Swarm Control in Unmanned Aerial Vehicles

Henry Hexmoor, Brian McLaughlan, and Matt Baker
University of Arkansas
Computer Science and Computer Engineering Department
Engineering Hall 313, Fayetteville AR 72701
{hexmoor, bpm02, mmb05} @uark.edu
Phone: 479 575 2420 Fax: 479 575 5339

Abstract

Swarm control is an open problem. We are taking steps for systematic flight control of large numbers of UAVs. In this paper, we present our design and implementation of UAV group control in a fully implemented simulated UAV environment. We explore issues of limited user intervention, novel agent personality, and multi-theater load balancing.

1. Introduction

Controlling UAV swarms via human supervision is of great interest to the US military. This paper reports on a project that models control of a fleet of unmanned aerial vehicles (UAVs) with minimal user intervention via simulation. The project simulates a situation where multiple UAVs must locate and track multiple ground vehicles. Once located, a ground vehicle must be "scanned" for a certain length of time to simulate gathering information about the target. If the target is lost and later reacquired, scanning may be picked up where left off. Any UAV may contribute to the scanning process. However, multiple UAVs scanning a single target at the same time will not speed up the scanning requirement. Once completely scanned, the target is of no further significance and is removed from consideration.

We have implemented a simulation of UAV swarms in Java. Figure 1 shows a snapshot of our simulated environment and its GUI control.

The system of UAV agents is essentially a hive mind organization controlled by the main loop of the program. However, individual decisions are controlled by the individual UAVs. The main loop is a turn-based system containing three main sub-phases. Each UAV is instructed in turn to perform a sensor pulse to locate nearby targets. Phase two consists of a negotiation where UAVs determine their best targets and settle on which UAV should be allowed to track which target. In the third and final phase, the action phase, each UAV is allowed in turn to perform its chosen objective. UAVs are capable of high speed and maneuverability relative to the target vehicles. The exact maximum speed is parametric. Since power consumption is not considered in this simulation, the UAV generally fly at their maximum speed. However, when nearing its chosen target, a UAV will slow down to avoid overshooting the target position. The UAV's movement is constrained to two degrees of freedom (the horizontal plane), but is quite maneuverable within those parameters. It is able to move in eight directions – North, South, East, West, Northeast, Southeast, Southwest, and Northwest – and can move in multiple directions each turn up to its movement limit. The method used to move the UAVs is rather simple. At the end of each turn, each UAV is given the opportunity to move toward its intended destination. The difference between the intended and actual positions is computed. The UAV then moves to minimize that difference. The UAV is only able to sense ground targets. The range is specified by a parameter and generates a circle of data around the UAV. Any target not completely within the sensor circle is not seen. The human operator is allowed to give two types of orders directly to individual UAVs. It can force a UAV to track a specific target. The UAV will then travel to the last known location of that target and attempt to acquire it. Second, the operator can ban a UAV from tracking a target. Regardless of its perceived appropriateness, the target cannot be tracked by the UAV until the ban is lifted. These commands are rarely issued. We found that the UAVs are generally quite capable of locating targets without any specific operator commands. The only times where these commands were required were when multiple UAVs were deadlocked in a bidding war or when a UAV became "too dedicated" to reacquiring a lost target. This latter situation is easily remedied via an appropriate UAV personality parameter adjustment.

Figure 1. Snapshot of our UAV simulation screen

We have explored several concepts including user suggestion and intervention, role negotiation, agent personality, and load balancing.

User suggestion is explored through the use of "hot" and "cold" regions that guide the swarm to areas the user deems productive or unproductive search areas. The individual UAVs are then free to incorporate these suggestions into their search planning algorithms.

Role negotiation is provided through a bidding process. In our simulation each UAV bids on its preferred target, and if it is outbid, it is allowed to bid on another target. To simulate disagreement, we adjust agent role negotiation algorithm. Instead of allowing the negotiation process to continue until all UAVs are satisfied, each agent is given only a single chance to bid on a proposed target at each turn. Being outbid causes the UAV to become frustrated and forces them to eventually take drastic action if it is not pacified. In addition, the fact that the negotiation may last multiple cycles means that the human operator must intervene when UAVs cease tracking due to impasse.

To parameterize UAV personality, we defined four characteristics. *Conformity* is the amount of consideration that the agent gives to user suggestions. *Sociability* is a statistic parameter that determines how UAVs clump together or spread apart. *Dedication* determines how quickly a UAV gives up after losing sight of its target. *Disposition* determines how quickly a UAV becomes frustrated with the negotiation process.

The simulation environment is abstracted in several ways. First, it is based on a rectangular grid. The grid parameters are adjustable at compile time. Second, the simulation is turn based. All vehicles must stay in the bounds of the grid. The only exception to this is when a UAV is transported to another theater of operation. In this case, it is removed from the current theater and sent via network to the destination. The environment contains simulated clouds that obscure the UAV sight. Any vehicle that moves under a cloud becomes hidden. The cloud movement is also under human control. The number and density of clouds can be controlled by the user.

The human operator is allowed two opposite general commands that designate a grid location as a hot or cold regions. When determining whether to move toward a random location to begin searching for targets, a UAV will consider the effects of all hot and cold spots on that location. The effects of hot and cold spots degrade exponentially. Thus, they maintain potency in close proximity to their location, but have little if any effect at any significant distance.

A hot spot near a potential search location will signify to the UAV that the operator believes travel to that location will yield significant tracking results – either in acquiring lost targets or discovering new ones. The UAV will weigh that location higher in proportion to the proximity to that location. Conversely, a cold spot will signify that the operator deems that searching near this location to be a futile effort. The UAV will weigh areas around this Spot lower than usual.

If a region lies on or near the path of a UAV, but the UAV destination is beyond the range of the location, the UAV will not divert its course toward or away from the spot. It also will not affect a UAV's desire to track a target that is known to reside within a spot. The region only affects the UAVs' desire to choose that location as its search destination. The UAV planning is flexible enough to allow the UAVs to change their minds halfway through an action if it determines that there exists a better one. For example, a UAV may decide that it will travel to a less populated area, but during its travels, another UAV locates a target that it does not want to track. If the first UAV determines that this new target is the most appropriate target, it will abandon its destination and choose to track the newly discovered target.

In the remainder of the paper, we begin by outlining related work. Next, we describe the overall simulation environment followed by details of the UAV model. We then describe our personality parameters and offer discussions of their use in our experimental testbed.

The ground vehicles simulate targets that the UAVs must track. They model rather dumb terrestrial vehicles that are not aware of each other or the UAVs. A ground vehicle randomly chooses a destination somewhere on the grid. Then during its movement phase, it moves one square closer to that location. When it reaches the destination, the process restarts. When a vehicle is tracked for ten time units, it is eradicated.

2. Related Work

Computational models of coordinated movements by collections of locally interacting individuals (swarms) include bird flocks [Reynolds, 1987, 1999], fish schools [Huth, et.al, 1992; Tu, et.al., 1994], social insect swarms [Bonabeau and Dorigo, 1999], and self-assembling molecules [Edwards, et.al, 1998] where numerous autonomous particles (e.g., birds, ants, vehicles, etc.) operate by reacting to local forces exerted on them by other nearby particles or their environment. Work on particle swarms [Parker, 1993], cultural algorithms [Chung and Reynolds, 1996], and bacterial chemotaxis algorithms [Muller, et.al., 2002] has generalized the idea for an abstract, n-dimensional cognitive spaces that make up self-organizing particle systems. Interactions between particles result in complex global behavior which emerges from the joint actions and relatively simple behaviors of the individual particles, thereby exhibiting self organization. These properties have been used in applications in computer graphics [Reynolds, 1987, 1999], multi-robot team control [Balch and Arkin 1998; Fredslund and Matartic 2002; Winder and Reggia, 2004; Vail and Veloso, 2003], and numerical optimization [Parker, 1993].

Research on UAV swarm control is a topic of interest in many research laboratories. The ability to control many remote entities with minimal user intervention has many military and commercial applications. Current techniques for controlling UAVs, which rely on centralized control and on the availability of global information, are not suited for the control of UAV swarms, owing to the complexity that arises from the interactions between swarm elements. Traditional, centralized approaches frequently lead to exponential increases in communication bandwidth requirements and in the size of the controlling swarm. In contrast, swarms of simple biological or artificial organisms can exhibit rich emergent behaviors without the need for centralized control or global communication [Bonabeau, Dorigo and Theraulaz, 1999].

In our testbed, the human operator is allowed to adjust four parameters that make up the UAV personality. Adjusting a parameter only affects UAVs that are currently present in the operator's battlefield. The personality parameters we will describe in the following sections are conformity, sociability, dedication, and disposition.

3.0 Conformity

Conformity determines how quickly the UAV reacts to operator suggestions. This does not affect specific commands such as forced targeting or banning. It only refers to user suggestions of "hot" and "cold" regions.

For instance, a UAV with a low conformity will not give much weight to operator suggestions. On the other hand, a UAV with high conformity will almost always follow the operator suggestions whenever possible.

This parameter is executed by scaling the hot and cold spots by a factor of the Conformity value. For instance, if a "hot" spot has a magnitude of 1 (the default value), a UAV with conformity of 50 will treat that spot as if it had a heat of 50.

Because conformity essentially causes the hot and cold regions to expand and to contract, the range of useful conformity values depend upon the size of the grid. For our typical 50 by 50 grid, scores between 10 and 50 were the most effective values.

A score of zero effectively neutralizes the hot and cold cues and is counterproductive. Similarly, a negative score is possible but not useful, as it essentially makes the UAVs openly belligerent to the operator suggestions.

4. Sociability

Sociability determines how gregarious the UAVs will be. A UAV with a positive sociability will tend to operate in proximity to other UAVs. Conversely, a negative sociability (anti-social behavior) will make the UAV shun others and operate independently. The magnitude of the sociability value determines the amplitude of behavior, i.e., the degree to which a behavior is exhibited.

This parameter causes a UAV to consider other UAVs as hot (attractor) or cold (repulsor). This depends on the sign of the score. Thus, when considering a possible destination, the UAV will weigh it higher or lower depending on the presence of other UAVs.

Sociability and Conformity have a relative relationship. If one parameter is significantly larger than the other, that parameter will essentially neutralize the other. If the two parameters are similar in magnitude, they will be given equal weight. However, the magnitude of the scores will still determine how likely a UAV will be to consider the score. For instance, consider a UAV that has both its Sociability and Conformity scores set at a very high levels relative to the size of the grid. Because the scores are equal, it will give equal weight to possible destinations that are near hot spots and other UAVs. However, because the scores are so high, it is almost assured that the UAV will choose a destination that is near at least one UAV or a hot spot.

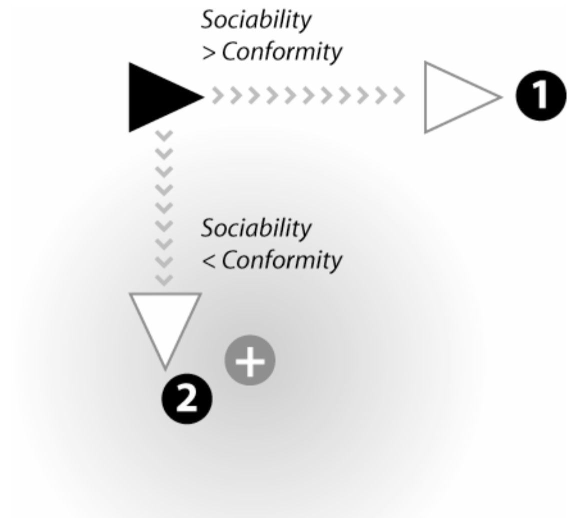

Figure 2: All things being equal, the two targets are just as likely. A UAV with a higher Sociability will prefer Location 1. Conversely, if the UAV has a higher Conformity, it will prefer Location 2.

Sociability scores can be set to any positive integer. However, high positive scores relative to the size of the grid tend to be counterproductive, as they cause the UAVs to remain in clumps and retard travel around the screen. In such, individual UAVs lose their ability to explore the battlefield and are doomed to circle their adversaries.

5. Commitment

Dedication (or commitment) determines how committed the UAVs are to reacquire lost targets. A UAV with a low dedication score will quickly give up if it loses sight of its target, while a high dedication score will cause the UAV to be more persistent. If a UAV decides not to search for a lost target, the UAV will announce to the group that the target's location has become unknown and will then seek a new target

A UAV that chooses to pursue its target will not announce the target's disappearance, but will instead circle around the last known location until a) it reacquires the target, b) another UAV locates the target and declares the location, or c) it finally gives up and chooses a new plan of action.

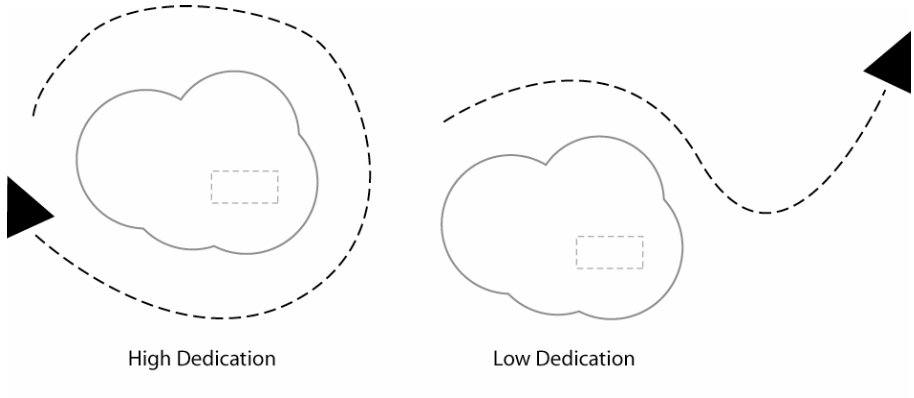

Figure 3: UAVs will circle a cloudbank and search for its target longer if it is Dedicated.

Dedication is expressed in terms of a percentage, which can be any integer from 0 to 100, and represents the probability of maintaining the search in each turn. A UAV with a dedication value of 0 never circles the target. A value of 100 will force the UAV to always circle.

Typically, the ideal dedication level is inversely proportional to the amount of cloud cover. A battlefield with substantial cloud-cover would not be suitable for high dedication levels since the targets often won't quickly reappear from hiding. UAVs may be forced to circle in one side of a huge cloudbank while the target remains hidden across the screen. Conversely, high dedication scores are beneficial when clouds are sparse since targets will typically only be hidden for a very short time. In this case, circling a cloudbank is more useful than abandoning the target. However, extremely high dedication scores (near 100%) are not useful, as multiple UAVs become locked in circling patterns with little recourse until their targets are located or they are forced to choose new targets through user intervention.

6. Disposition

We use disposition (as a measure of temperament) to set the rate of frustration among UAVS for target acquisition. This is formally also known as neuroticism. Often, a UAV's plans of a are subverted by another UAV during the negotiation phase. This is most likely due to the fact that the second UAV is more qualified to track the target that the first UAV. The first UAV must wander until the next negotiation phase.

Repeated subversion will cause a UAV to become frustrated. A UAV with a low disposition value will quickly become upset and refuse to relinquish its chosen target. In this situation, the other UAVs will leave it alone until the neurotic UAV voluntarily abandons the target (typically due to destroying the target track, losing it in the clouds, or finding a more appropriate target).

Disposition is any non-negative integer value. It represents the number of time units that the UAV will remain frustrated without giving up its target.

The simulation is equipped with the ability to send UAVs across a network to other theaters of operating theatre. This gives the operator the ability to balance the utilization of UAVs as resources across semi-isolated locations. Also, it allows the interaction between multiple types of UAVs. It is possible that UAVs with different personalities and performance capabilities could provide better utilization than a single type of UAV.

To simulate data messaging and synchronization, our model utilizes a group memory. When a UAV sees a target or decides on a course of action, it updates the appropriate list. Conversely, when a UAV wants to learn information about the known targets or the other UAVs, it consults the available lists for data.

7. Target Bidding

During the objective determination phase, each UAV is given the opportunity to bid on its most desired target. The target desirability is primarily determined by proximity and seconly by the amount of tracking already accomplished against that target.

Each UAV determine its best target and submits it to the group. If two or more UAVs choose the same target, the closest) UAV wins the bid. Each UAV is only given only one chance to bid. Thus, the UAVs that lose the bidding process are forced to perform a wandering search until the next negotiation phase.

While bidding, it may become apparent that a UAV is not going to be able to pick a target. This may be due to any combination of lack of targets, unknown target locations, or losing a bid. In any of these situations, a default action is selected.

A UAV first selects a number of random destinations. The exact number of destinations depends on the size of the coordinate grid and is computed at 2% of the total grid locations. The UAV determines at which of these locations it is most likely to find unknown targets and sets that location as its intended destination.

The UAV then determines likelihood of finding an unknown target by weighing the location's proximity to hot and cold regions as well as other UAVs. The weights that it gives to these objects are determined by its conformity and sociability personality parameters. The location that is determined most likely to contain uncontested targets is chosen as the UAV's next destination.

It is debatable as to whether cooperation will yield better results due to the increased communication overhead involved.

7. Discussion and Conclusions

Human operators will ultimately control UAV swarms at a high level. Our work takes a modest step toward a multi-tiered architecture that provides a high level control using social and personality parameters. The transition from strictly remote-controlled vehicles to intelligent, goal-directed, cooperative agents requires a well-defined hierarchy of autonomies.

The Commitment parameter (we also called it dedication) could be further automated since it has a fairly straightforward use. There may be circumstances where the operator would want to override the internal adjustments.

Second, Temperament (we also called it disposition) appears to be a difficult parameter. Intuitively, one would assume that a low disposition would be appropriate in a setting where there is an overabundance of UAVs. This "bad temper" would allow the UAV to stick up for itself in a situation where there would most likely be many other UAVs that would try to take its target. However, we were unable to determine any significant benefit from adjusting the disposition values.

Finally, it would appear that Sociability and Conformity allow a wide range of adjustments. We have determined that positive sociability and negative conformity can be detrimental in many situations. In addition, fine-tuning the settings does not appear to have a significant effect. We found the best settings to be either relatively small or large positive values (i.e., no middle value) for Conformity and negative values for Sociability. Regarding reducing argumentativeness, we were unable to produce enough arguments to warrant user interdiction. With dedication, conformity, and sociability characteristics, UAVs perform a subset of behaviors that would normally require exhaustive human supervision.

References

T. Balch & R. Arkin. Behavior-based Formation Control for Multi-robot Teams, IEEE Robotics and Automation, 1998, 14(6),926-939.

E. Bonabeau, M. Dorigo & G. Theraulz. Swarm Intelligence, Oxford Univ. Press, 1999.

C. Chung and R. Reynolds, A Testbed for Solving Optimization Problems Using Cultural Algorithms, Evolutionary Programming, 1996, 225-236.

L. Edwards, Y. Peng and J. Reggia, Computational Models for the Formation of Protocell Structures, Artificial Life 4 (1998), 61-77.

J. Fredslund and M. Mataric, A General Algorithm for Robot Formations Using Local Sensing and Minimal Communication, IEEE Transactions on Robotics and Automation, 2002, 18(5):837-846.

A. Huth and C. Wiesel, Simulation of Movement of Fish Schools. J. Theoret. Biol. 1992, 156:365-385.

S. Muller, J. Marchetto, S. Airaghi and P. Kourmoutsakos, Optimization Based on Bacterial Chemotaxis. IEEE Transactions on Evolutionary Computation, 2002, 6:16-29.

L. Parker, Designing Control Laws for Cooperative Agent Teams, Proc. IEEE International Conference on Robotics and Automation, 1993, 3:582-587.

X. Tu, and D. Terzopoulos, Artificial Fishes: Physics, Locomotion, Perception, Behavior, Proc. ACM SIGGRAPH, 1994, 42-48.

D. Vail and M. Veloso, Multi-Robot Dynamic Role Assignment and Coordination Through Shared Potential Fields, In Multi-Robot Systems. A. Schultz, L. Parker, and F. Schneider (editors), Kluwer, 2003.

S. Muller, J. Marchetto, S. Airaghi and P. Kourmoutsakos, Optimization Based on Bacterial Chemotaxis. IEEE Transactions on Evolutionary Computation, 2002, 6:16-29.

D. Vail and M. Veloso, Multi-Robot Dynamic Role Assignment and Coordination Through Shared Potential Fields, In Multi-Robot Systems. A. Schultz, L. Parker, and F. Schneider (editors), Kluwer, 2003.

R. Winder and J. Reggia, Distributed Partial Memories to Improve Self-Organizing Collective Movements, IEEE Trans. SMC B, 2004, 34:1967-1707.

C. Reynolds, Flocks, Herds and Schools, Computer Graphics, 1987, 21:25-34.

C. Reynolds, Steering Behaviors for Autonomous Characters, Proc. Game Developers Conf. 1999, 763-782.