POCS method [6, 7] is given by the following algorithm.

1. Initialization: θ^0 is arbitrary.
2. Iterative step: For all $k \geq 0$,

$$\theta^{k+1} = \theta^k + \lambda_k \left[\mathcal{P}_{D_{\kappa(k)}}(\theta^k) - \theta^k \right] \qquad (6)$$

where $\{\lambda_k\}_{k \geq 1}$ is a sequence of relaxation parameters satisfying for all k, $\epsilon_1 \leq \lambda_k \leq 2 - \epsilon_2$ for some $\epsilon_1, \epsilon_2 > 0$, $\kappa(k) = k \bmod L$, and $\mathcal{P}_S(x)$ is defined in (5).

Usually the projection operator is the most computationally demanding element of POCS. In our application, however, a closed form expression is available for (6). Clearly, if $||\theta - r_l|| \leq \left[\frac{A}{\overline{y}^l} \right]^{1/\beta}$ then $\theta \in D_l$ and $\mathcal{P}_{D_l}(\theta) = \theta$, otherwise,

$$\mathcal{P}_{D_l}(\theta^k) = r_l + [\alpha \cos(\phi),\ \alpha \sin(\phi)]^T \qquad (7)$$

where $\alpha = \sqrt[\beta]{A / \overline{y}^l}$, and $\phi = \operatorname{atan}(\theta^k(2) - r_l(2), \theta^k(1) - r_l(1))$, where $\operatorname{atan}(\cdot, \cdot)$ is the four quadrant inverse tangent function, and for a vector vector $x \in \mathbb{R}^2$, $x(1)$ and $x(2)$ denote its first and second coordinates, respectively.

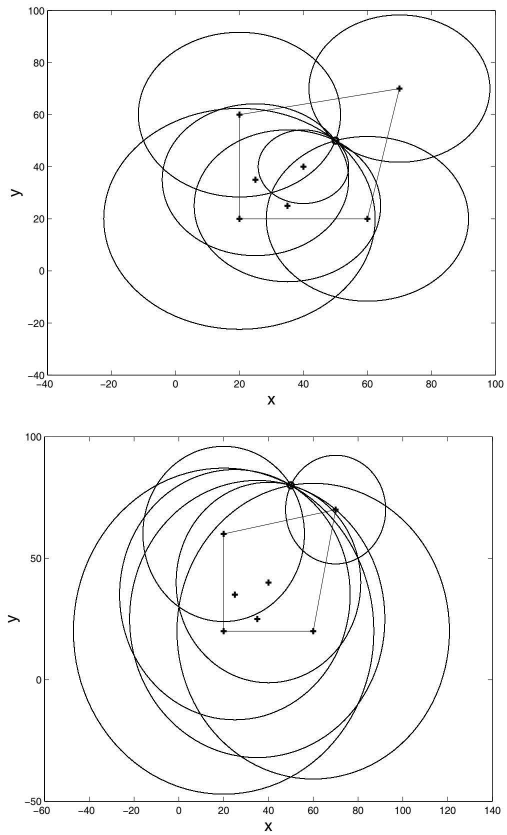

Fig. 2. Source, denoted by a black dot, is located inside (top) and outside of (bottom) the convex hull $\mathcal{H}$ of the sensors' locations, denoted by crosses.

The relaxation parameters λ_k play an important role in the convergence of the method. At the first phase of the implementation of the POCS method, the relaxation parameters are set to 1. As the method progresses, a convergence criterion is repeatedly checked. If convergence to a single point is detected, it is concluded that $D \neq \emptyset$, and the final estimate $\widehat{\theta}$ is set to the limit point. If convergence to a limit cycle is detected, i.e., each sensor converges to a different value, it is concluded that $D = \emptyset$ and the method enters phase two. At phase two the relaxation parameters are decreased at a rate of $1/k$. In [9], it is shown that this relaxation sequence leads to convergence to the point x that minimizes the sum of squared distances to the sets D_l, that is, to $\widehat{\theta}$ defined in (4). It should be noted that if the transition to phase two occurs prematurely, this convergence result still holds. The effect will be a slowdown of convergence. A sub-optimal but computationally cheaper alternative to phase two is to approximate $\widehat{\theta}$ by the arithmetic mean of the points in the limit cycle. This simple approach was used in the simulation reported in Sec. 4. Due to its global convergence properties, the estimate resulting from the POCS method could also be used to trigger a local search for the nonlinear least squares estimator such as the one in [2].

Note that all the information required for the computation of (7) (or (6)) is available at sensor l and hence a distributed implementation is possible. Following [8], assume without loss of generality that the indices $l = 1, \ldots, L$ correspond to a cycle through the network. Let sensor 1 be initiated with a pre-specified initial value θ^0. Sensor 1 generates θ^1 through (6) and transmits θ^1 to sensor 2. Upon receiving θ^k from sensor $\kappa(k)$, sensor $\kappa(k+1)$ calculates θ^{k+1} and transmits it to sensor $\kappa(k+2)$. The information cycle continues until the detection of convergence to either a limit point or a limit cycle. The convergence detection criteria can be easily implemented in a distributed manner as well. Phase two can be implemented in a similar way.

4. SIMULATION RESULTS

This section presents a simulation of a sensor network of $L = 5000$ nodes, distributed in a 100×100 field. Each sensor collects a single measurement of the acoustic source located at $\theta^* = [50, 50]^T$, which emits a signal with A set to 100. The measurement noise variance is $\sigma^2 = 1$ and the attenuation coefficient is $\beta = 2$. Following [3, 5], not all sensors participate in the estimation task. At an acquisition phase, each sensor decides whether or not a source is present using a simple threshold test. Only those sensors whose acoustic reading is above 10 participate. In the realization presented here, $L_1 = 31$ sensors detected the source and entered the estimation phase.

A realization of the objective function associated with the nonlinear least squares estimation method (1) is shown in Fig. 1. To optimize the viewing angle, the negative log of the objective function is presented. Hence, the optimum point is the global maximum rather than the minimum, which appears close to the true location of the source. The objective function has multiple local optima and saddle points, which impose difficulties on any local search method. In Fig. 3, the paths taken by the steepest descent (SD) method initiated from multiple points on a grid are presented on top of the contour plot of the nonlinear objective function (1). The SD method could also be implemented in a distributed manner, e.g., distributed Fisher scoring [10]. The initial points are depicted by crosses, followed by a line which follows the path taken by the algorithm, and ends at the convergence points depicted by circles. It is seen that only when the method is initiated close to the global optimum at the center of the plot, does convergence to the global optimum occur. The method mostly stagnates at local optima or saddle points.

In contrast to this shortcoming of the local search method, the proposed POCS method converges to the vicinity of the global optimum regardless of the initial point. In Fig. 4 the paths taken by the POCS method are presented. The order of the sensors in the information cycle described in Sec. 3 was selected randomly. A better illustration of the method is presented in Fig. 5, in which four representative paths are superimposed on top of the convex sets (discs) (2). At each iteration the sequence generated by the algorithm is projected onto a different disc, unless it is already inside it. It is seen that the convergence is extremely fast; after as few as three sub iterations (6), the sequence reaches the vicinity of the global optimum.

5. CONCLUSIONS

The problem of distributed acoustic source localization using a wireless sensor network was formulated as a convex feasibility problem and solved via the POCS method. The solution has global convergence properties with fast convergence rates. Finally we note that this concept can be generalized to other problems in

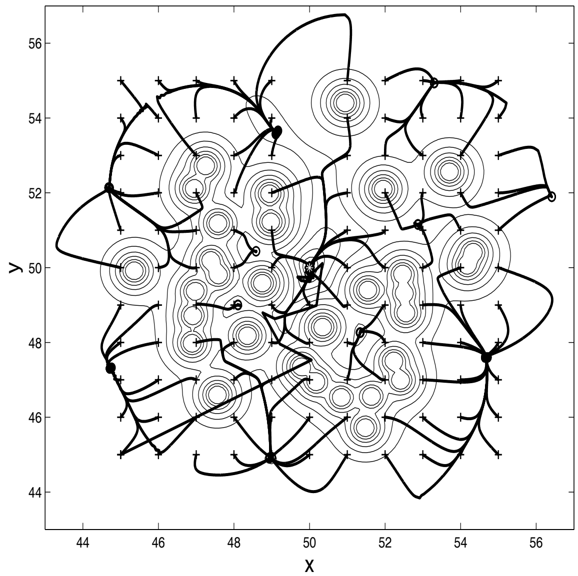

Fig. 3. Paths taken by the steepest descent method.

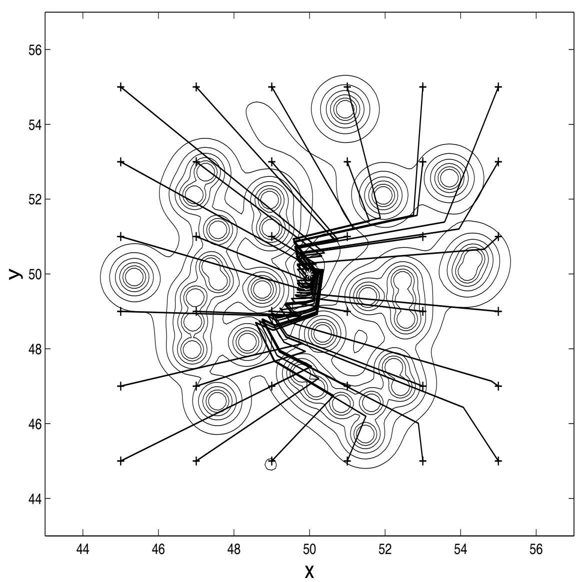

Fig. 4. Paths taken by the POCS method.

which the objective function depends on the parameters through terms of the form $||\theta - c_l||$, where $c_l, l = 1, \ldots, L$ are data dependent terms. In particular, this concept can be easily generalized to the three dimensional case. Generalization to non-isotropic media, unknown source amplitude, non-Gaussian noise, and the effect of channel noise are worthy of additional study.

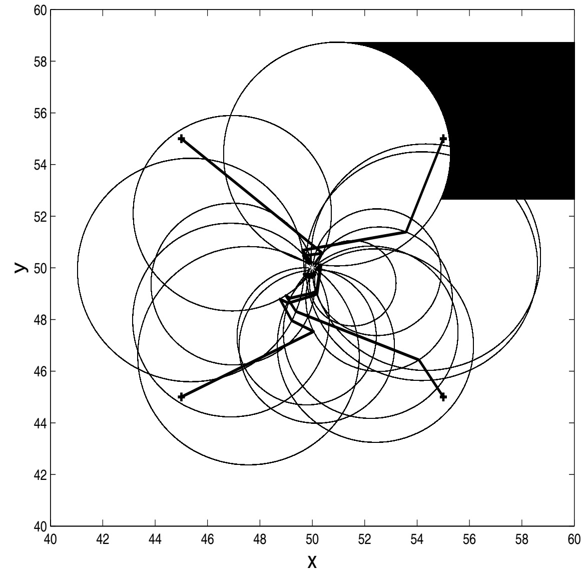

Fig. 5. Paths taken by the POCS method on top of the convex sets.

6. REFERENCES

[1] D. Li, K. D. Wong, Y. H. Hu, and A. M. Sayeed. Detection, classification, and tracking of targets. *IEEE Signal Processing Magazine*, 19(2):17–29, March 2002.

[2] M. G. Rabbat and R. D. Nowak. Decentralized source localization and tracking. In *Proceedings of the 2004 IEEE International Conference on Acoustics, Speech, and Signal Processing*, Montreal, Canada, May 2004.

[3] X. Sheng and Y. H. Hu. Energy based acoustic source localization. In *Information Processing in Sensor Networks, Second International Workshop*, Palo Alto, California, April 2003.

[4] D. Blatt and A. Hero. A convergent incremental gradient method with constant step size. *Submitted to: SIAM J. Optim.*, 2004.

[5] M. G. Rabbat and R. D. Nowak. Distributed optimization in sensor networks. In *Proceedings of the Third International Symposium on Information Processing in Sensor Networks*, Berkeley, California, April 2004.

[6] L. G. Gubin, B. T. Polyak, and E. V. Raik. The method of projections for finding the commmon point of convex sets. *USSR Computational Mathematics and Mathematical Physics*, 7:1–24, 1967.

[7] Y. Censor and S. A. Zenios. *Parallel Optimization: Theory, Algorithms, and Applications*. Oxford University Press, New York, 1997.

[8] R. D. Nowak. Distributed EM algorithms for density estimation and clustering in sensor networks. *IEEE Trans. Signal Processing*, 51(8):2245–2253, August 2003.

[9] Y. Censor and A. R. De Pierro. Steered sequential projections for the inconsistent convex feasibility problem. *Nonlinear analysis: theory, methods, and application, series A*, 59:385–405, 2004.

[10] D. Blatt and A. Hero. Distributed maximum likelihood for sensor networks. In *Proceedings of the 2004 IEEE International Conference on Acoustics, Speech, and Signal Processing*, Montreal, Canada, May 2004.