ADAPTIVE DISTRIBUTED SENSING FOR EMITTER LOCALIZATION WITH AUTONOMOUS UAV TEAM COOPERATION

Gerald Fudge, Paul Deignan,* Joshua Anderson,* Emanuel Owoye, Paul Scerri,** & Robin Glinton***

* L-3 Communications / Integrated Systems
** School of Computer Science, Carnegie Mellon University

ABSTRACT

An adaptive distributed sensing approach for geolocation of ground-based radio frequency emitters by an autonomous unmanned aircraft system (UAS) is described. The UAS consists of a team of autonomous unmanned aerial vehicles (UAVs) with received signal strength indicator and video sensors under the control of Machinetta intelligent agents. In the presented methodology, each UAV collaboratively estimates a local Bayes filter with a corresponding entropy map for use in a modified forward Rapidly-exploring Random Tree path planner to determine informational optimal non-overlapping routes. The UAV team also collaboratively performs a video search of likely emitter locations. In this paper we present Monte Carlo simulated as well as live flight experimental results and discuss the methodology as it relates to distributed compressed sensing.

1 INTRODUCTION

Geo-location of radio frequency (RF) emitters is a requirement that arises in a number of defense and commercial applications. Traditional approaches to RF emitter geo-location often rely on a small number of precision sensors to measure angle of arrival by triangulation or time difference of arrival. An alternative approach lies in the emerging field of distributed sensor networks [1]. The primary motivation is the potential of fusing information from a large number of spatially distributed sensors that are small, low-power, low-cost, possibly mobile, and equipped with wireless communications. Although the capability of any single sensor may not be significant, the sensors can form an intelligent network with very high aggregate capability.

Problematic aspects of mobile distributed sensor networks include communications limitations (bandwidth, link, network), received signal imperfections (noise, multi-path, etc.), network dynamics (sensor failure, or new sensors entering the network), platform dynamics, and environmental issues (wind, rain, etc.). Size, weight, and power constraints exacerbate the situation. Rosencrantz, Gordon, and Thrun [2] propose a distributed particle filter, but their approach requires a model of what other teammates know in order to choose what to send and does not scale well to large teams because of the required pair-wise information exchanges with neighbors. Beard et al. [3] present a general cooperative UAV framework for coordinated maneuvering, surveillance of known target locations, and boundary patrolling, but [3] does not address RF emitter localization or true adaptive distributed sensing. DeLima et al. [4] describe a UAV team that autonomously locates a group of RF emitters where each UAV broadcasts all sensor readings to all teammates. In order to alleviate the exponential communications cost, the UAVs form sub-teams which each maintain a separate sub-team posterior probability distribution, thus providing optimization within small sub-teams only.

We are interested here in locating RF emitters using extremely low quality sensor information with autonomous UAVs that are cooperating as a team to form an adaptive distributed sensing network. The work described in this paper proceeds from that presented in [5], which focuses on localization of RF emitters using received signal strength indicator (RSSI) sensors. Since RSSI sensors typically provide only an approximate indication of the signal strength and do not provide range information directly, significant spatial diversity is required for emitter localization. Spatial diversity can be achieved through mobility of a team of sensors. The approach in [5] uses a combination of techniques, including a Bayesian estimator [6] maintained by each UAV, local selection of sensors readings to share, entropy map to plan paths through areas of maximum entropy, and a modified rapidly-expanding random tree (RRT) path planner coupled to a cost map which includes entropy, terrain, and paths of other UAVs. The algorithms and receivers are encapsulated within Machinetta intelligent agent proxies [7] so that a variety of sensors, platforms, team strategies, and mission requirements may be incorporated within a unified framework.

2 TEST METHODOLOGY

Testing is performed in a sequence of steps, including algorithm testing in Matlab, medium fidelity UAV simulation, high fidelity OpNet-Machinetta simulation (Figure 1), and live flight tests. Both the medium and high fidelity environments use the same proxy software as with live flights, thus allowing algorithms and most of the software to be de-bugged prior to live flight. We have also been able to use the simulation environment to reproduce unexpected behavior observed during live flight testing. In conjunction with the terrain modeling module DTED data, we use the Terrain Integrated Rough Earth Model (TIREM) propagation delay model. TIREM predicts the RF propagation loss from 1 MHz to 40 GHz, over land and water, and includes the following parameters in simulating RSSI measurements: ground conductivity, relative permittivity, humidity, surface refractivity and resolution (distance between terrain samples.)

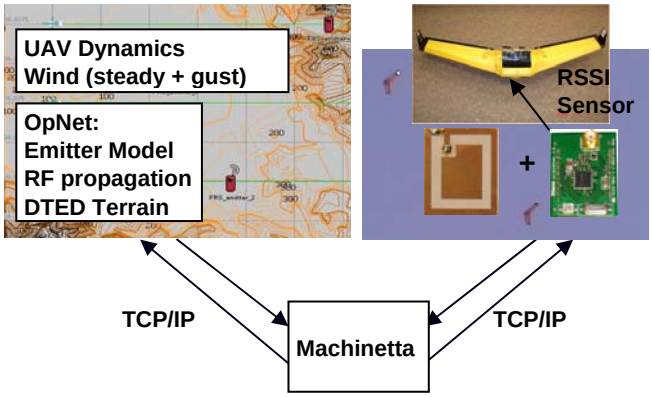

Figure 1: OpNet- Machinetta Co-Simulation Environment

In our current model, the UAV model only includes turn radius since the turn radius is the most significant UAV dynamics parameter based on comparison with live flight tests. Note that this test environment allows testing with a combination of real and simulated UAVs.

3 EMITTER LOCALIZATION

3.1 Joint EO/IR & RSSI Team Behavior

While RSSI sensors are effective at determining the approximate location of RF emitters, for most practical applications, additional sensors, such as EO/IR, are required to positively and accurately identify the source and location of the signal. However, the number of EO/IR images must be kept to a minimum to avoid overloading human operators, bandwidth, and processor capability.

Given a spatial probability distribution over emitter location, we must next determine where to send UAVs with EO/IR sensors. Typically, this distribution will have spatial areas of higher probability surrounded by larger areas of much lower probability. The aim is to find these areas of high probability and send EO/IR UAVs to these locations. While this is a relatively straightforward problem, there are several practical challenges. First, early in the RSSI search, there is more uncertainty so that clustering tends produce overly large clusters, implying a longer search with wasted resources. Waiting for the RSSI process to create small, clearly defined clusters means that it will be longer before emitters are positively identified. Second, due to the nature of the sensors and signals, a single emitter may be split into a group of non-contiguous clusters. Finally, once a cluster has been identified as a likely emitter location, an appropriate UAV needs to be found and tasked to provide imagery. We next describe an initial approach to this problem and simulation results.

The Bayesian estimator represents the posterior probability $P_t^i(e_1 \cdots e_n \mid z_{to}^i \cdots z_t^i)$ over emitter locations by the set $P = \{p_1,....p_{MN}\}$, where p_k is the probability that cell k in an *M x N* grid of the region G contains an emitter.

The first step is to identify cells where the probability of an emitter is above some defined level *b*. The result is a set of cells $P_H \subset P$, where $P_H = \{k \in G : p_k > b\}$. In most cases, P_H will contain groups of adjacent cells that represent the uncertainty about the precise emitter location.

A standard clustering algorithm finds contiguous sets of cells where the probability exceeds the threshold *b*. Small gaps in a group of cells are permitted, allowing small artifacts in the spatial probability distribution. Each such cluster is added to a list, $L_{EO/IR}$ for further processing.

For each cluster in $L_{EO/IR}$, the algorithm checks the entropy in the cluster. If the cluster entropy drops below a threshold γ then a UAV should be sent to provide video of the area. The threshold γ allows a balance between quickly sending EO/IR UAVs to possible emitter locations and use of UAVs for other purposes such as RSSI search.

To actually perform the tasking, we have chosen the LA-DCOP task allocation algorithm [8], which works by creating a *token* representing the responsibility to perform the task. The token is passed around the team until accepted by a proxy agent team member available to perform the task and considering itself sufficiently capable of performing the task for it to be of benefit of the team. The use of such a flexible task allocation algorithm allows the team to have multiple EO/IR UAVs or UAVs with both RSSI and EO/IR sensors.

3.2 Joint Search & Communications Behavior

The Bayes filter may be updated directly through new

measurements taken by sensors under the control of the local agent and indirectly through the communication of measurements from other agents. Both sources of information are practically limited and may have different characteristics so it is useful to condition their ability to update the local Bayesian estimator maps separately. However, since the objective of each agent is to improve its local Bayes estimate efficiently to represent the true state of the environment, there must be a common cost metric for both communicated and direct measurements.

Since the priors used by the agents to optimize their searches and communications are not necessarily in agreement, the objective function could be based on Shannon's information rather than a function of the various posterior emitter probabilities alone. Now, the objective of collaborative emitter searches can be formulated for each agent as the individual maximization of the rate of information gained relative to that agent's local Bayesian estimator map over an indexed two-dimensional region, i.e.

$$\max \sum_{i=1}^{N}\sum_{j=1}^{M} H_{i,j}^{m}(t) + H_{i,j}^{c}(t) - I_{i,j}^{m,c}(t) \quad \text{for } t \in [t_0, T)$$

where $H_{i,j}^{m}(t) = -\left(\left|p_{i,j}^{m}(t) - p_{i,j}^{m}(t_0)\right|\right)\log\left(\left|p_{i,j}^{m}(t) - p_{i,j}^{m}(t_0)\right|\right)$ is the measured entropy difference with respect to the agent's *a priori* probability estimate at time t_0, $H_{ij}^{c}(t)$ is the communicated entropy difference with respect to the transmitting agent's *a priori* estimate at time t_0 , and

$$I_{i,j}^{mc}(t) = -\left(\left|p_{i,j}^{mc}(t) - p_{i,j}^{mc}(t_0)\right|\right)\log\left(\frac{\left|p_{i,j}^{mc}(t) - p_{i,j}^{mc}(t_0)\right|}{\left|p_{i,j}^{m}(t) - p_{i,j}^{m}(t_0)\right|\left|p_{i,j}^{c}(t) - p_{i,j}^{c}(t_0)\right|}\right)$$

is the mutual information difference of the joint measured and communicated agents' *a priori* probability estimate at time t_0. Note that the objective function can be rewritten simply as the maximization of a joint entropy difference from time t_0. However, the formulation given is more practical to estimate for this application. While the objective of the team is to measure and disseminate new information at an optimal rate, the result is that the joint entropy of these maps is minimized as the information is absorbed into the local Bayesian estimates.

In order that learning occurs despite potential communication lags between agents, the *a priori* Bayesian estimator maps are updated relative to a common clock. Therefore all shared measurements must include a timestamp. Measurements that are received with timestamps occurring before the new t_0 are disregarded. This will be rare if measurements are transmitted in near real time and if the time interval of update is significantly greater than the agent's average processing time per update. However, the time interval of update should be fast enough to measure any desired dynamics of the environment. If the agents have unique station identifications to associate with their measurements, it is possible to also allow agents to retransmit measurements from agent to agent without redundancy within the current update interval.

To address bandwidth limitations, measurements can be ordered for transmission according to the contemporaneous differential information of the particular agent which performed the original measurement. Thus, the estimated measured entropy difference can also be included with the measurement packet strictly for network management purposes. The rule of thumb here is that measurements that are locally important to an agent are more likely to be important in general.

4 TEST RESULTS

Prior test results [5] focused on the RSSI-only case and examined information sharing & agent messaging efficiency, probability of collision, Kullback Leibler divergence used in the cost function, and limited live flight test results. This section presents the joint EO/IR & RSSI case and includes additional live flight test results.

4.1 Joint EO/IR & RSSI Simulation Results

In this section, we present simulation results of the joint EO/IR & RSSI behavior approach described earlier. In these simulations, the RSSI sensor range is 2500m to 5000m and the area of interest is 50km by 50 km. Each scenario also includes an EO/IR UAV which is tasked to provide video footage of the emitters once the cluster is small enough. Each video (or picture) task order contains the coordinates of the believed location of the emitter and these coordinates are compared against the true location to estimate average error and average distance traveled by the RSSI UAV up to that point. Any picture orders outside a 1000m of the true location are considered false positives. The simulator and proxies are spread out over up to 15 desktop computers and communication is via multi-cast UDP resulting in around 3% message loss. Figure 2 shows the average distance from the picture order to the true emitter location. Note that increasing the number of RSSI UAVs improves the accuracy and increasing the number of emitters decreases the accuracy from ambiguity in the received signal due to overlap. Figure 3 shows the average distance traveled per RSSI UAV and Figure 4 shows the total distance traveled by all the RSSI UAVs before all emitters were found. In the 20 UAV cases the UAVs traveled less distance each, but more in total than in the 10 UAV case, suggesting a speed/fuel tradeoff when determining how many UAVs to deploy.

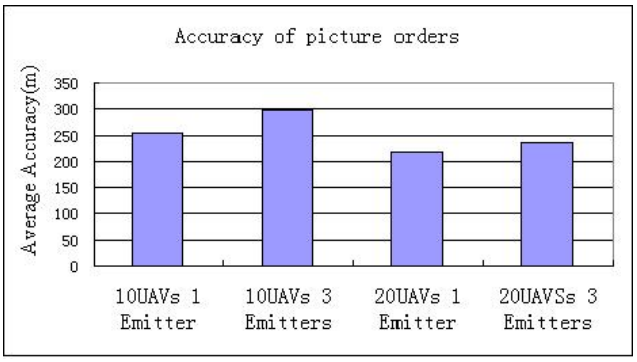

Figure 2: Average Accuracy of Video Tasking Coordinates

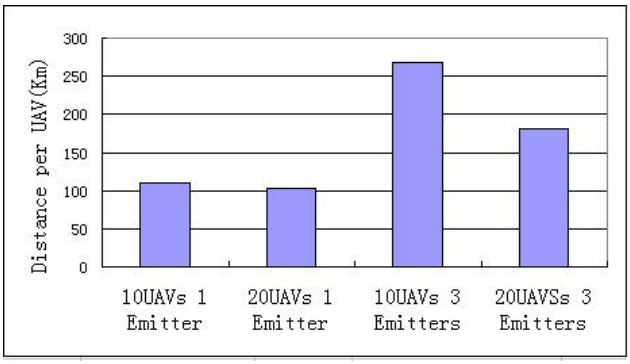

Figure 3: Distance Traveled per RSSI UAV.

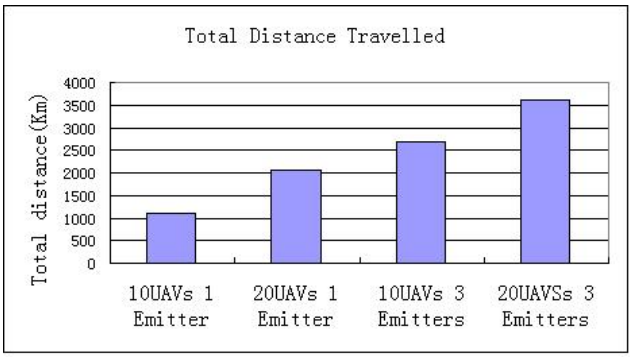

Figure 4: Total Distance Traveled by UAVs.

In the experiments described above, we observed that in some cases, the entropy of a cluster was slightly too high to trigger an EO/IR UAV to go to that cluster. However, the RSSI UAVs did not travel to that area to reduce the entropy because there were other parts of the environment where there is substantially more entropy. This is particularly problematic behavior when there is a high entropy cluster in the middle of a large area where there is very unlikely to be an emitter, since there is very low incentive for a UAV to revisit that area. As an experiment, we artificially manipulated the entropy maps to provide a very large incentive for a UAV to visit a cluster with slightly too much entropy. The result is shown in Figure 5. Notice that the time to find a single emitter was substantially reduced, if this incentive is provided, but the impact dramatically decreases as more emitters are added to the environment.

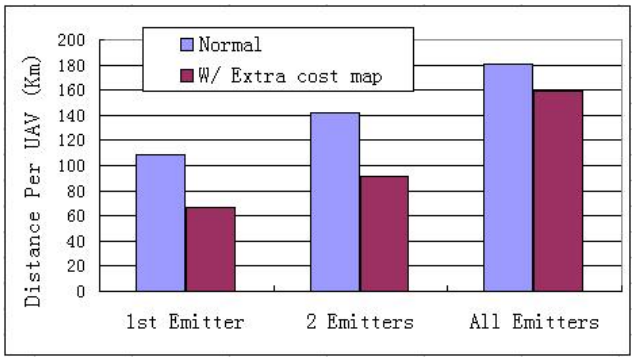

Figure 5: Comparison of Performance With and Without Additional Incentive to Reduce Cluster Entropy.

4.2 Live Flight Test Results

Figure 6 below illustrates an typical UAV flight path and RSSI values from one of the live flights conducted on 13 December, 2006. In this test, the emitter was a Family Radio Service (FRS) radio transmitting at 462.6375 MHz (channel 4) located on the south edge of the area of interest. The search area is roughly 1km by 1km. The signal strength was set to achieve a detection radius of about 50m from an altitude of about 50m – note that the live flight results can be scaled to a larger search area with stronger signals. The UAV path is shown in red. The RSSI measurements are shown on top of the UAV path, with both height and color (blue = weaker signal, red = stronger signal) to indicate the RSSI values. Note that the RSSI values are not gathered continuously because of data link limitations. Also, note that the RSSI values are very noisy, with some stronger measurements occurring farther from the emitter than some of the weaker measurements. Figure 7 shows the corresponding Bayesian estimator map for this flight. Although the map has gaps, enough information to start the EO/IR search is available at about half-way through the flight test. Figure 8 shows another view of the RSSI measurements as a function of range to the emitter. Note that the RSSI measurements have an extremely large variance as a function of range – this is caused by a combination of factors, including noise, multi-path, and relative aspect of the antenna with respect to the emitter. Despite this large variance, the Bayes estimate converges successfully to the true emitter location in live flight tests.

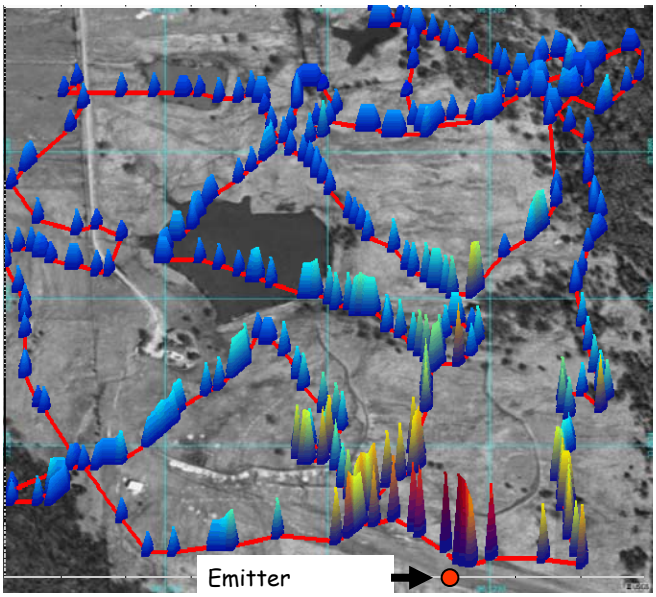

Figure 6: Live Flight RSSI Measurements Overlaid With Imagery, UAV Path, & Emitter Location.

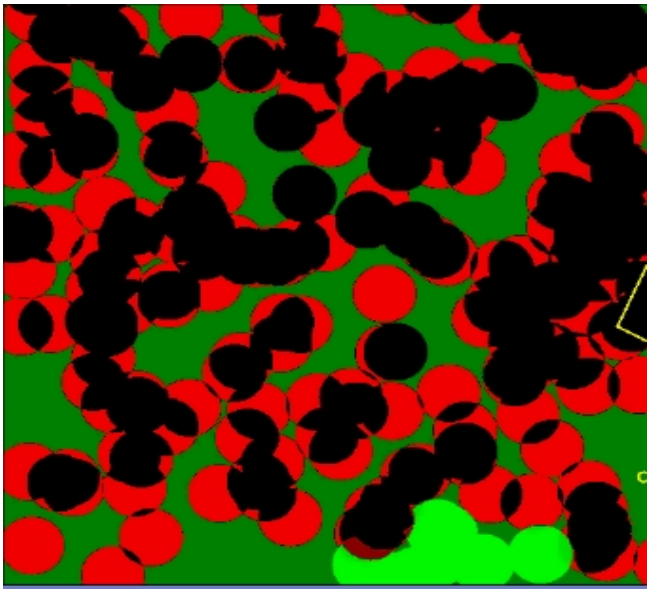

Figure 7: Live Flight Bayesian Estimator Map.

The live flight test results agree closely with the simulations; in fact, the high fidelity simulation has proven accurate enough to evaluate and predict behavior observed during live flight tests. Based on this, we have confidence in the performance predicted by the Monte Carlo trials.

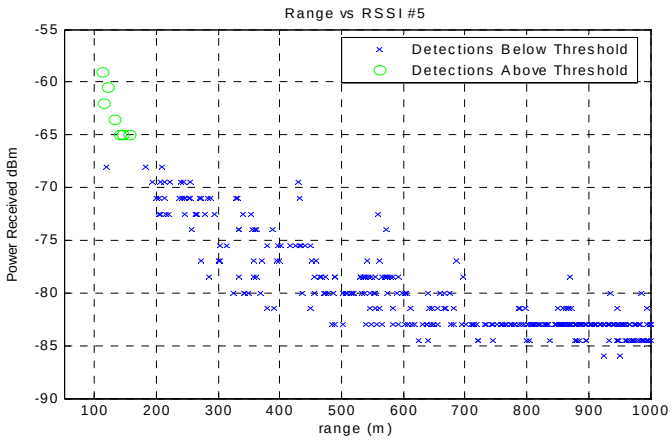

Figure 8: Live Flight RSSI Measurements as a Function of Distance From Emitter.

5 FUTURE RESEARCH & CONCLUSIONS

One phenomenon we have observed during high fidelity UAV simulations as well as live flights is that if a UAV group is simultaneously launched from the same location, the UAVs will occasionally head towards the same local area within the search region. Although they eventually spread out, this clearly is inefficient. Since it is not always practical to launch from geographically diverse points, this issue needs to be addressed. While path planning can be formulated for each agent over an update period through maximization of the expected measured entropy difference between the current *a priori* probability map and an estimated probability map of grids to be potentially visited, it is important to also estimate the new information expected to be gained from communication with other agents in order to prevent a suboptimal "group think" condition. Thus an agent must discount the expected differential communicated entropy from other agents in calculating its path so that the agents will automatically seek to explore diverse regions within the search area while potentially having very similar local probability maps.

The area of compressive sensing (CS) [9][10][11] has recently emerged based on the idea that a signal with a sparse representation in a basis Ψ can be recovered from a relatively small number of projections onto a basis set Φ that is incoherent with the first basis Ψ, even in the presence of noise; furthermore if the projections are random, then the incoherence property is generally met. Since the true emitter signal strength can be modeled as a smoothly varying spatial function with noise, the sparsity requirement is met. With the spatial random sampling provided by the modified RRT, it should be possible to localize the emitters with a reduced measurement set. A related area is active

learning [12], where the sampling process adapts based on the measurements to improve the information per sample, and to thus allow improved performance over CS alone [13]. Applying existing results of CS and active learning theory, one can show that the Bayes filter / entropy map / RRT approach is efficient. In addition, performance bounds can be formally developed using CS and active learning theory. Figure 9 illustrates a three emitter simulation. The left-hand plot shows the true RSSI intensity map and a UAV path driven by the RRT path planner with measurement points. The right-hand plot shows the reconstructed map using the discrete cosine transform dictionary with the Stanford CS software library [14][15], with the true emitter locations marked by triangles. This example is not definitive, but serves to illustrate the potential of CS applied to distributed sensing.

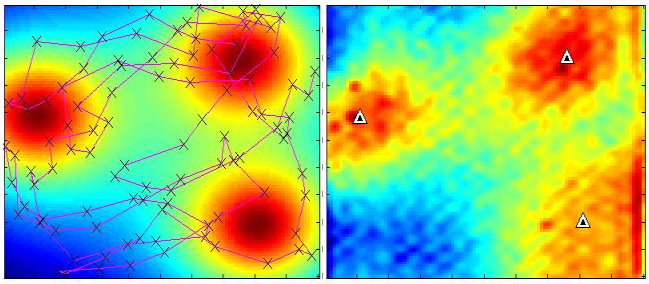

Figure 9: RSSI Intensity Map With UAV Path & Measurement Points (Left); Reconstruction (Right)

We have presented an adaptive distributed sensing approach for locating RF emitters using RSSI and EO/IR UAVs cooperating autonomously as a team. Simulations and live flight experiments show the effectiveness of the overall approach. Future research areas include application of compressive sensing and adaptive learning techniques, automatic video processing, and faster convergence of the emitter locations.

Acknowledgements: The authors would like to thank the following L-3 Communications research collaborators on this effort: Dan Rutherford, Ken Stroud, Jon Brown, Alex Yeh, Tracy Von Gonten, Michael Custer, Scott Burkart, Vaughan Bowman, and David Reid.

6 REFERENCES

[1] *IEEE S.P. Magazine*, Special Issue on Distributed Signal Processing in Sensor Networks, July 2006.

[2] M. Rosencrantz, G. Gordon, and S. Thrun, "Decentralized sensor fusion with distributed particle filters," *Conference on Uncertainty in AI* (UAI), 2003.

[3] R. Beard, T. Mclain, D. Nelson, and D. Kingston, "Decentralized cooperative aerial surveillance using fixed-wing miniature UAVs," *IEEE Proceedings: Special Issue on Multi-Robot Systems*, July 2006.

[4] P. DeLima, G. York, and D. Pack, "Localization of ground targets using a flying sensor network," *IEEE Intl. Conf. on Sensor Networks, Ubiquitous, and Trustworthy Computing* (SUTC 1), pp. 194–199, 2006.

[5] Paul Scerri, Robin Glinton, Sean Owens, Katia Sycara, "Geolocation of RF Emitters with Teams of UAVs," *7th International Conference on Cooperative Control and Optimization*, January 31 – February 2, 2007.

[6] H. P. Moravec, "Sensor fusion in certainty grids for mobile robots," AI Magazine, 9(2):71-74, 1998.

[7] P. Scerri, D. V. Pynadath, L. Johnson, P. Rosenbloom, N. Schurr, M. Si, M. Tambe, "A Prototype Infrastructure for Distributed Robot-Agent-Person Teams," *2nd International Joint Conference on Autonomous Agents and Multiagent Systems*, 2003.

[8] P. Scerri, A. Farinelli, S. Okamoto, M. Tambe, "Allocating Tasks in Extreme Teams," *AAMAS*'05, 2005.

[9] D.L Donoho, "Compressed sensing," *IEEE Transactions on Information Theory*, April 2006

[10] E. J. Candes, T. Tao, "Near-optimal signal recovery from random projections: universal encoding strategies?," *IEEE Trans. Info. Theory*, Dec. 2006

[11] Marco F. Duarte, Michael B. Wakin, Dror Baron, and Richard G. Baraniuk, "Universal Distributed Sensing via Random Projections," Proceedings of IPSN'06, April 2006

[12] R. Castro and R. Nowak, "Upper and Lower Bounds for Active Learning", in Proceedings of the 44th Annual Allerton Conference on Communication, Control and Computing, Allerton House, University of Illinois, September 2006

[13] R. Castro, J. Haupt, and R. Nowak, "Compressive sensing vs. active learning," in *Proc. IEEE Intl. Conf. on Acoustics, Speech, and Signal Processing* (ICASSP), May 2006

[14] Scott Chen, David Donoho , Michael Saunders, "Atomizer for Matlab5.x," Stanford University, http://www-stat.stanford.edu/~atomizer/

[15] David Donoho, Mark Duncan, Xiaoming Huo, Ofer Levi, "Wavelab802 for Matlab5.x,", Stanford University, http://www-stat.stanford.edu/%7Ewavelab/