

Michael J. Hirsch
Panos M. Pardalos
Robert Murphey
Don Grundel (Eds.)

Advances in Cooperative Control and Optimization

Proceedings of the
7th International Conference on
Cooperative Control and Optimization



Springer

0232-53
C778
2007

Michael J. Hirsch, Panos M. Pardalos,
Robert Murphey, Don Grundel (Eds.)

Advances in Cooperative Control and Optimization

Proceedings of the 7th International Conference
on Cooperative Control and Optimization



 Springer



E2007003629

Series Advisory Board

F. Allgöwer, P. Fleming, P. Kokotovic,
A.B. Kurzhanski, H. Kwakernaak,
A. Rantzer, J.N. Tsitsiklis

Editors

Dr. Michael J. Hirsch

Raytheon, Inc.
P.O. Box 12248
St. Petersburg, FL 33733 USA
Email: Michael_J_Hirsch@Raytheon.com

Dr. Robert Murphey

GNC Branch, Munitions Directorate
AFRL
101 W. Eglin Blvd., Ste. 331
Eglin AFB, FL 32542 USA
Email: Murphey@eglin.af.mil

Dr. Panos M. Pardalos

Department of Industrial and Systems Engineering
University of Florida
P.O. Box 116596
Gainesville, FL 32611 USA
Email: Pardalos@ufl.edu

Dr. Don Grundel

AAC/ENA
101 W. Eglin Blvd., Ste. 383
Eglin AFB, FL 32542 USA
Email: Don.Grundel@eglin.af.mil

Library of Congress Control Number: 2007932965

ISSN print edition: 0170-8643

ISSN electronic edition: 1610-7411

ISBN-10 3-540-74354-5 Springer Berlin Heidelberg New York

ISBN-13 978-3-540-74354-5 Springer Berlin Heidelberg New York

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilm or in any other way, and storage in data banks. Duplication of this publication or parts thereof is permitted only under the provisions of the German Copyright Law of September 9, 1965, in its current version, and permission for use must always be obtained from Springer. Violations are liable for prosecution under the German Copyright Law.

Springer is a part of Springer Science+Business Media
springer.com
© Springer-Verlag Berlin Heidelberg 2007

The use of general descriptive names, registered names, trademarks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

Typesetting: by the authors and SPS using a Springer L^AT_EX macro package

Printed on acid-free paper SPIN: 11864295 89/SPS 5 4 3 2 1 0

Lecture Notes in Control and Information Sciences 369

Editors: M. Thoma, M. Morari

Preface

Optimal control of cooperative systems continues to be at the forefront of research initiatives in the military sciences. Recently, cooperative system research has expanded from the military domain to other engineering disciplines, including drug design and disaster recovery. While there exist many powerful techniques for optimal cooperative control problems, this area is still considered one of the most difficult in the applied sciences. Thus, there must be continual improvements and new insight directed to the modeling and analysis of optimal cooperative control problems. This present volume, as well as volumes from previous years, clearly illustrate novel solutions from some of the best and brightest optimal cooperative control researchers.

This volume represents the most recent in a series of publications discussing recent research and challenges in the field of optimal cooperative control. Most of the chapters in this book were presented at the Seventh International Conference on Cooperative Control and Optimization, which took place in Gainesville, Florida, January 31 – February 2, 2007. It is our belief that this book will be an invaluable resource to faculty, researchers, and students in the fields of optimization, control theory, computer science, and applied mathematics.

We gratefully acknowledge the financial support of the Air Force Research Laboratory, The Center for Applied Optimization at The University of Florida, and Raytheon, Inc. We thank the contributing authors, the anonymous referees, and Springer Publishing for making the conference so successful and the publication of this book possible.

Michael J. Hirsch
Panos M. Pardalos
Robert Murphey
Don Grundel

Lecture Notes in Control and Information Sciences

Edited by M. Thoma, M. Morari

Further volumes of this series can be found on our homepage:
springer.com

Vol. 369: Hirsch M.J.; Pardalos P.M.;
Murphey R.; Grundel D.
Advances in Cooperative Control and
Optimization
423 p. 2007 [978-3-540-74354-5]

Vol. 368: Chee F.; Fernando T.
Closed-Loop Control of Blood Glucose
157 p. 2007 [978-3-540-74030-8]

Vol. 367: Turner M.C.; Bates D.G. (Eds.)
Mathematical Methods for Robust and Nonlinear
Control
444 p. 2007 [978-1-84800-024-7]

Vol. 366: Bullo F.; Fujimoto K. (Eds.)
Lagrangian and Hamiltonian Methods for
Nonlinear Control 2006
398 p. 2007 [978-3-540-73889-3]

Vol. 365: Bates D.; Hagström M. (Eds.)
Nonlinear Analysis and Synthesis Techniques for
Aircraft Control
360 p. 2007 [978-3-540-73718-6]

Vol. 364: Chiasso A.; Ferrante A.;
Pinzoni S. (Eds.)
Modeling, Estimation and Control
356 p. 2007 [978-3-540-73569-4]

Vol. 363: Besançon G. (Ed.)
Nonlinear Observers and Applications
224 p. 2007 [978-3-540-73502-1]

Vol. 362: Tam T.-J.; Chen S.-B.;
Zhou C. (Eds.)
Robotic Welding, Intelligence and Automation
562 p. 2007 [978-3-540-73373-7]

Vol. 361: Méndez-Acosta H.O.; Femat R.;
González-Álvarez V. (Eds.):
Selected Topics in Dynamics and Control of
Chemical and Biological Processes
320 p. 2007 [978-3-540-73187-0]

Vol. 360: Kozłowski K. (Ed.)
Robot Motion and Control 2007
452 p. 2007 [978-1-84628-973-6]

Vol. 359: Christophersen F.J.
Optimal Control of Constrained
Piecewise Affine Systems
190 p. 2007 [978-3-540-72700-2]

Vol. 358: Findeisen R.; Allgöwer
F.; Biegler L.T. (Eds.): Assessment and Future Di-
rections of Nonlinear
Model Predictive Control
642 p. 2007 [978-3-540-72698-2]

Vol. 357: Queinnec I.; Tarbouriech
S.; Garcia G.; Niculescu S.-I. (Eds.):
Biology and Control Theory: Current Challenges
589 p. 2007 [978-3-540-71987-8]

Vol. 356: Karatkevich A.:
Dynamic Analysis of Petri Net-Based Discrete
Systems
166 p. 2007 [978-3-540-71464-4]

Vol. 355: Zhang H.; Xie L.:
Control and Estimation of Systems with
Input/Output Delays
213 p. 2007 [978-3-540-71118-6]

Vol. 354: Witczak M.:
Modelling and Estimation Strategies for Fault
Diagnosis of Non-Linear Systems
215 p. 2007 [978-3-540-71114-8]

Vol. 353: Bonivento C.; Isidori A.; Marconi L.;
Rossi C. (Eds.)
Advances in Control Theory and Applications
305 p. 2007 [978-3-540-70700-4]

Vol. 352: Chiasson, J.; Loiseau, J.J. (Eds.)
Applications of Time Delay Systems
358 p. 2007 [978-3-540-49555-0]

Vol. 351: Lin, C.; Wang, Q.-G.; Lee, T.H.; He, Y.
LMI Approach to Analysis and Control of
Takagi-Sugeno Fuzzy Systems with Time Delay
204 p. 2007 [978-3-540-49552-9]

Vol. 350: Bandyopadhyay, B.; Manjunath, T.C.;
Umapathy, M.
Modeling, Control and Implementation of Smart
Structures 250 p. 2007 [978-3-540-48393-9]

Vol. 349: Rogers, E.T.A.; Galkowski, K.;
Owens, D.H.
Control Systems Theory
and Applications for Linear
Repetitive Processes 482
p. 2007 [978-3-540-42663-9]

Vol. 347: Assawinchaichote, W.; Nguang, K.S.;
Shi P.
Fuzzy Control and Filter Design
for Uncertain Fuzzy Systems
188 p. 2006 [978-3-540-37011-6]

Vol. 346: Tarbouriech, S.; Garcia, G.; Glattefeld, A.H. (Eds.)
Advanced Strategies in Control Systems
with Input and Output Constraints
480 p. 2006 [978-3-540-37009-3]

Vol. 345: Huang, D.-S.; Li, K.; Irwin, G.W. (Eds.)
Intelligent Computing in Signal Processing
and Pattern Recognition
1179 p. 2006 [978-3-540-37257-8]

Vol. 344: Huang, D.-S.; Li, K.; Irwin, G.W. (Eds.)
Intelligent Control and Automation
1121 p. 2006 [978-3-540-37255-4]

Vol. 341: Commault, C.; Marchand, N. (Eds.)
Positive Systems
448 p. 2006 [978-3-540-34771-2]

Vol. 340: Diehl, M.; Mombaur, K. (Eds.)
Fast Motions in Biomechanics and Robotics
500 p. 2006 [978-3-540-36118-3]

Vol. 339: Alamir, M.
Stabilization of Nonlinear Systems Using
Receding-horizon Control Schemes
325 p. 2006 [978-1-84628-470-0]

Vol. 338: Tokarzowski, J.
Finite Zeros in Discrete Time Control Systems
325 p. 2006 [978-3-540-33464-4]

Vol. 337: Blom, H.; Lygeros, J. (Eds.)
Stochastic Hybrid Systems
395 p. 2006 [978-3-540-33466-8]

Vol. 336: Pettersen, K.Y.; Gravdahl, J.T.;
Nijmeijer, H. (Eds.)
Group Coordination and Cooperative Control
310 p. 2006 [978-3-540-33468-2]

Vol. 335: Kozłowski, K. (Ed.)
Robot Motion and Control
424 p. 2006 [978-1-84628-404-5]

Vol. 334: Edwards, C.; Fossas Colet, E.;
Fridman, L. (Eds.)
Advances in Variable Structure and Sliding Mode
Control
504 p. 2006 [978-3-540-32800-1]

Vol. 333: Banavar, R.N.; Sankaranarayanan, V.
Switched Finite Time Control of a Class of
Underactuated Systems
99 p. 2006 [978-3-540-32799-8]

Vol. 332: Xu, S.; Lam, J.
Robust Control and Filtering of Singular Systems
234 p. 2006 [978-3-540-32797-4]

Vol. 331: Antsaklis, P.J.; Tabuada, P. (Eds.)
Networked Embedded Sensing and Control
367 p. 2006 [978-3-540-32794-3]

Vol. 330: Koumoutsakos, P.; Mezic, I. (Eds.)
Control of Fluid Flow
200 p. 2006 [978-3-540-25140-8]

Vol. 329: Francis, B.A.; Smith, M.C.; Willems,
J.C. (Eds.)
Control of Uncertain Systems: Modelling,
Approximation, and Design
429 p. 2006 [978-3-540-31754-8]

Vol. 328: Loria, A.; Lamnabhi-Lagarrigue, F.;
Panteley, E. (Eds.)
Advanced Topics in Control Systems Theory
305 p. 2006 [978-1-84628-313-0]

Vol. 327: Fournier, J.-D.; Grimm, J.; Leblond, J.;
Partington, J.R. (Eds.)
Harmonic Analysis and Rational Approximation
301 p. 2006 [978-3-540-30922-2]

Vol. 326: Wang, H.-S.; Yung, C.-F.; Chang, F.-R.
 H_∞ Control for Nonlinear Descriptor Systems
164 p. 2006 [978-1-84628-289-8]

Vol. 325: Amato, F.
Robust Control of Linear Systems Subject to
Uncertain
Time-Varying Parameters
180 p. 2006 [978-3-540-23950-5]

Vol. 324: Christofides, P.; El-Farra, N.
Control of Nonlinear and Hybrid Process Systems
446 p. 2005 [978-3-540-28456-7]

Vol. 323: Bandyopadhyay, B.; Janardhanan, S.
Discrete-time Sliding Mode Control
147 p. 2005 [978-3-540-28140-5]

Vol. 322: Meurer, T.; Graichen, K.; Gilles, E.D.
(Eds.)
Control and Observer Design for Nonlinear Finite
and Infinite Dimensional Systems
422 p. 2006 [978-3-540-27938-9]

Vol. 321: Dayawansa, W.P.; Lindquist, A.;
Zhou, Y. (Eds.)
New Directions and Applications in Control
Theory
400 p. 2005 [978-3-540-23953-6]

Vol. 320: Steffen, T.
Control Reconfiguration of Dynamical Systems
290 p. 2005 [978-3-540-25730-1]

Vol. 319: Hofbauer, M.W.
Hybrid Estimation of Complex Systems
148 p. 2005 [978-3-540-25727-1]

Vol. 318: Gershon, E.; Shaked, U.; Yaesh, I.
 H_∞ Control and Estimation of State-multiplicative
Linear Systems
256 p. 2005 [978-1-85233-997-5]

Vol. 317: Ma, C.; Wonham, M.
Nonblocking Supervisory Control of State Tree
Structures
208 p. 2005 [978-3-540-25069-2]

Vol. 316: Patel, R.V.; Shadpey, F.
Control of Redundant Robot Manipulators
224 p. 2005 [978-3-540-25071-5]

Vol. 315: Herbordt, W.
Sound Capture for Human/Machine Interfaces:
Practical Aspects of Microphone Array
Signal Processing
286 p. 2005 [978-3-540-23954-3]

¥1075.00元

Table of Contents

Locating RF Emitters with Large UAV Teams	1
<i>Paul Scerri, Robin Ginton, Sean Owens, and Katia Sycara</i>	
Out-of-Order Sigma-Point Kalman Filtering for Target Localization Using Cooperating Unmanned Aerial Vehicles	21
<i>Gregory L. Plett, Dimitri Zarzhitsky, and Daniel J. Pack</i>	
Multi-cumulant Control for Zero-Sum Differential Games: Performance-Measure Statistics and State-Feedback Paradigm	45
<i>Khanh D. Pham</i>	
Decentralized Cooperative Optimization for Multi-criteria Decision Making	65
<i>Cristinca Fulga</i>	
Simultaneous Localization and Planning for Cooperative Air Munitions	81
<i>Andrew J. Sinclair, Richard J. Prazenica, and David E. Jeffcoat</i>	
Second-Order Cone Programming (SOCP) Techniques for Coordinating Large-Scale Robot Teams in Polygonal Environments	95
<i>Jason C. Derenick and John R. Spletzer</i>	
UAV Splay State Configuration for Moving Targets in Wind	109
<i>Derek Kingston and Randal Beard</i>	
A Risk-Based Approach to Sensor Resource Management	129
<i>Dimitri Papageorgiou and Maxim Raykin</i>	
Constructing Optimal Cyclic Tours for Planar Exploration and Obstacle Avoidance : A Graph Theory Approach	145
<i>Abhishek Tiwari, Harish Chandra, Jacob Yadegar, and Junxian Wang</i>	
An Analysis and Solution of the Sensor Scheduling Problem	167
<i>Mesut Yavuz and David Jeffcoat</i>	
Cooperative Vision Based Estimation and Tracking Using Multiple UAVs	179
<i>Brett Bethke, Mario Valenti, and Jonathan How</i>	

Waypoint Selection in Constrained Domains (for Cooperative Systems)	191
<i>Yongjie Zhu, Yongling Zheng, and Ümit Özgüner</i>	
Cooperative Formation Flying in Autonomous Unmanned Air Systems with Application to Training	203
<i>Hongliang Yuan, Vivian Gottesman, Mark Falash, Zhihua Qu, Etyan Pollak, and Jiangmin Chunyu</i>	
Virtual Leader Based Formation Control of Multiple Unmanned Ground Vehicles (UGVs): Control Design, Simulation and Real-Time Experiment	221
<i>Wenchuan Cai, Liguang Weng, R. Zhang, M. Zhang, and Y.D. Song</i>	
Cooperative Control of Multiple Agents and Search Strategy	231
<i>Vitaliy A. Yatsenko, Michael J. Hirsch, and Panos M. Pardalos</i>	
Real-Time Optimal Time-Critical Target Assignment for UAVs	265
<i>Yoonsoo Kim, Da-Wei Gu, and Ian Postlethwaite</i>	
Sequential Inspection Using Loitering	281
<i>Meir Pachter, Swaroop Darbha, and P. Chandler</i>	
Distributed Cooperative Systems with Human Operator-in-the-Loop	293
<i>Pavlo A. Krokhmal and David E. Jeffcoat</i>	
Decentralized Extremum-Seeking Control of Nonholonomic Vehicles to Form a Communication Chain	311
<i>Cory Dixon and Eric W. Frew</i>	
An Adaptive Sequential Game Theoretic Approach to Coordinated Mission Planning for Aerial Platforms	323
<i>Dan Shen, Genshe Chen, Jose B. Cruz Jr., and Khanh D. Pham</i>	
Characteristics of the Distribution of Hamming Distance Values Between Multidimensional Assignment Problem Solutions	339
<i>Alla R. Kammerdiner, Pavlo A. Krokhmal, and Panos M. Pardalos</i>	
Robust Cooperative Visual Tracking: A Combined NonLinear Dimensionality Reduction/Robust Identification Approach	353
<i>Vlad I. Morariu, Octavia I. Camps, Mario Sznajder, and Hwasup Lim</i>	
A Lagrangian-Based Algorithm for a Combinatorial Motion Planning Problem	373
<i>Sai K. Yadlapalli, Waqar A. Malik, Swaroop Darbha, and Siva Rathinam</i>	

A Random Keys Based Genetic Algorithm for the Target Visitation Problem	389
<i>Ashwin Arulseivan, Clayton W. Commander, and Panos M. Pardalos</i>	
Cooperative Rendezvous Between Active Autonomous Vehicles	399
<i>Yechiel J. Crispin and Marie E. Ricour</i>	
Author Index	423

Locating RF Emitters with Large UAV Teams

Paul Scerri, Robin Glington, Sean Owens, and Katia Sycara

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213, USA

{pscerri, rglinton, owens, katia}@cs.cmu.edu

Gerald Fudge and Joshua Anderson
L-3 Communications Integrated Systems
Greenville, TX 75402, USA

{Gerald.L.Fudge, Joshua.D.Anderson}@L-3com.com

Abstract. This chapter describes a principled, yet computationally efficient way for a team of UAVs with Received Signal Strength Indicator (RSSI) sensors to locate radio frequency emitting ground vehicles in a large environment. Such a capability has a range of both civilian and military applications. RSSI sensor readings are noisy and multiple emitters will cause ambiguous, overlapping signals to be received by the sensor. Generating a probability distribution over emitter locations requires integrating multiple signals from different UAVs into a Bayesian filter, hence requiring cooperation between the UAVs. To build a coherent distributed picture given communication limitations, the UAVs share only those sensor readings that induce the largest changes in their local filter. Each UAV translates its probability distribution into a map of information *entropy* and then plans a path that will maximize the reduction in entropy (or conversely provides the highest information gain.) Planned paths are shared with a subset of other UAVs to minimize overlapping search. Experiments in a medium fidelity simulation environment show the approach to be lightweight and effective. Live flight results with lightweight Class I UAVs validate our approach.

1 Introduction

The rapidly improving availability of small, unmanned aerial vehicles (UAVs) and their ever decreasing 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 meters 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 chapter, we look at how teams of UAVs can use very small Received 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 technologies 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 [8], 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 [12] or planning cooperative paths [11], 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 [9].

Our approach to this problem has three key elements that enable locating RF emitters with large teams of lightweight UAVs. The first key element is a distributed filter to localize RF emitters in the environment. Each UAV has a Binary Bayesian Grid Filter [7] 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 teammates know in order to choose what to send [9], we started from the assumption that if some information leads to large local information gain, it will probably do so for much of the team. We investigated two information gain based heuristics for choosing which readings to share with teammates. 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 teammates 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 area, the entropy will remain high and the UAV will consider re-traversing 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 dynamism 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.

2 Problem

This chapter presents a method for localizing an unknown number of RF emitters using a team of UAVs. UAVs are outfitted with RSSI sensors which detect the power of an RF signal at a position in space. The UAVs must maintain a belief over the state of all emitters in the environment in a decentralized manner.

The emitters are represented by the set: $E = \{e_1 \dots e_n\}$ where n is not known to the team of UAVs. Emitters are all assumed to be emitting at a single known frequency.¹ Emitters are mobile and emit intermittently. The homogeneous UAVs are represented by the set: $U = \{u_1 \dots u_m\}$. Each u_i flies a path given by $\mathbf{u}^i(t)$. During flight a UAV takes sensor readings, $z_t(\mathbf{loc})$ which are the received signal power at a location $\mathbf{loc} = \{x, y, z\}$ where $\{x, y, z\}$ gives the Euclidean coordinates of a point in space relative to a fixed origin. The power of the signal received is a result of three components. The first component, $\Gamma(\mathbf{loc}, e_i) = \frac{e_{const}}{dist(\mathbf{loc}, e_i)^2}$, where $dist(\mathbf{loc}, e_i)$ is the Euclidean distance between \mathbf{loc} and e_i and e_{const} is a constant that gives the power at $dist(\mathbf{loc}_{e_1}, e_i) = 0$, is due to the sources themselves. The second component, $EN(\mathbf{loc}, E)$, is due to multi-path and attenuation of the signal due to environmental factors. Multi-path occurs when a reflected component of the signal arrives at a receiver and in combination with an attenuated direct signal results in a perturbed source

¹ This will be relaxed in future work.

location estimate. Finally ϵ gives typical zero-mean normally distributed sensor noise. The total power received at a location (loc) in space is then given by:

$$z_t(\mathbf{loc}) = \sum_{e_i \in E} \Gamma(\mathbf{loc}, e_i) + EN(\mathbf{loc}, E) + \epsilon \sim \mathcal{N}(0, \sigma)$$

Figure 1 shows some signals that will be received at different distances from a single emitter (i.e., no overlap). This is the basic signal model used in the simulation results below and closely represents real data collected from RSSI sensors on a physical UAV. The x-axis shows the distance and the y-axis shows the signal strength in dB (which is a log scale.) There are two important things to notice about this signal. First, it is very noisy, with high variation at all distances from the emitter, with some background noise high enough to represent being close to the emitter. Second, it has a very long “tail”, i.e., at a reasonable distance from the emitter there is still useful information in the signal. Figure 2 shows the sensor readings when the UAV flies near one emitter and then another. Notice the overlap in the signals between the emitters, which are about 350m apart.




Fig. 1. Sensor readings taken from different distances from an RF emitter

The sensor readings taken by the i th UAV, up until time t are $z_{t_0}^i \dots z_t^i$. Each UAV maintains a posterior distribution P over emitter locations given by $P_t^i(e_1 \dots e_n | z_{t_0}^i \dots z_t^i)$. The UAVs proactively share sensor readings to improve each other’s posterior distribution. At time t each u_i can send some subset of locally sensed readings: $z_t^i \subset z_{t_0}^i \dots z_t^i$.

The true configuration of the emitters in the environment at time t is represented as a distribution Q such that

$$Q_t(e_1 \dots e_n) = 1$$

when $e_1 \dots e_n$ gives the true configuration of the emitters at t . The objective is to minimize the divergence between the team belief and the true state of the emitters, while minimizing the cost of UAV flight path, and minimizing the total

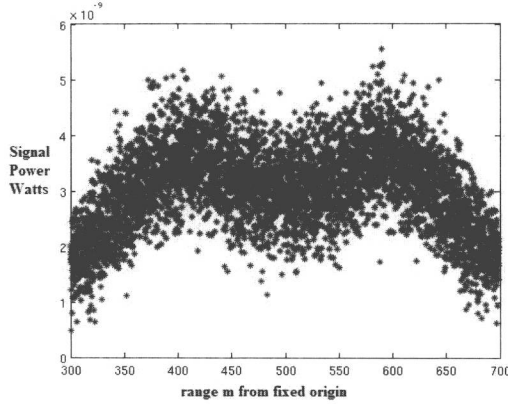


Fig. 2. Sensor readings taken when flying between two emitters, first near one, then near the other

number of messages shared between UAVs. The following function expresses this mathematically:

$$\min_{\mathbf{u}^i} \sum_t \sum_{u_i \in U} \beta_1 \text{Cost}(\mathbf{u}^i(t)) + \beta_2 D_{\text{KL}}(P_t^i \| Q) + \beta_3 |\mathbf{z}_t^i|$$

where D_{KL} denotes the Kullback Leibler divergence and β_1, \dots, β_3 are weights which control the importance of the individual factors in the optimization process.

3 Algorithm

The most important feature of the overall algorithm is the tight integration of all the key elements to maximize performance at a reasonable computational and communication cost. A Binary, Bayesian Grid Filter (BBGF) maintains an estimate of the current locations of any RF emitters in the environment. This distribution is translated into a map of the entropy in the environment. The entropy is captured in a *cost map*. UAVs plan paths with a modified Rapidly-expanding Randomized Tree (RRT) planner that maximize the expected change in entropy that will occur due to flying a particular path. The most important incoming sensor readings, as computed by the KL information gain they cause, are forwarded to other members of the team for integration into the BBGFs of other UAVs. Planned paths are also shared so that other UAVs can take into account the expected entropy gain of other UAVs when planning their own paths. The paths of other UAVs are also captured in a *cost map*. Additional cost maps, perhaps capturing results of terrain analysis or no-fly zones, can be easily added to the planner.

3.1 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 prac-

tical 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 physical UAV or on a UAV ground station, depending on the vehicle. In the experiments below, the simulations use *exactly* the same proxy code as the live flight experiments with physical UAVs. Figure 3 shows the main components and information flows from the perspective of one UAV-proxy.

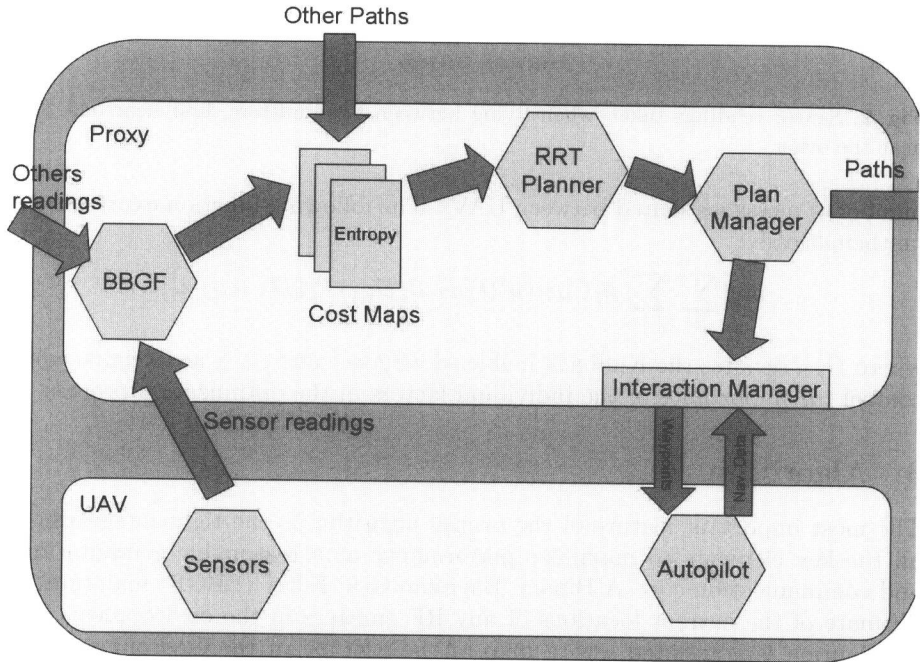


Fig. 3. Block diagram of architecture

4 Distributed State Estimation

In this section, we describe the filter used to estimate the locations of the emitters and the decisions individual UAVs make about sending information to one another.

4.1 Binary, Bayesian Grid Filter

The filter uses a grid representation, where each cell in the grid represents the probability that there is an emitter in the area on the ground corresponding to

that location.² For a grid cell c the probability that it contains an emitter is written $P(c)$. The grid as a whole acts as the posterior $P_t^i(e_1 \dots e_n | z_{t_0}^i \dots z_t^i)$.

To make calculations efficient, we represent probabilities in *log odds* form, i.e., $l_t = \log P(i)$. Updates on grid cells are done in a straightforward Bayesian manner.

$$l_t = l_{t-1} + \log \frac{P(e_i | z_t)}{1 - P(e_i | z_t)} - \log \frac{P(e_i)}{1 - P(e_i)}$$

where $P(e_i | z_t)$ is an inversion of the signal model, with the standard deviation extended for higher powered signals, i.e.,

$$P(e_i | z_t) = \begin{cases} \frac{1}{\sqrt{2\pi(\sigma_1^2)}} e^{-\frac{1}{2}(z_t - \Gamma)^2} & \text{if } z_t \geq \Gamma \\ \frac{1}{\sqrt{2\pi(\sigma_2^2)}} e^{-\frac{1}{2}(z_t - \Gamma)^2} & \text{otherwise} \end{cases}$$

where $\sigma_1 > \sigma_2$ scales the standard deviation on the noise to take into account structural environmental noise and overlapping signals. Intuitively, overlapping and other effects might make the signal stronger than expected, but they are less likely to make the signal weaker than expected. Figure 4 shows a plot of the (log) probability (y-axis) of a signal of a particular strength (x-axis) when the emitter is 500 m from the sensor.




Fig. 4. Mapping between probability and signal strength

Notice that there is no normalization process across the grid because the number of emitters is not known. If the number of emitters were known, a normalization process might be able to change the probability of emitters even in areas where no sensor readings had been taken. Initial values of grid cells are set to values reflecting any prior knowledge or some small uniform value if no knowledge is available.

² A quad-tree or other representation might reduce memory and computational requirements in very large environments, but the algorithmic complexity is not justified for reasonable sized domains.