

Collaborative Sensing by Unmanned Aerial Vehicles

W. T. Luke Teacy, Jing Nie,
Sally McClean and Gerard Parr
School of Computing & Information Engineering
University of Ulster
Cromore Road, Coleraine, BT52 1SA, UK
{l.teacy,j.nie,gp.parr,s.mcclean}@ulster.ac.uk

Stephen Hailes and Simon Julier
Department of Computer Science
University College of London
London WC1E 6BT, UK
{s.hailes, s.julier}@cs.ucl.ac.uk

Niki Trigoni and Stephen Cameron
Computing Laboratory
University of Oxford
Oxford OX1 3QD, UK
{niki.trigoni,stephen.cameron}@comlab.ox.ac.uk

ABSTRACT

In many military and civilian applications, Unmanned Aerial Vehicles (UAVs) provide an indispensable platform for gathering information about the situation on the ground. In particular, they have the potential to revolutionize the way in which information is collected, fused and disseminated. These advantages are greatly enhanced if swarms of multiple UAVs are used, since this enables the collection of data from multiple vantage points using multiple sensors. However, enhancements to overall operational performance can be realised only if the platforms have a high degree of autonomy, which is achieved through machine intelligence.

With this in mind, we report on our recently launched project, SUAAVE (Sensing, Unmanned, Autonomous, Aerial VEhicles), which seeks to develop and evaluate a fully automated sensing platform consisting of multiple UAVs. To achieve this goal, we will take a multiply disciplinary approach, focusing on the complex dependencies that exist between tasks such as data fusion, ad-hoc wireless networking, and multi-agent co-ordination. In this position paper, we highlight the related work in this area and outline our agenda for future work.

Categories and Subject Descriptors

I.2.9 [Artificial Intelligence]: Robotics—*Autonomous vehicles*

General Terms

Theory, Algorithms, Experimentation

Keywords

UAV, UAV swarms, sensor networks, decentralized control

1. INTRODUCTION

In many military and civilian applications, an aerial view is invaluable to gain information about the situation on the ground [13]. Such applications include search and rescue, perimeter surveillance, crowd control and situation awareness in natural disasters. In manned flight, such scenarios place a heavy burden on pilots, requiring long hours of

monotonous flight at high-levels of concentration. Moreover, in combat situations, there may be a significant risk to the pilot if a mission must be carried out over hostile territory.

Increasingly, however, advances in airframe design and control technology mean that using Unmanned Aerial Vehicles (UAVs) for such tasks is becoming a viable option. Small, inexpensive aircraft are now commercially available, and using GPS technology, exhibit a high degree of stability in the air (see Section 6). Existing applications of these systems typically involve small numbers of UAVs working in isolation, where each one is under the constant control of a single user on the ground. Thus, operating current UAVs can still require a significant number of man hours, and there is limited technological support to co-ordinate the actions of multiple UAVs effectively.

For this reason, there are many potential benefits to be gained by increased autonomy and co-ordination in the control of multiple UAVs. In particular, we identify the following three as significant.

1. Fully autonomous UAVs require less human intervention, and therefore can increase the number of UAVs that can be operated by a single user.
2. By working together, autonomous UAVs can split up to reduce redundancy when performing a search task, and so can cover a large area efficiently with minimal resources.
3. Multiple UAVs can increase accuracy in sensing tasks by fusing information gathered from different viewpoints.

Realizing this potential is the aim behind our recently launched project, SUAAVE (Sensing, Unmanned, Autonomous, Aerial, VEhicles). The SUAAVE consortium is an interdisciplinary group in the fields of computer science and engineering. Its focus is to investigate and elucidate the principles underlying the control of clouds of networked resource-limited UAVs acting as sensor platforms, that are targeted towards achieving a global objective in an efficient manner. We consider

a number of application scenarios, for example, disaster relief after an earthquake, flood prediction and monitoring, homeland security and surveillance, and search and rescue.

Although extensive research relevant to this vision has been done before (See [13] for a full review), published works suffer from at least one of three limitations:

1. they focus on a small part of the sensing problem in isolation, such as image processing [6] or communication [9];
2. they are developed and evaluated in simulated environments [10], which make simplifying assumptions that may not hold in real-life applications; or
3. they are not directly applicable to UAV control, because they are aimed either at stationary or ground based mobile sensor networks [8].

In SUAAVE, our intention is to address these limitations by adopting a practical engineering approach, building fully integrated hardware and software systems, which we shall evaluate in a real environment. In this way, we not only intend to make novel contributions to individual areas of research, but shall investigate the complex interplay that exists between tasks such as data fusion, wireless communication, and distributed search strategies. Moreover, by evaluating our solutions as fully deployed systems, we aim to establish results on performance characteristics that are directly applicable to real world scenarios.

In the following sections, we shall outline in more detail the key issues we wish to investigate with an emphasis on the interplay between the problems and solutions they entail. More specifically, in the remainder of this paper, we highlight our proposed research in the areas of wireless ad-hoc networking (Section 2), distributed data fusion (Section 3), decentralized optimization (Section 4) and safety management (Section 5). In Section 6, we discuss the hardware platform that we shall use to conduct experiments. Finally, we conclude in Section 7, by summarising the main research issues that we aim to address in SUAAVE, and how we propose to do so.

2. COMMUNICATION

In most communication networks, routing and transmission protocols are loosely coupled with the applications that they support. This approach makes sense in networks, such as the Internet, that must support a wide variety of applications, and so should not be optimized to suit one purpose at the expense of another. However, in mobile sensor networks, there is generally a case for tighter integration between routing and application processes. There are two reasons for this:

1. Sensor networks are generally designed with a specific set of tasks in mind. We therefore have more information that can be used to optimize communication algorithms to suit the application.
2. Radio Frequency (RF) characteristics can be greatly affected by the current position of a UAV relative to

other UAVs and features on the ground. This can lead to trade-offs between moving UAVs into a good location for sensing, and moving into a good location for communication.

Together, these two features mean that there is both the need and opportunity to develop tightly coupled sensing and communication protocols, which can optimize the overall performance of the sensor network by taking into account the dependencies between communication and sensing tasks.

In SUAAVE, we plan to achieve this by following three lines of investigation. First, using the hardware platform outlined in Section 6, we shall collect data about how factors such as atmospheric conditions, and UAV position relative to the ground affect RF characteristics. Although some relevant data about RF characteristics does exist [9, 4], the number of variables involved mean that results do not always generalize well, and so there is no substitute for performing platform specific observations. Second, using this data, we shall develop probability models of RF characteristics as a function of space and time. These will be used to inform optimization algorithms when deciding how a UAV should behave to maximise its overall performance. Third, we shall develop both unicast and multicast routing protocols, which will take into account the needs and behaviour of high-level processes to optimize communication performance.

3. DISTRIBUTED DATA FUSION

Although UAVs provide several advantages for gathering information about various phenomena, the nature of this platform poses unique challenges for data analysis and fusion. In particular, although it is in the nature of all sensors to generate noisy output, by mounting them on a moving object whose position is difficult to determine with accuracy introduces an extra level of uncertainty. For example, even small changes to the pitch of the craft can make significant changes to the angle that it makes with the horizon, thus making it difficult to associate images with specific regions on the ground.

Another important issue is the location where data fusion takes place. The conventional approach is to use centralised data fusion — information is sent to a central site to be fused. However, such architectures require significant bandwidth and are potentially fragile: a failure in the link with the central site means that no fusion can take place. An alternative approach is to use distributed data fusion (DDF) [1], in which fusion occurs in nodes throughout the network. Such networks can be inherently robust (failures lead to gradual degradation), scalable (nodes need only know local network topology) and modular (new nodes can be introduced to enhance sensing capabilities). However, distributed fusion introduces a number of challenges. One of the most significant problems is double counting [5]: the system overestimates the amount of information available, leading to implausibly accurate estimates.

In the SUAAVE project, we shall investigate DDF algorithms that overcome these issues, and can also optimize the location where data is fused and processed based on factors such as the available processing capacity of each node,

and their proximity to the data. We will then implement field versions of these algorithms that take into account positional information and RF characteristics, as discussed in the previous section.

4. MULTI-OBJECTIVE DECENTRALIZED OPTIMIZATION

For most non-trivial applications, a UAV's actions must be guided by multiple objectives and are subject to multiple constraints. For example, in a search and rescue scenario, the optimal action for a UAV is not necessarily to cover the search area using the shortest possible path. Instead, it may also need to avoid obstacles, return to base to refuel, or move to a different position to acquire a reliable communication link with a ground control unit or other UAVs.

The complexity of this issue is further increased when UAVs act as part of a collective to achieve a common set of goals. In this case, actions of multiple UAVs must be co-ordinated to achieve the best results. For instance, by flying in a specific formation, UAVs may cover a search area more efficiently, achieve more reliable communication links or minimise the chance of collision. More significantly, by taking into account all of these aspects in parallel the collective may be able to achieve the best trade-off between all three.

In such cases, it may be possible to achieve trade-offs in a centralised way by sending all information to be processed at a single location (for example, a control station on the ground). However, as with data fusion, adopting a centralised approach can incur an unacceptable communication overhead and leave the system vulnerable if the single point of control fails or becomes unavailable. For this reason, we plan to investigate the use of decentralized control and optimization algorithms for co-ordinating multiple UAVs.

To achieve this, we plan to exploit and build on existing work on decentralized co-ordination in multi-agent systems. In particular, agent co-ordination problems can often be framed as Distributed Constraint Optimization Problems (DCOPs) for which a number of optimal solutions exist [11] [12]. However, existing optimal algorithms for DCOPs scale poorly as the number of agents in a system increases and are prone to failure if messages between agents are lost or the topology of the network changes. However, more robust results have been shown using the max-sum algorithm [7]. Although this is not guaranteed to produce optimal solutions in all cases, it has demonstrated near optimal performance empirically in a wide variety of cases, scales well in large systems containing many agents, and is robust against message loss and failure of individual agents.

Although these approaches are promising, one of their drawbacks is that they do not explicitly deal with cases in which agents must co-ordinate their future actions in an uncertain environment. For example, during search and surveillance, UAVs may need to plan their future joint actions based on current information, and then re-plan should an unexpected obstacle turn up in their flight path. Problems such as these are addressed by multi-agent reinforcement learning algorithms. However, existing approaches can suffer from limitations, such as lack of convergence guarantees. In our work, we shall attempt to address these issues by looking at how



Figure 1: Ascending Technologies Hummingbird UAV

DCOP algorithms can be used to share information between UAVs to achieve more effective results.

5. SAFETY AND COLLISION AVOIDANCE

As discussed in previous sections, we aim to develop sophisticated autonomous UAVs that can take into account multiple objectives when deliberating over their actions. However, increased autonomy and sophistication comes with increased risk. This is particularly true if UAVs are to operate in or near to populated areas where there is the potential for collision resulting in personal injury or damage to property, or legal issues regarding flight in prohibited areas. Thus, to ensure public acceptance of such technology, it is essential that we have appropriate safety mechanisms in place to minimise both the probability of a malfunction occurring and the damage caused if one does occur.

For this reason, we shall develop a safety protocol and associated architecture that will fulfill three roles. First, it will provide real time information and predictive models to inform high-level decision processes about the current and future state of critical resources. Most significantly, this will include information about a UAV's current energy supply and how this will change in response to future actions. Second, it will verify that each planned action does not move the UAV into a position of adverse risk, such as into the flight plan of other UAVs, or close to known obstacles or no flight zones. Third, it will continually monitor the current state of the UAV to detect malfunctions if and when they occur. If a fault is detected the UAV will attempt a set of pre-programmed behaviours to recover from the failure if possible (for example, by invoking collision avoidance algorithms [2, 3]), or abort the mission in the safest possible way.

6. IMPLEMENTATION AND EXPERIMENTS

Although software simulation can be a useful tool for exploring the properties of proposed algorithms, only so much can be learnt without applying a system to its intended purpose. This is particularly true for UAV applications, since simplifying assumptions are often made during algorithm or simulation development, which may give an inaccurate portrayal of how a system will perform in a real environment.

For this reason, we are adopting a practical approach on the SUAAVE project, developing fully integrated systems deployed on real UAVs. Initially, this will be based around As-

ending Technologies¹ Hummingbird platform, which consists of a four rotor airframe with on-board battery, GPS receiver, and wireless communication capabilities (see Figure 1).

At an approximate diameter of 53cm, this device is relatively small and inexpensive and can carry a payload of up to 200 grams. Out-of-the-box, it is capable of maintaining its position using its GPS receiver; and accepting instructions by remote control, or from on board software developed using the vendor's APIs. In each case, commands can either be issued in terms of low level maneuvers (for example, by varying thrust to alter the yaw, pitch or roll of the aircraft) or as instructions to visit GPS way-points.

From a research perspective, these features give us two advantages. First, the available APIs and potential payload allow us to add our own devices, such as additional processors and sensors, to interface with the main system to support additional functionality. Second, since most low level control issues are taken care of by the platform, we can concentrate on adding high-level functionality, such as co-ordination and data fusion algorithms.

In our initial trails, we plan to use four of these UAVs with on board cameras, to gather realistic data for image processing tasks and radio transmission characterization. This will be used to carry out initial experiments and analysis in the lab, which will inform the development of our communication protocols and high-level control algorithms. Finally, we aim to evaluate a complete system consisting of up to ten UAVs, including high-level co-ordination, safety and sensing mechanisms.

7. CONCLUSIONS

In many applications, an aerial view is indispensable for improving situation awareness about various phenomena on the ground. Until recently, this could only be achieved by manned flight, which is expensive and places a high burden on the pilot, with long hours of potentially dangerous activity at high levels of concentration. UAVs have the potential to alleviate these problems, but their maximum benefit can only be achieved through a high level of automation and co-ordination between multiple UAVs.

In SUAAVE, we aim to achieve this vision by developing fully integrated hardware and software systems. We shall investigate the complex dependencies that exist between tasks such as ad-hoc wireless communication, data fusion and multi-agent co-ordination; and so produce mechanisms that optimize overall system performance.

Moreover, to establish performance results that are directly applicable to the real world, we shall evaluate our solutions using field trails of multiple UAVs operating together to achieve a common set of goals.

Acknowledgments

The SUAAVE project is a collaboration involving academic partners from University College London, the University of Oxford, and the University of Ulster. This work is funded by

¹<http://www.asctec.de>

the UK Engineering and Physical Sciences Research Council (EPSRC), grant reference EP/F06358X/1.

8. REFERENCES

- [1] D. Akselrod, A. Sinha, and T. Kirubarajan. Hierarchical markov decision processes based distributed data fusion and collaborative sensor management for multitarget multisensor tracking applications. In *the IEEE International Conference on Systems, Man and Cybernetics*, pages 157 – 164, 2007.
- [2] E. Boivin, A. Desbiens, and E. Gagnon. Uav collision avoidance using cooperative predictive control. In *the 16th Mediterranean Conference on Control and Automation*, pages 682 – 688, 2008.
- [3] F. Borrelli, T. Keviczky, and G. Balas. Collision-free uav formation flight using decentralized optimization and invariant sets. In *the 43rd IEEE Conference on Decision and Control*, volume 1, pages 1099 – 1104, 2004.
- [4] C. Cheng, P. Hsiao, H. Kung, and D. Vlah. Performance measurement of 802.11 a wireless links from uav to ground nodes. In *the 15th International Conference on Computer Communications and Networks*, 2006.
- [5] C.-Y. Chong, S. Mori, W. Barker, and K.-C. Chang. Architectures and algorithms for track association and fusion. *IEEE Aerospace and Electronic Systems Magazine*, 15(1):5 – 13, Jan 2000.
- [6] G. Conte and P. Doherty. An integrated uav navigation system based on aerial image matching. In *the IEEE Aerospace Conference*, pages 1 – 10, 2008.
- [7] A. Farinelli, A. Rogers, A. Petcu, and N. Jennings. Decentralised coordination of low-power embedded devices using the max-sum algorithm. In *the 7th International Conference on Autonomous Agents and Multi-Agent Systems*, pages 639–646, 2008.
- [8] E. Ferranti, N. Trigoni, and M. Levene. Brick&mortar: an on-line multi-agent exploration algorithm. In *the IEEE International Conference on Robotics and Automation*, 2007.
- [9] D. Hague, H. Kung, and B. Suter. Field experimentation of cots-based uav networking. In *the Military Communications Conference*, Jan 2006.
- [10] R. Henriques, F. Bacao, and V. Lobo. Uav path planning based on event density detection. In *International Conference on Advanced Geographic Information Systems & Web Services*, pages 112 – 116, 2009.
- [11] P. J. Modi, W. Shen, M. Tambe, and M. Yokoo. Adopt: Asynchronous distributed constraint optimization with quality guarantees. *Artificial Intelligence Journal*, 161:149–180, 2005.
- [12] A. Petcu and B. Faltings. A scalable method for multiagent constraint optimization. In *the 19th International Joint Conference on Artificial Intelligence*, pages 266–271, 2005.
- [13] A. Ryan, M. Zennaro, A. Howell, R. Sengupta, and J. Hedrick. An overview of emerging results in cooperative uav control. In *the 43rd IEEE Conference on Decision and Control*, volume 1, pages 602 – 607, 2004.