

Global Optimization for Multiple Transmitter Localization

Jill K. Nelson
Department of Electrical and
Computer Engineering
George Mason University
Fairfax, VA 22030
Email: jnelson@gmu.edu

Megan U. Hazen and Maya R. Gupta
Department of Electrical Engineering
University of Washington
Seattle, WA 98105
Email: megan@apl.washington.edu
gupta@ee.washington.edu

Abstract—We propose a global optimization approach to locating multiple transmitters within a geographic area. A set of sensor nodes are assumed to be present in the region and to measure total power received at their respective locations. These measurements are communicated to a processing node, which uses particle swarm optimization to find the transmitter locations that minimize the difference between the true received power and the estimated power based on the chosen propagation model. Clustering is used to generate initial estimates of the transmitter locations, thereby increasing the likelihood that the particle-based optimizer reaches the global minimum. Simulation results show that global optimization is an effective method for multiple transmitter localization and that generating “smart” initial conditions via clustering can yield an average performance improvement of over 25% compared to random initial conditions.

I. INTRODUCTION

In this paper, we address the problem of locating multiple transmitters within a region based on a set of received power measurements. While the general problem is relevant in a range of applications, perhaps the most notable application of current interest is to cognitive radio systems [1], [2]. A cognitive radio system is designed to operate concurrently with primary (conventional) radio systems without disrupting the operation of the primary systems. A cognitive radio system may be characterized as coordinated or uncoordinated with primary systems [3]. In the former case, the cognitive radio system cooperates with primary systems to establish transmission and general usage protocols. We are interested in the latter case, in which, rather than being assigned to certain regions of the wireless frequency spectrum, the cognitive radio system operates in an opportunistic manner, communicating in spectral regions that are assigned to primary systems but are under-utilized. When operating in an opportunistic fashion, the cognitive radio system must identify available spectrum it can occupy without disrupting the primary systems; these regions of unused spectrum are generally referred to as “spectral holes.”

Note that a spectral hole includes three dimensions: frequency, space, and time. (When sensors observe waveforms rather than only power, one may also consider signal parameters to be a dimension of the spectral hole [4].) Spectral usage

will vary across the range of wireless transmission frequencies, across the spatial region of interest, and possibly across time. Assuming slow variation with time, one may approach the identification of spectral holes by searching across possible frequencies at a particular location in the geographic region of interest or by searching across a region for holes at a particular frequency. In this paper, we take the latter approach. We consider a geographic region in which the members of a cognitive radio system wish to communicate using a particular frequency band F_0 . It is possible, however, that conventional systems with Federal Communications Commission (FCC) rights to the band F_0 are present and active in the region. It is the responsibility of the cognitive radio system to identify the area(s) within the geographic region in which primary transmitters are active and to avoid interfering with these systems.

Assuming that no prior knowledge of the location and activity of primary transmitters is available, the nodes within the cognitive radio system must identify the areas in which the primary systems are active based on measurements of received power in the band F_0 . A common method for declaring a spectral hole at a location is by simply setting a threshold on the maximum power observed at that location. For example, if a cognitive radio node at location \mathbf{p}_1 measures received power below P_{\max} , a spectral hole is declared at that location. The node at \mathbf{p}_1 may communicate with a node at location \mathbf{p}_2 if the measured power at \mathbf{p}_2 is also below P_{\max} . Erroneous detection results due to shadowing can be mitigated by incorporating cooperation, i.e., sharing information among the nodes in the cognitive radio system [5], [6].

Recent work has argued that spectral hole identification must move beyond simple detection-based methods [7], since such methods limit cognitive radio nodes to very conservative transmission. Approximate knowledge of transmitter locations can allow more aggressive spectrum exploitation. For example, consider cognitive radio nodes A and B pictured in Figure 1. Nodes A and B wish to communicate in a particular frequency band, and both sense the same power in that band. Whether or not A and B can communicate in the band depends on the location(s) of the source(s) from which the sensed power emanate(s). In the top scenario of Figure 1, a single transmitter

midway between A and B generates the observed power; A and B would not be able to transmit in the band of interest without causing interference. In the bottom scenario of Figure 1, the power observed by nodes A and B is the sum of power emanating from several transmitters. Because the transmitters are sufficiently distant from both A and B, the cognitive nodes can transmit without causing interference.

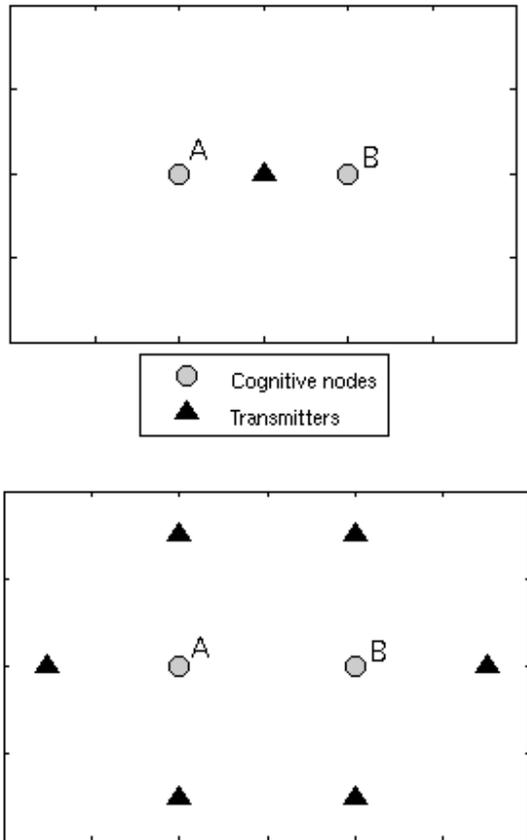


Fig. 1. Consider two scenarios in which cognitive nodes A and B sense the same power in a frequency band of interest. Top: The nodes cannot communicate in the chosen band without interfering with the transmitter between them. Bottom: The sensed power emanates from several transmitters, all of which are sufficiently distant from A and B to prevent interference. Hence, A and B can transmit in the chosen band.

To better exploit the spectrum, we consider the use of cooperative sensing to determine the approximate locations of the transmitters using the power levels observed by the cognitive radio nodes. Once transmitter locations have been estimated, FCC regulations can be employed to determine the radius around each transmitter in which opportunistic communication must be avoided. Cognitive radio nodes outside these circles can safely transmit without disrupting primary systems. The proposed technique identifies occupied spectrum and communicates around it rather than searching for unused spectrum.

An additional benefit of using transmitter localization to identify spectral holes arises when cognitive radio nodes

are mobile. While detection-based methods require nodes to repeatedly measure observed power as they move, knowledge of transmitter locations allows mobile nodes to determine whether or not transmission is allowable from a particular location without requiring that a measurement be obtained at that point. A natural extension would be to consider scenarios in which the primary transmitters are also in motion. A spectral-hole identification system based on transmitter localization lends itself to a simple tracking of the transmitters' locations, while the nodes in a detection-based system must continually update observed power measurements, even if the nodes are not mobile.

The problem of estimating transmitter locations based on a set of observed power measurements at arbitrary locations does not admit an obvious solution. In the absence of noise, power measurements at three locations determine a single transmitter location via trilateration. In the proposed scenario, however, multiple transmitters are present and contribute to the observed power at each sensor node. Additionally, error will likely be introduced both by imperfections in the sensors and by approximations in the signal propagation model. To address these challenges, we propose to find the transmitter locations by global optimization using “smart” initial conditions generated via clustering. The organization of the paper is as follows. The system model is defined in Section 2. In Section 3, we present the details of the proposed algorithm for transmitter localization. Section 4 contains simulated results of the proposed technique, and Section 5 concludes the paper.

II. SYSTEM MODEL

We consider a scenario in which M transmitters and N sensors are positioned within a square region of unit area, i.e., a 1×1 square. The transmitter localization procedure proposed in this paper extends easily to regions of other shapes and sizes. The locations of the sensor nodes are known but arbitrary. The locations of the transmitters and sensors are represented in Cartesian coordinates. The transmitter locations are given by an (unknown) $M \times 2$ -dimensional vector \mathbf{L}_T , where the i th row of \mathbf{L}_T contains the x and y coordinates of the i th transmitter. Similarly, the locations of the sensors are given by an $N \times 2$ -dimensional vector \mathbf{L}_S . As an example, a diagram of the square region containing a single ($M = 1$) transmitter and $N = 5$ receivers is shown in Figure 2.

For the work presented here, we consider line-of-sight channels from each transmitter to each sensor, and hence we assume a free space path loss model. Power is assumed to decay at a rate inversely proportional to the square of distance traveled according to

$$P_R = P_T(P_0 d)^{-2},$$

where d denotes Euclidean distance between the transmitter and sensor, P_T denotes transmit power, P_R denotes received power, and P_0 is a constant chosen to according to the carrier frequency and antenna structure [8]. For our simulations, the transmitters are assumed to send signals with known, equal transmit powers P_T .

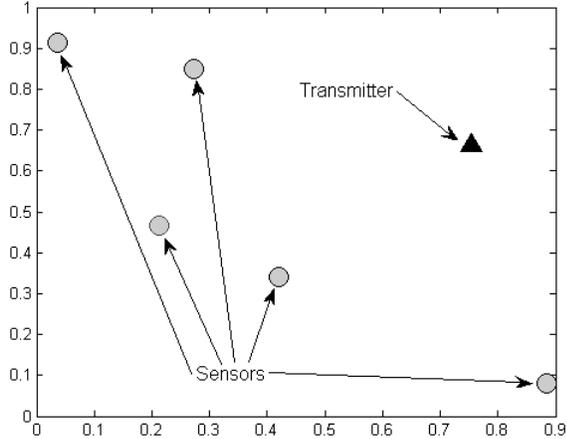


Fig. 2. $M = 1$ transmitter (black triangle) and $N = 5$ sensors (gray circles) arbitrarily placed within a 1×1 square. Power measurements at the sensors are used to locate the transmitter within the unit area.

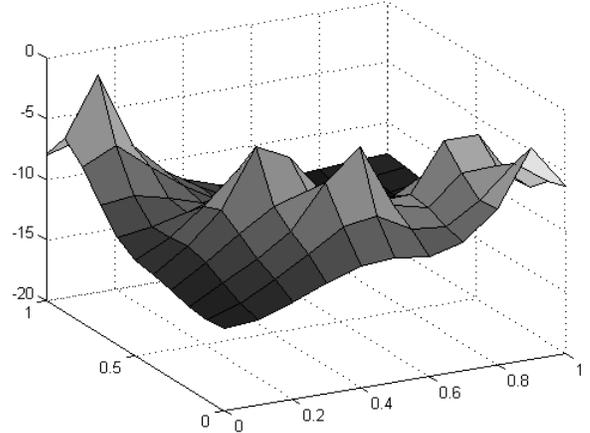


Fig. 3. Surface defined by the cost function (1) for the transmitter and sensor locations given in Figure 2.

We have chosen a simple path loss model, as it facilitates clear presentation of the proposed algorithm and associated optimization function. In practice, one would likely consider a path loss model that accounts for multipath, shadowing, and other effects. More sophisticated models can be incorporated in the proposed technique by appropriately modifying the optimization cost function.

III. TRANSMITTER LOCALIZATION ALGORITHM

We frame the transmitter localization task as a global optimization problem by defining a cost function and searching for the transmitter locations that minimize that function. The cost function is the sum of the squared differences between the true observed power at each sensor and the predicted received power based on the estimated transmitter locations:

$$F(\hat{\mathbf{L}}_T) = \sum_{i=1}^N \left(P_R(i) - \sum_{j=1}^M \frac{P_T}{(P_0 d(i, j))^2} \right)^2, \quad (1)$$

where $\hat{\mathbf{L}}_T$ denotes an $M \times 2$ -dimensional matrix containing the estimated transmitter locations, $P_R(i)$ denotes the observed power at the i th sensor, and $d(i, j)$ denotes the Euclidean distance between the i th sensor and the estimated location of the j th transmitter. The above cost function is not convex with respect to the estimated transmitter locations, even when only a single transmitter is present in the region of interest. As an example, the cost function corresponding to the transmitter and sensor locations in Figure 2 is shown in Figure 3. To avoid local minima solutions, we consider a global optimization algorithm.

It is important to note that the cost function defined represents the transmitter locations in terms of their power contributions to the sensors, but does not measure the error in the transmitter locations directly. This, of course, is impossible

when the transmitter locations are unknown. The correct transmitter locations minimize the error between the estimated and actual power levels at the sensors. However, there may be false minima in the defined objective function where the estimated power levels are close to the measured levels, but the corresponding transmitter locations are incorrect. Figure 4 shows the positions of the transmitters and receivers for one simulated scenario, along with a plot of the error functions for some estimated transmitter locations.

A. Initial Condition Generation via Clustering

The global optimization algorithm used for this work is particle swarm optimization (PSO) [9]. This technique employs a set of particles, or agents, that traverse the space in search of the location at which the cost function achieves its minimum. Initial conditions define the starting locations for each of the agents. Since the nonconvex optimization function will exhibit local minima in which agents may become trapped, the initial locations at which the agents are placed can have a significant impact on the final performance of the optimization algorithms. Seeding the optimization algorithm with good initial conditions increases the likelihood that the area of attraction for the global minimum will be found with few function evaluations.

We propose the use of spatial clustering of the receivers to generate smart initial locations for the agents. To compute these initial estimates, assume a simplified model such that the power observed at each sensor is received entirely from the transmitter that is closest to the sensor. First, spatial clustering divides the receivers based on their location into M clusters. For each cluster, the best estimate of the location of a single transmitter that would produce the power observed at each sensor in the cluster is generated. Note that a cluster must include at least three sensors to fully specify the location of an associated transmitter via trilateration. If there are fewer than three sensors in a cluster, one of the non-unique solutions

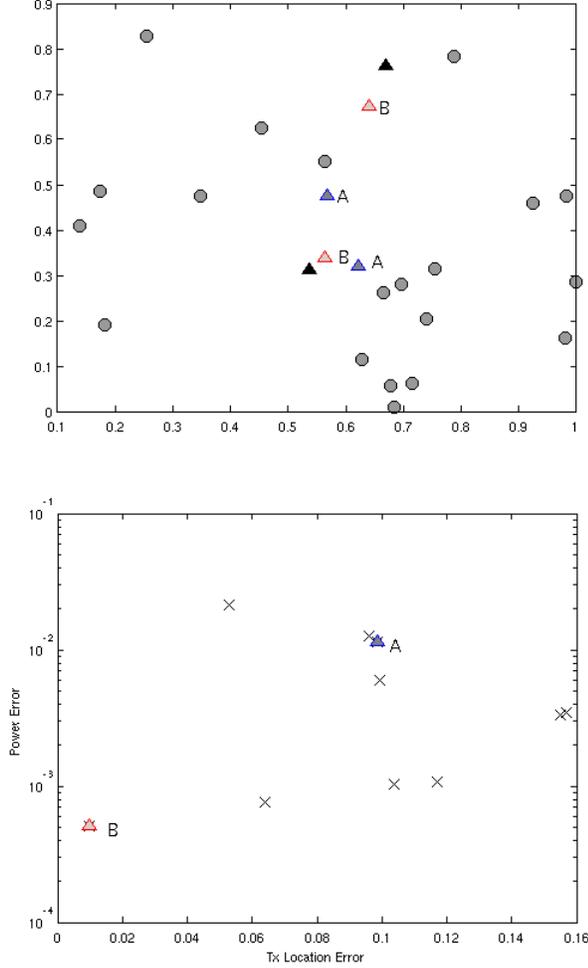


Fig. 4. Top: A simulated scenario is shown, with two transmitters (black triangles) and twenty receivers (gray circles). Two of the smart initial estimates of the transmitter locations are shown (gray triangles) labeled A and B. Bottom: Each of the smart initial estimates is plotted in terms of each estimate's actual transmitter location error and each estimate's power error as measured by the cost function (1). B is a better estimate than A. Some estimates have greater power error than A but lower transmitter location error. This occurs because the cost function is asymmetric, nonconvex, and has local minima. The bottom figure also shows a sample of the quality of the smart initial estimates.

for the transmitter locations is chosen. If more than three sensors are present in a cluster, the location of the associated transmitter will be overdetermined. In such cases, the single transmitter location (x_j, y_j) for the j th cluster is chosen to solve

$$(x_j, y_j) = \arg \min_{x,y} \sum_{i \in I_j} \left(P_R(i) - \frac{P_T}{(P_0 d(i, j))^2} \right)^2,$$

where I_j denotes the set of indices of all sensors in the j th cluster. The M initial transmitter locations generated by the clustering and simplified "single transmitter" model are then used as initial estimates for the global optimization, as described further in III-B.

The spatial clustering uses the k-means clustering algorithm

[10]. Although other clustering methods could be used, k-means clustering can provide multiple clusterings. K-means begins by randomly choosing M initial centroids within the search space and associates each sensor with its nearest centroid, thus performing the first clustering. The centroid associated with each cluster is then recomputed as the spatial mean of the sensors in that cluster. Using the new centroid locations, the step of assigning each sensor to the nearest centroid is repeated. The process continues to iterate until the centroids are stationary, or until a pre-defined number of iterations has been performed. The goal of the k-means algorithm is to minimize the sum of squared distances between sensors and associated centroids, i.e.

$$\min \sum_{j=1}^M \sum_{i \in I_j} d(i, c_j)^2,$$

where $d(i, c_j)$ denotes the Euclidean distance between the i th sensor and the centroid of the j th cluster. Clustering generally has local minima, and thus the k-means algorithm is not guaranteed to find the clustering that minimizes the sum of squared distances. Different clustering minima will occur depending on the initial centroid locations. In this application, the multiple clusterings are an advantage: each different clustering results in a different initial set of estimated transmitter locations. An optimal clustering of the sensors is not needed, because even one iteration of k-means will produce clusters of sensors that are geographically adjacent. To ensure a multiplicity of clusterings, the k-means algorithm is run for only two iterations for each random start.

B. Global Optimization Techniques

In choosing a global optimization algorithm with which to perform transmitter localization, we are conscious of the complicated nature of the cost function (1). We choose to avoid algorithms that require gradients, as such calculations may be costly and misleading. We focus our attention on the popular modern optimization algorithm PSO, in which multiple agents are employed to search the $2M$ -dimensional space of vectors of transmitter locations.

PSO was originally proposed in [9] and has received substantial attention in the global optimization community. PSO operates using a swarm of particles, each of which has a particular initial location and travels within the search space with a velocity (direction and speed) determined by the algorithm. The velocity at which a particle travels is updated according to the best (lowest cost) position visited by the individual particle, as well as the best location visited by any member of the particle swarm. The position of a particle at iteration i is determined as

$$\bar{p}_i = \bar{p}_{i-1} + \bar{v}_i,$$

where \bar{v}_i denotes the velocity of that particle at time i and is defined as

$$\bar{v}_i = c_I \bar{v}_{i-1} + c_C \rho (\bar{p}_{i-1} - \bar{p}_{i-1}) + c_S \rho (\bar{g}_{i-1} - \bar{p}_{i-1}),$$

where \overline{pb}_{i-1} and \overline{gb}_{i-1} denote the best position of the particle at time $i - 1$ and the globally best position (best across the swarm) at time $i - 1$, respectively, and ρ is a uniform random number in the interval $[0, 1]$. The parameters c_C and c_S are termed the cognitive and social scaling parameters, respectively. As one can see in the structure of (III-B), c_C controls the degree to which the individual particle's best position affects its velocity vector, while c_S controls the degree to which the swarm's overall best position affects the velocity vector. The parameter c_I is an inertial weight, and is often set to one. There are a plethora of variations to the PSO algorithm, and any number of successful applications. Further PSO information can be found at [11].

For our application, we employ constriction PSO, an enhancement of PSO in which the velocity of each particle is multiplied by a constriction factor at each update. Recently, Schutte and Groenwold compared a number of PSO variants and recommended constriction particle swarm along with the set of parameters used in this work [12]. The effect of the constriction factor is to reduce the velocity of the particles as the search iterates, thereby increasing the likelihood that the particle will converge to a minimal location. In our PSO implementation, the search is constrained to a known search area by moving any particle that is outside that area to the closest boundary.

To begin the PSO algorithm, each particle must be placed in an initial location. These initial locations could be determined in a random fashion, but simulation results will show that using k-means clustering to generate smart initial locations for the particles yields significant performance improvement. More discussion of the initialization schemes used in this work may be found in IV-A. The initial PSO agents also require a starting velocity. In all cases this initial velocity is assigned randomly.

IV. SIMULATION RESULTS

To evaluate the effectiveness of the proposed transmitter localization technique, we simulate its performance for four to twenty sensors and two to four transmitters. Increasing the number of transmitters and receivers does increase the computational load of the simulation, but, the proposed approach has no inherent limits to the dimensionality it can handle.

A. Initialization Sets

PSO operates with a number of agents, and each agent must have a position from which to start the search. In the absence of other information, the starting positions are generated with a random distribution across the search space. In this paper, random initialization is compared to a variety of initializations seeded with smart positions.

PSO will be started with $(2M + 1)^2$ agents. The recommended number of agents for PSO, according to Schutte [12], is 25, but some results suggest that using $(2M + 1)^2$ agents works as well or better [13]. In order to generate initial condition sets, k-means is run $(2M + 1)^2$ times with random initial centroid locations. This results in multiple sensor clusterings of which only the unique clusters are retained. Transmitter

locations are then estimated for each clustering. As the number of unique clusters is usually smaller than the desired $(2M + 1)^2$ the estimated transmitter locations are augmented to provide the desired number of initial conditions.

Three methods of augmenting the estimated solutions are used in this work. In the first, the smart estimates are simply repeated as many times as necessary to generate the correct number of initial conditions. This set, known as *smart replicated*, provides the correct number of guesses, but the fact that some of the guesses are duplicates is expected to decrease the performance of the global optimizer, even though the replicated initial guesses will begin with different randomly drawn velocities. The second augmentation method, *smart replicated plus noise*, adds noise to the replicated vector:

$$\overline{x}_{SRPN} = \overline{x}_{SR} + \overline{n}_U,$$

where the elements of \overline{n}_U are i.i.d. distributed uniformly between 0 and 0.1. The third augmentation is generated by concatenating the smart locations with a vector of randomly generated locations. This *smart plus random* vector is able to use the good estimates, but also has more variance in the initial locations than the replicated vectors do.

B. Discussion

The performance of the localization techniques is given in terms of the sum of the squared distances between the true transmitter locations and the estimated transmitter locations. The error curves show the average error over 1000 runs of the simulation.

The results in Figure 5 clearly show that the global optimization routine is selecting estimated transmitter locations that are close to the actual transmitter locations. These results also demonstrate the usefulness of including smart estimates in the initial agent settings. The mean errors for the smart initial conditions are significantly lower than those for the randomized initialization. The median error plots show that the optimization performs quite well on most of the scenarios. The mean errors are higher due to occasional outlying solutions. Our hypothesis is that these outliers are due to PSO becoming trapped in local minima. This is supported by the fact that the mean errors for the randomized initialization are higher than those for the smart initializations. The random starting conditions mean that the algorithm may start with no agent close to the true solution, and PSO may never be able to find the global minimum's area of attraction. In contrast, starting PSO with some smart guesses increases the likelihood that an agent is near the global minimum's area of attraction.

The results for localization of $M = 3$ and $M = 4$ transmitters are similar to the results for the two transmitter case and are shown in Figures 6 and 7, respectively. One difference in these results is that the *smart replicated with noise* conditions perform slightly worse. This may be because the scale of the uniform noise is better suited to the characteristics of the objective function for two transmitters than for three or four transmitters. Future work using noisy initial estimates will need to refine the type of noise added to the vector.

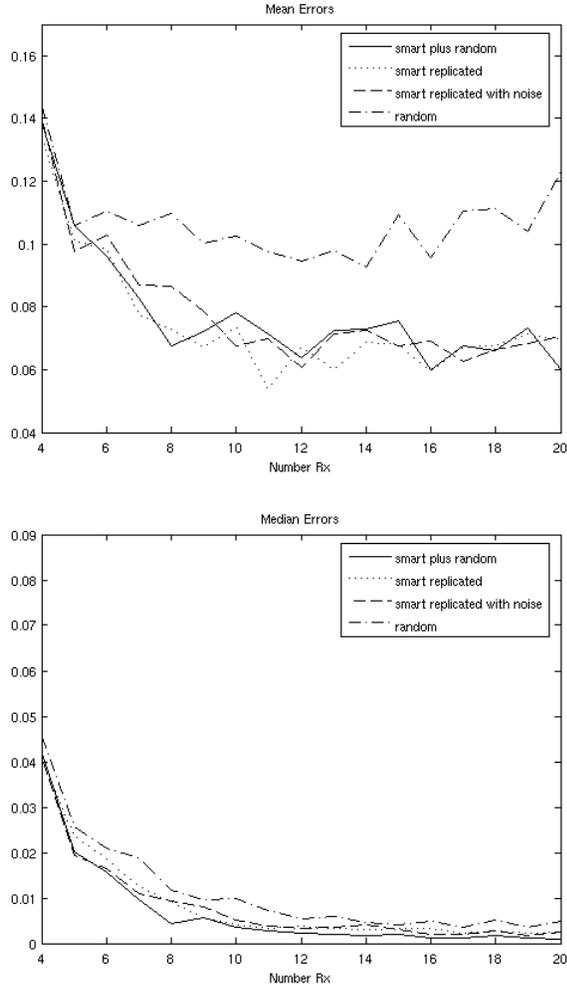


Fig. 5. Mean (top) and median (bottom) Euclidean distance error between true and estimated transmitter locations for 1000 simulations of $M = 2$ transmitters and $N = 4$ to $N = 20$ sensors, or receivers.

Looking at all three figures together allows us to draw some conclusions about the effects of the number of transmitters and receivers on overall performance. As the number of receivers increases we see that performance improves. This is due to the fact that more receivers supply greater information, and thus less ambiguity, about the scenario. It is worth noting that the performance does level off as the number of receivers grows, as is expected due to diminishing returns of adding a new receiver. If two receivers are very close together, the information from the second one will be largely redundant. Finally, if we consider that each receiver affects the objective function for some surrounding diameter, it stands to reason that the search space may be more complicated with a greater number of receivers. It is less likely to have false global minima but perhaps more likely to have a larger number of local minima, rendering the optimization problem more difficult.

As the number of transmitters is increased the performance characteristics degrade somewhat. This is because the

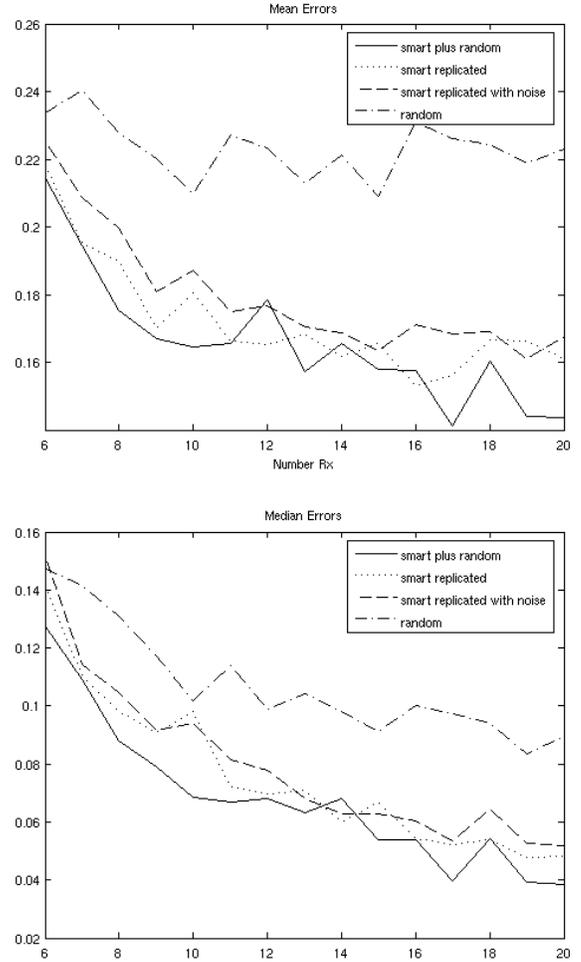


Fig. 6. Mean (top) and median (bottom) Euclidean distance error between true and estimated transmitter locations for 1000 simulations of $M = 3$ transmitters and $N = 6$ to $N = 20$ receivers.

increased dimensionality results in a more difficult global optimization problem. The curse of dimensionality is well known in global optimization research. In our work, an attempt was made to scale the algorithm with dimensionality both by increasing the number of agents and increasing the pre-set number of function evaluations allowed per run. Still, it is difficult to determine whether the algorithm is converging early to a poor solution, or perhaps not converging at all. Future work will address this issue.

V. CONCLUSION

Localization of multiple transmitters within a geographic region based on received powers is a challenging problem that can be solved by a non-convex minimization. We have shown that PSO global optimization of the cost function (1) can yield effective estimates, and that these estimates are substantially improved by using clustering to generate smart initial conditions.

These promising results warrant further exploration of the proposed technique for generating smart initial conditions

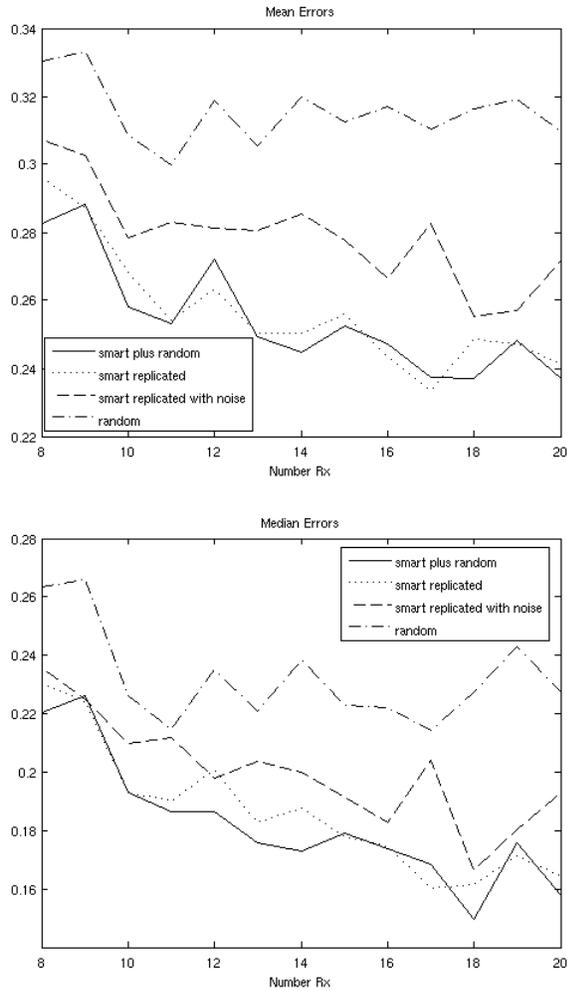


Fig. 7. Mean (top) and median (bottom) Euclidean distance error between true and estimated transmitter locations for 1000 simulations of $M = 4$ transmitters and $N = 8$ to $N = 20$ receivers.

and localizing multiple transmitters. Future work to refine the estimated smart positions and to improve the global optimization may render this technique even more valuable. In this work, a standard PSO variant was used to perform the optimization; while this algorithm is known to perform well on many problems, future work may modify the optimization algorithm to perform better for the transmitter localization problem. This may be particularly important when the number of expected transmitters increases, as most global optimization algorithms, PSO included, falter in high dimensions.

In addition to modifying the optimization technique, we believe it would be beneficial to explore alternative cost functions for the transmitter localization problem. While the chosen cost function reaches a global minimum at the true location of the M transmitters, the function is nonlinear in the transmitter location error, and hence misleading results (such as those shown in Figure 4) may be generated. An alternative to changing the cost function would be to allow the optimizer to return multiple (dissimilar) estimates that appear correct.

The cognitive radio system could then evaluate the possible locations in light of other available information and draw an informed conclusion.

A natural extension of this work is to consider situations in which the number of transmitters and/or the power generated by each transmitter, is unknown. The current objective function is based on comparing an estimated power level to the measured power level for each sensor. In the case where transmitter power is unknown, methods for estimating the power level will be refined. Many of these issues may be addressed with a multi-step global optimization which can select for both transmitter location and transmitter properties.

REFERENCES

- [1] J. Mitola, "Cognitive radio: An integrated agent architecture for software defined radio," PhD Thesis, Royal Institute of Technology (KTH), Stockholm, Sweden, 2005.
- [2] S. Haykin, "Cognitive radio: Brain-empowered wireless communications," *IEEE Journal on Selected Areas in Communications*, vol. 23, no. 2, pp. 201–220, February 2005.
- [3] P. Kolodzy, "Application of cognitive radio technology across the wireless stack," *IEICE Transactions on Communications*, vol. E88-B, no. 11, pp. 4158–4162, November 2005.
- [4] J. Poston, W. Horne, M. Taylor, and F. Zhu, "Ontology-based reasoning for context-aware radios: Insights and findings from prototype development," in *IEEE Symposium on New Frontiers in Dynamic Spectrum Access Networks*, November 2005, pp. 634–637.
- [5] A. Ghasemi and E. Sousa, "Collaborative spectrum sensing for opportunistic access in fading environments," in *IEEE Symposium on New Frontiers in Dynamic Spectrum Access Networks*, November 2005, pp. 131–136.
- [6] G. Ganesan and Y. Li, "Cooperative spectrum sensing in cognitive radio networks," in *IEEE Symposium on New Frontiers in Dynamic Spectrum Access Networks*, November 2005, pp. 137–143.
- [7] D. Cabric, S. Mishra, D. Willkomm, R. Brodersen, and A. Wolisz, "A cognitive radio approach for usage of virtual unlicensed spectrum," in *Proceedings of the 14th IST Mobile and Wireless Communications Summit*, June 2005.
- [8] A. Goldsmith, *Wireless Communications*. New York, NY: Cambridge University Press, 2005.
- [9] R. Eberhart and J. Kennedy, "A new optimizer using particle swarm theory," in *Sixth International Symposium on Micro Machine and Human Science*. IEEE, October 1995, pp. 39–43.
- [10] T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning*. New York: Springer-Verlag, 2001.
- [11] <http://www.swarmintelligence.org/tutorials.php>, Internet Publication.
- [12] J. Schutte and A. Groenwold, "A study of global optimization using particle swarm," *Journal of Global Optimization*, vol. 31, pp. 93–108, 2005.
- [13] M. Hazen and M. Gupta, "A multiresolutional estimated gradient architecture for global optimization," *Proc. IEEE Congress on Evolutionary Computation*, 2006.