Locating RF Emitters with Large UAV Teams

Paul Scerri, Robin Glinton, Sean Owens, Steven Okamoto and Katia Sycara
School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213, USA
{pscerri, rglinton, owens, sokamoto, katia}@cs.cmu.edu

Keywords

Algorithms, Unmanned Aerial Vehicles

1. INTRODUCTION

The rapidly improving availability of small, unmanned aerial vehicles (UAVs) and their ever reducing cost is leading to considerable interest in multi-UAV applications. However, while UAVs have become smaller and cheaper, there is a lack of sensors that are light, small and power efficient enough to be used on a small UAV yet are capable of taking useful measurements of objects often several hundred metres below them. Static or video cameras are one option, however image processing normally requires human input or at least computationally intensive offboard processing, restricting their applicability to very small UAV teams. In this paper, we look at how teams of UAVs can use very small Relative Signal Strength Indicator (RSSI) sensors whose only capability is to detect the approximate strength of a Radio Frequency (RF) signal, to search for and accurately locate such sources. RSSI sensors give at most an approximate range to an RF emitter and will be misleading when signals overlap. Applications of such UAV teams range from finding lost hikers or skiers carrying small RF beacons to military reconnaissance operations. Moreover, the core techniques have a wider applicability to a range of robotic teams that rely on highly uncertain sensors, e.g., search and rescue in disaster environments.

Many of the key technogies required to build a UAV team for multi-UAV applications have been developed and are reasonably mature and effective[1, 2]. However, for large UAV teams with very noisy sensors, key problems remain, specifically, much previous work is formally grounded but impractical[3]. Often the coordination and planning algorithms and the representations of the environment are not appropriate for more than two or three UAVs and targets. For example, some solutions require a UAV to know the planned paths of all other UAVs in order to plan its own path[6], but this is infeasible (both in terms of communication and computation) when the number of UAVs is large. Other approaches only solve part of the problem, e.g., estimating locations from sensor readings[10] or planning cooperative paths[9], but do not combine these elements in an integrated solution, although there are some exceptions[4]. Signal processing techniques for creating probability distributions from noisy signals have been extensively studied, but rarely have distributed filters versions been created and those that have been do not scale to larger teams[7].

Our approach to this problem has three key elements. The first key element is a distributed filter to localize RF emitters in the environment. Each UAV has a Binary Bayesian Grid Filter[5] where a value of a cell in the grid represents the probability that there is an emitter in the corresponding location on the ground. Due to limitations on available communication bandwidth, it is infeasible for UAVs to share their entire distribution, instead they share a small subset of their sensor readings with others in the team. Hence, departing from previous approaches that elicited a model of what team mates know in order to choose what to send[7], we started from the assumption that if some information leads to large local information gain, it will probbably do so for much of the team. We investigated two information gain based heuristics for choosing which readings to share with team mates. The first heuristic is to send sensor readings that have the greatest impact on the UAV's local probability distribution. The second heuristic is to create a parallel probability distribution based purely on readings received from team mates and send sensor readings that have the biggest impact on that distribution. Intuitively, the first heuristic sends readings that were most important for the local UAV, while the second sends sensor readings that are most important to the team, given a local model of what the team knows. Experiments show that the first heuristic results in better team behavior than sending random messages, but the second heuristic performs worse than random.

The second element of the approach is to tightly couple estimates of the current locations of the emitters to the UAV path planning process. Specifically, a probability distribution over emitter locations is translated into a map of the information *entropy* in the environment. UAVs plan paths through areas of maximum entropy, hence maximizing expected information gain. The UAVs plan only a relatively short distance ahead in each planning cycle. This approach allows the UAVs to be reactive to new information, which is critical when sensors are highly uncertain and the domain is dynamic. For example, if a UAV traverses an area, but the sensor readings do not provide an accurate picture of that

AAMAS'07 May 14–18 2007, Honolulu, Hawai'i, USA.

area, the entropy will remain high and the UAV will consider retraversing the area. Notice that the entropy map coupled with the path planner looking to maximize information gain provides an integrated way for trading off between going to the locations where there will be most information gain and locations that can be quickly reached.

The third key element of the approach is a very lightweight, computationally inexpensive method for cooperative path planning. The important application feature underlying the approach is that due to the high uncertainty and dynamicism in the environment, some overlap of paths is acceptable (or even desirable), provided that the UAVs mainly spread out and search areas of maximum entropy. Our approach is for each UAV to share its planned path with some other members of the team. When planning, each UAV estimates the change in entropy that would be induced by those paths being flown by others and plans on the resulting entropy map. If the most current path of a particular UAV is not known the most recent location is used to roughly estimate where that UAV might be searching.

1.1 Implementation

The overall, integrated process aims to balance the desire to have a principled, formally grounded approach, yet be lightweight and robust enough to be practical for a team of UAVs. The hardware independent components (planners, filters, etc.) are isolated from the hardware specific components (sensor drivers, autopilot) to allow the approach to be quickly integrated with different UAVs or moved from simulation to physical UAVs. The hardware independent components are encapsulated in a *proxy* which will either be on the physcial UAV or on a UAV ground station, depending on the vehicle. In the experiments below, *exactly* the same proxy code is used in simulation as will be used in tests with physical UAVs. Figure 1 shows the main components and information flows from the perspective of one UAV-proxy.

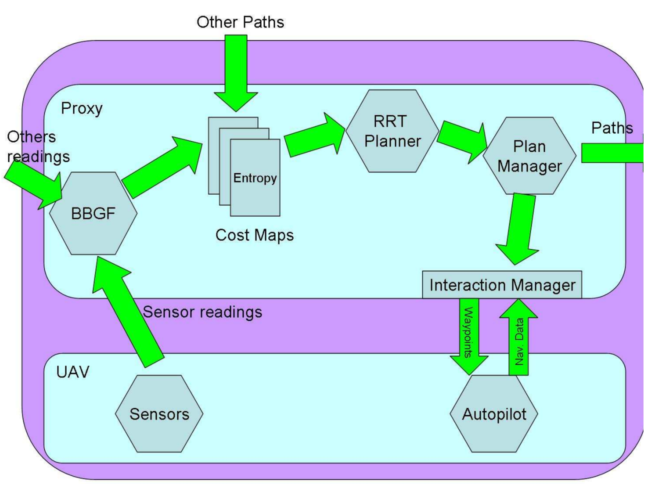

Figure 1: Block diagram of architecture.

2. EXPERIMENTS

In this section, we present empirical simulation results of the approach described above. The approach is implemented with the Machinetta proxy[8] framework integrated with the Sanjaya UAV simulation environemnt. The signal model is derived from real data from an RSSI sensor flown on a real UAV. The code is used is exactly the same code as being used in an ongoing flight test, with the exception of the code between the proxy and the autopilot. The simulated environment is 50km by 50km and the results below represent several hundred hours of simulated flying time, with each data point an average of five runs. Unless otherwise stated there were four RF emitters in the environment. 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. These experiments are conducted in simulation due to the practical difficulty of conducting experiments with large numbers of physical UAVs. With an industrial partner this approach was validated with four physical UAVs in a series of tests in late 2006 and early 2007.

Information Sharing Experiments.

In the first experiment, we looked at the three different information sharing heuristics. Figure 2 shows the average KL-divergence from the ground truth over time. Ground truth is modeled as tight $\frac{1}{r^2}$ probability around the real emitter location. The figure shows that all the information sharing algorithms were effective at determining the location of the emitters, but that H_LOCAL_KL was substantially better than the other heuristics. Interestingly, sending random sensor readings, H_RAND, around the team was clearly better than H_TEAM_KL, sending readings according to a model of the team. Figure 3 gives one possible reason for this, i.e., that H_TEAM_KL sent very few readings around. H_LOCAL_KL gives a low number of messages along with its good KL-divergence, showing it to be clearly the best heuristic.

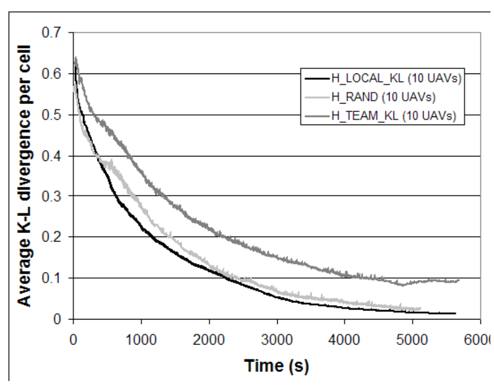

Figure 2: The KL-divergence over time for three different information sharing algorithms.

Number of UAVs and Number of Emitters.

The second experiment varied both the number of emitters and number of UAVs in the environment. Figure 4 shows that more UAVs led to a faster decrease in the KL-divergence, showing that the additional UAVs were useful. Interestingly, more UAVs actually made reducing the KL-divergence faster. We hypothesize that this was because the UAVs were able to take use the additional signals in the

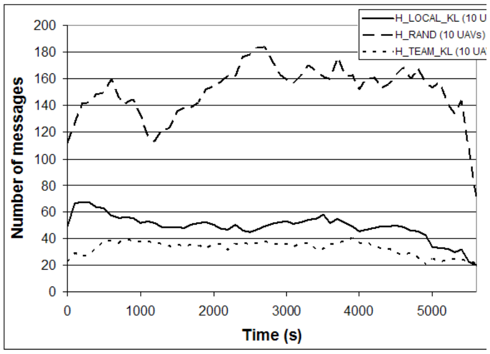

Figure 3: The number of messages sent between UAVs for three different information sharing algorithms.

environment to quickly identify RF emitter locations.

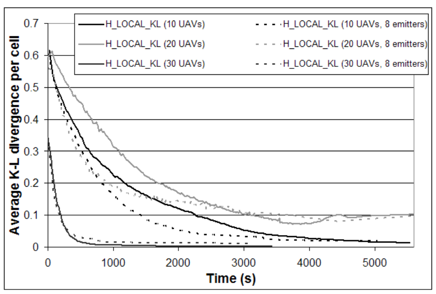

Figure 4: The impact on KL-divergence of changing the number of UAVs and the number of RF emitters.

Intermittent Signals.

The final experiment varied how often the RF emitters were giving off a signal that could be detected, see Figure 5. The four emitters had periods ranging from 5 seconds to 30 minutes, then the percentage of that period that they were on for was varied from 25% to 100%. Curiously, the KL-divergence appears better when the emitter is *off* more. However, this is only due to a quirky interaction between the KL-divergence measurement and the very noisy sensors. Specifically, the noisy sensors do not allow the UAVs to very precisely locate the emitters, so believing that they were not there at all could actually lower the KL-divergence. Figure 6 shows an oscillation in the number of messages sent between UAVs as emitters turn on and off.

Acknowledgements

This research has been sponsored in part by AFOSR FA9550-07-1-0039, AFOSR FA9620-01-1-0542, L3-Communications (4500257512) and NSF ITR IIS-0205526.

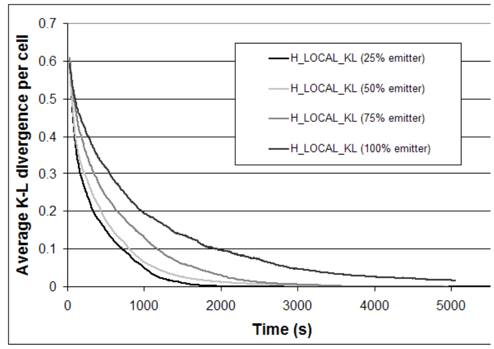

Figure 5: The KL-divergence over time for four different percentages of time the RF emitters emit.

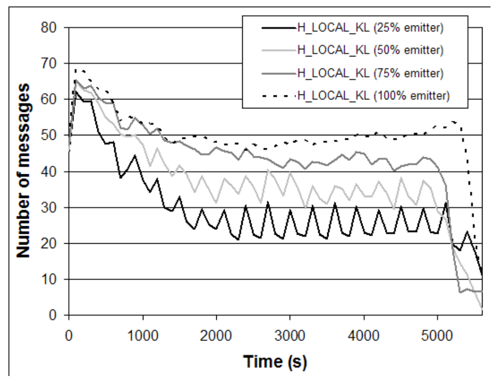

Figure 6: The number of messages over time, for different levels of intermittency.

3. REFERENCES

[1] 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*, to appear.

[2] L. Bertuccelli and J. How. Search for dynamic targets with uncertain probability maps. In *IEEE American Control Conference*, 2006.

[3] S. Butenko, R. Murphey, and P. Pardalos, editors. *Recent Developments in Cooperative Control and Optimization*. Springer, 2003.

[4] P. DeLima, G. York, and D. Pack. Localization of ground targets using a flying sensor network. In *SUTC (1)*, pages 194–199, 2006.

[5] H. P. Moravec. Sensor fusion in certainty grids for mobile robots. *AI Magazine*, 9(2):61–74, 1998.

[6] D. Pack, G. York, and G. Toussaint. Localizing mobile RF targets using multiple unmanned aerial vehicles with heterogeneous sensing capabilities. In *IEEE International Conference on Networking, Sensing, and Control*, 2005.

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

[8] P. Scerri, D. V. Pynadath, L. Johnson, P. Rosenbloom, N. Schurr, M. Si, and M. Tambe. A prototype infrastructure for distributed robot-agent-person teams. In *The Second International Joint Conference on Autonomous Agents and Multiagent Systems*, 2003.

[9] T. Schouwenaars, J. P. How, and E. Feron. Multi-vehicle path planning for non-line of sight communication. In *IEEE American Control Conference*, 2006.

[10] M. Wax. *Detection and estimation of superimposed signals*. PhD thesis, Stanford Univ, 1985.