Evolving Large Scale UAV Communication System

Adrian Agogino
UCSC at NASA Ames
Mail Stop 269-3
Moffett Field, CA 94035
Adrian.K.Agogino@nasa.gov

Chris HolmesParker
Oregen State University
204 Rogers Hall
Corvallis, OR 97331
holmespc@onid.orst.edu

Kagan Tumer
Oregen State University
204 Rogers Hall
Corvallis, OR 97331
Kagan.Tumer@oregonstate.edu

ABSTRACT

Unmanned Aerial Vehicles (UAVs) have traditionally been used for short duration missions involving surveillance or military operations. Advances in batteries, photovoltaics and electric motors though, will soon allow large numbers of small, cheap, solar powered unmanned aerial vehicles (UAVs) to fly long term missions at high altitudes. This will revolutionize the way UAVs are used, allowing them to form vast communication networks. However, to make effective use of thousands (and perhaps millions) of UAVs owned by numerous disparate institutions, intelligent and robust coordination algorithms are needed, as this domain introduces unique congestion and signal-to-noise issues. In this paper, we present a solution based on evolutionary algorithms to a specific ad-hoc communication problem, where UAVs communicate to ground-based customers over a single wide-spectrum communication channel. To maximize their bandwidth, UAVs need to optimally control their output power levels and orientation. Experimental results show that UAVs using evolutionary algorithms in combination with appropriately shaped evaluation functions can form a robust communication network and perform 180% better than a fixed baseline algorithm as well as 90% better than a basic evolutionary algorithm.

Categories and Subject Descriptors

I.2.11 [**Computing Methodologies**]: Artificial Intelligence—*Multiagent systems*

General Terms

Algorithms, Performance, Reliability

Keywords

UAVs, Evolution, Multiagent Systems

GECCO'12, July 7-11, 2012, Philadelphia, Pennsylvania, USA.

1. INTRODUCTION

Traditionally, unmanned aerial vehicles (UAVs) have been used for targeted surveillance and military operations. Whether they are powered for many hours by jet engine, or for a fraction of an hour by electric motor, UAV missions tend to have the same profile: They are launched, they carry out a mission, then they return [5]. Such missions are naturally time-limited, labor intensive and logistically challenging, especially since the repeated takeoff and landings need to be coordinated with air traffic control. Currently the number of simultaneous UAV missions is in the single digits. However, technology may soon change this. For more than a decade, development has progressed on solar powered UAVs with battery reserves, allowing them to fly for a certain amount of time after dark. Rapid industry advancements in battery capacity, electric motor technology and solar cell efficiency, is allowing the progression of solar powered UAVs to proceed even faster [17]. Once UAV technology goes over the critical hump, where they can fly all night on charges received during the day, the mission profiles of UAVs will change radically: they will go up, and stay up. UAV missions could last for months if not years, with just a small turnover for maintenance.

In fact, at the margins, we are already there as the solar powered QinetiQ's Zephyr has been demoed, flying for more than two weeks [15]. Permanently flying UAVs allow aircraft to be used in domains currently monopolized by satellite and ground based-systems, including two-way communications, continuous surveillance and broadcasts. In addition the technologies in solar UAVs are relatively cheap, with continual price declines. Soon these aircraft will be available to nearly all nations, institutions and perhaps even individuals. Also since they do not need oxygen to burn fuel, they can be designed to fly at altitudes much higher than conventional aircraft can fly, allowing them be used independently of the current air traffic system. Due to these factors, there will likely be an explosion in the number and uses of UAVs in the future [3].

We believe that genetic and evolutionary algorithms in combination with multiagent techniques will have a key role in the complex problem of controlling such a large diverse set of aircraft. Multiagent systems match the distributed nature of the problem, while evolution addresses the messy nonlinear control and coordination issues. While the potential applications of these UAVs are vast, we believe that evolution and multiagent systems are applicable to a wide range of UAV applications that share these common properties:

- Long distance communication is easy due to UAVs being line-of-sight to each other and to ground.
- Communication congestion is severe due to UAVs being line-of-sight.
- Coordination of UAVs should be distributed due to their large numbers, their geographic separation, and UAVs being owned by different institutions.
- Hardware failures and integration failures will be common, due to differing age, types, manufactures and sheer volume of UAVs.

These properties make efficient control difficult with a top-down approach, but are a natural for adaptive and distributed approaches such as evolutionary algorithms and multiagent systems.

In this paper we apply these multiagent techniques combined with evolution to the specific domain of creating an air-to-ground communication network over a single channel. In some ways this model is related to current WiFi networks that can share many users over a single channel. However, the dynamics of having the signals come from high altitude UAVs make the problem considerably different than that of terrestrial communication: Huge numbers of UAVs may be accessible at once, congestion may be severe, and many failed or uncooperative UAVs will be line-of-sight and will have to be dealt with. In these senses, this domain shares many of the qualities that we expect future UAV problems to have. In addition this domain is important in itself as it allows UAV based communication networks to be used for such purposes as voice communication and data networks without the need of expensive and difficult to maintain ground-based hardware. In addition it allows such networks to be created in an adhoc way, where different aircraft are owned by different institutions. If setup properly, this paradigm may open the path for UAV based communications to be as ubiquitous for long-range communication, as WiFi is for short-range communication today.

The main contribution of this paper is to present an application of evolutionary algorithms to the domain of UAV communication that is both useful for downlink communication and shows the potential for swarms of high-altitude, low-cost UAVs likely to be common in the future. In Section 2 we describe current and future UAV systems as well as multiagent communication. In Sections 3 and 4 we describe details of the UAV communication domain used in this paper. In Section 5 we discuss how the evolving agents can be used to maximize customer communication bit rates. In Section 6 we present experimental results, showing how the evolving agents are able to perform well under varying conditions. Finally in Section 7 we provide a discussion on the impact of this application and the implications of the results.

2. BACKGROUND

In the near future, solar UAVs will play critical roles in the military, industrial, scientific, and academic communities [14, 17, 18]. These devices have seemingly limitless applications including communications, reconnaissance missions, space launch platforms, and wireless power beaming [10, 14]. Recent missions including NASA's Pathfinder-Plus and QinetiQ's Zephyr (which remained airborne for over two weeks nonstop) have advanced the state of the art in solar powered UAVs, taking them from limited mission life and endurance to the point they can remain operational for weeks at a time [10, 15]. As a result of the increasing capabilities and availability of these devices coupled with their falling costs, a plethora of novel domains and applications will emerge to utilize the newly developed technological capabilities of these platforms [14].

Methods of controlling and coordinating networks of UAVs have been researched including genetic and evolutionary algorithms and reinforcement learning methods[9, 2, 16, 13, 11, 7, 6]. Genetic algorithms have been used to evolve decision trees that allow UAV teams to collaborate in search missions, and to facilitate UAV task assignments [16, 19]. In addition evolutionary algorithms have been shown to be effective in single player and multi-player UAV path planning domains [8]. Genetic programming has also been shown effective in UAV multi-task allocation domains when there is limited communication [4].

In addition to evolutionary and genetic algorithms, swarm techniques have also been used to coordinate UAVs. In the cooperative hunters domain a swarm of UAVs, using hand coded optimization methods, is used to search for one or more "smart" targets [2]. Another UAV control problem focuses on an NP-complete task allocation problem which assigned tasks to swarms of UAVs [12]. Frequently UAVs are utilized in reconiassance tasks involving Automatic Target Recognition (ATR). In such domains it is desired to have a balance between high coverage of discovered targets and broad area coverage. One high-performing approach to solving this coordination and control problem utilizes ant-based swarm methods [12].

3. UAV TO EARTH COMMUNICATION

As we progress into the information age, communication becomes an increasingly critical component of every day life. Today, cellular phones, laptops, hand held computers, and other wireless electronic devices have changed the way we see and interact with the world. At the core of these advancements is a well designed wireless communication network, which handles the workload and facilitates information sharing between devices connected to the network. Current networks rely on a series of radio towers to facilitate this information sharing work load. Traditional towers have worked well to date, but they have several key drawbacks:

1. They are expensive to build.
2. They are expensive to maintain.
3. They have limited communication due to obstructions (cannot communicate "around" obstructions).
4. They have static placement (holes in coverage areas).

In this paper we focus on a subset of this domain where there is a set of UAVs that are flying at fixed locations (flying in small circles) for long periods of time (perhaps months or years) and are transmitting data to a set of customers below (see Figure 1). UAVs have an advantage in sending data from high altitude in that they can have line-of-sight communication to many customers. In addition by virtue of being overhead, such UAVs can focus on what areas of the surface they will project most of their signal power to, allowing for better coverage.

In this domain, each UAV can communicate to multiple customers. In addition communication is done over a shared channel (over the same frequency band) analogous to the way WiFi networks transmit data. Using a shared channel allows the system to be very adhoc, where UAVs can come and go, and can decide whether or not to participate in the system without any need for channel arbitration. Note that for simplicity we only look at the download problem, where UAVs are sending information down to customers. Also we make no assumptions on how the UAVs get their data feeds. We believe that this half of the problem is the most important, as typical internet use tends to be dominated by download traffic. Although the uplink problem is fairly similar as long as it is done on a different channel than the downlink.

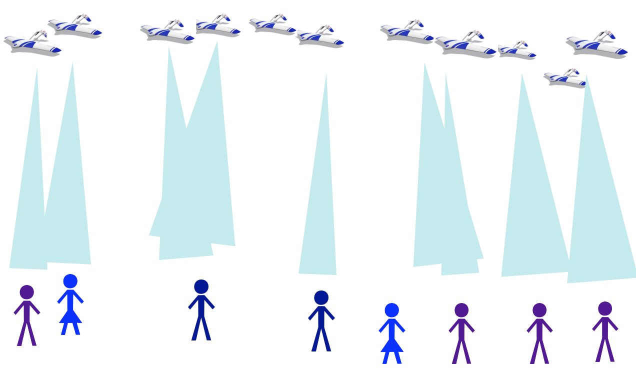

Figure 1: UAV Communications. A set of UAVs at high altitude transmit data to a set of customers on ground over a single communication channel. The task of the system is to maximize average bitrate customers receive. Multiple UAVs may communicate to single customer. A UAV communicates to at most one customer.

4. SIGNAL DYNAMICS

We assume that the UAVs are all at similar altitudes and communicate through directional antennas pointed towards the ground. The amount of area on the ground that is covered by the UAV is determined by the gain of its antenna. Antennas with low gain, transmit over a wider area, but within that area the strength of the signal is lower (see Figure 2). Antennas with high gain, have more signal power in the center of their area, but transmit over a smaller area. The maximum signal received from a UAV is proportional to the inverse square of the gain radius for the antenna:

$$S_j^{max} = aP_j/r_j^2 \tag{1}$$

where a is a constant, P_j is the power transmitted from UAV j, and r_j is the signal half-power radius for UAV j. S_j^{max} is the amount of signal received directly at the center of the transmission. Further from the center, the amount of signal received decreases exponentially according to the signal radius:

$$S_{i,j} = e^{-b\frac{r_{i,j}}{r_j}} S_j^{max} \tag{2}$$

where b is a constant and $r_{i,j}$ is the distance from customer i and the center of UAV j's transmission. The noise received by customer i is simply the sum of the signal from all the UAVs it is not communicating with:

$$N_i = \sum_{j \notin J_i} S_{i,j} + k \ , \tag{3}$$

where J_i is the set of UAVs customer i is communicating with and k is a constant for background noise. The maximum communication rate for customer i can then be estimated from the signal-to-noise ratio using Shannon's law:

$$C_{i,j} = B \log_2(1.0 + S_{i,j}/N_i) \ , \tag{4}$$

where B is the bandwidth of the channel in Hz [1]. The total data rate for customer i is the sum of the data rates for each UAV the customer is communicating with:

$$C_i = \sum_{j \in J_i} C_{i,j} \ . \tag{5}$$

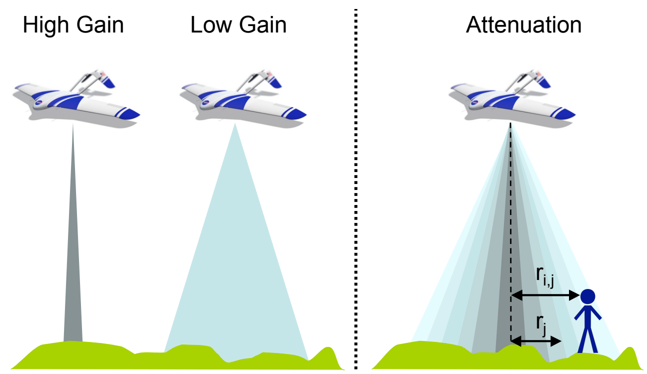

Figure 2: Signal Dynamics. UAVs with high-gain antennas throw a strong signal over a small area. UAVs with low-gain antennas throw weaker signal over larger area (Left). The strength of the signal depends on how far the customer is away from the center of the signal cone (Right).

5. SYSTEM EVALUATION FUNCTION AND AGENTS

The objective of this problem is to maximize the average data rate of each customer:

$$G = \frac{1}{n} \sum_{i=1}^{n} C_i \ , \tag{6}$$

where there are n customers, and G is the system evaluation function. Combining equation 1, 2, 3, 4, 5, we obtain:

$$G = \frac{1}{n} \sum_{i=1}^{n} \sum_{j \in J_i} B \log_2 \left(1.0 + \frac{\frac{aP_j}{r_j^2} e^{-b\frac{r_{i,j}}{r_j}}}{\sum_{j \notin J_i} \frac{aP_j}{r_j^2} e^{-b\frac{r_{i,j}}{r_j}} + k} \right) \ , \tag{7}$$

putting our global objective in terms of our control variables: UAV power level, P_j, and indirectly, $r_{i,j}$, through the orientation of the UAV.

[1]For simplicity, certain factors, such as relative inverse square distance signal attenuation are ignored that were determined to have little impact on performance.

These controls allow us to change the signal-to-noise characteristics at different locations on the ground. However, this is a difficult problem as increasing the signal for one customer may increase the noise for another. It is especially difficult, since we want this communication network to be adhoc, where it is controlled in a distributed way: UAVs are entering and leaving the system, and some UAVs fail to cooperate or operate correctly. Fortunately, evolutionary algorithms and multiagent techniques are a natural match to this problem.

There are many possible agent definitions and controls for the UAVs in our domain, including altitude, antenna gain, power levels and antenna angle. Here we focus on the last two: adjusting the power level P_j and orientation (the direction the transmitter points to) of each UAV (see Figure 3). We control each of these actions through agents. The solution to the full problem consists of the power level and orientation values for all the UAVs. However to simplify the problem, we break the task into a multiagent system, where a single agent controls both power level and orientation for each UAV. To perform control, each agent makes discrete actions. For adjusting power, the action is scaled exponentially to the action:

$$P_j = Pe^{z_j^p} \ , \tag{8}$$

where P is the base power and z_j^p is the action of the agent for UAV j controlling its power. To control orientation, an agent chooses one of nine directions: either straight down, or one of eight cardinal directions around the UAV. The angle of the pointing is fixed so that the center of the new orientation is moved a distance of r what it would have been if it had pointed strait down (see Figure 3 - right). As described in the next section, the values of the controls for each agent are determined by an evolutionary algorithm.

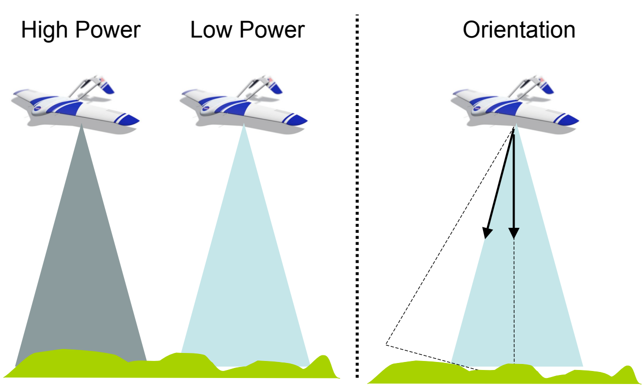

Figure 3: Agent Actions. An agent can choose power level of UAV within certain range. An agent can also choose orientation of antenna. The Agent must choose power levels and orientations to balance giving more signal to their customers and less noise to other customers.

5.1 Evolving Agents

The objective of each agent is to evolve the best values of power level and orientations that will lead to the best system fitness evaluation function, G. The value of power level and orientation is determined by the agent's current policy. Each policy contains a discrete value of z_j^p determining power level and a discrete value of z_j^o determining orientation direction. The power level is discretized to 10 different values. Along with the 9 possible orientations each policy for each agent has a total of 90 different possible values.

An agent's policy is determined by an evolutionary algorithm, where an agent evolves a population of policies. At every time step, an agent chooses a policy from its population of policies using an epsilon-greedy selector, where the best policy is chosen with probability $1 - \epsilon$ and a random policy is chosen with probability ϵ. The chosen policy table determines the power and orientation of the UAV for that time step. Once all the power levels and orientations are chosen the performance of the system is evaluated (see Figure 4). The types of evaluations that can be used is discussed in Section 5.2.

Once the current choice of policy has been evaluated, the evaluation of the policy is updated with a learning rate alpha: New Value = $(1 - \alpha)$ * Old Evaluation + α * New Evaluation. Each agent then updates its population by eliminating the lowest value member of the population, and then copies the highest value member of population and mutates it. Mutation is applied by taking a random table entry and setting each value to a number between 0 and 9 taken from a uniform random distribution. Various other forms of mutation were tried including mutating more table entries at each evolution step, but they did not improve performance.

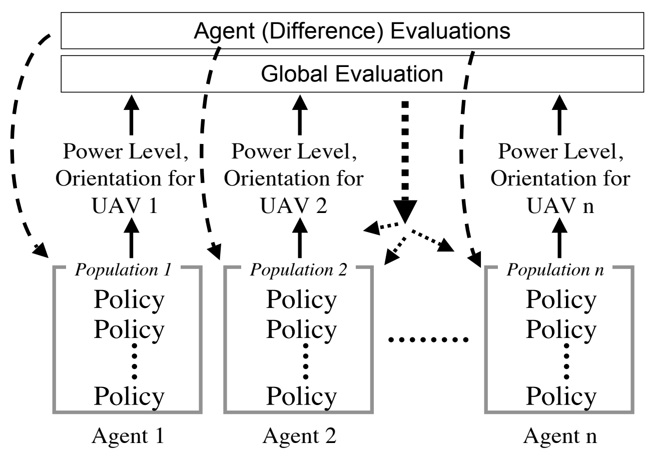

Figure 4: Evolving Populations of Agents. Each agent has its own population of policies. At every step a policy is chosen from the population, which determines the power level and orientation of a single UAV. The choice of power level and orientation can then be evaluated in two different ways: 1) Global evaluation looks at the value of all of the actions of all of the agents and returns the same value for all agents, 2) Shaped agent-specific evaluations (implemented with the "difference evaluation" in this paper) make a separate evaluation for each agent using information about all of the actions from all agents.

5.2 Agent Evaluation Functions

The final issue that needs to be addressed is selecting the fitness evaluation function for the evolving agents. The first and most direct approach is to let each agent receive the system performance as its evaluation function. However, in many domains using such an evaluation function leads to slow evolution. We will therefore also set up a second evaluation function based on an agent-specific evaluation. Given that agents aim to maximize their own evaluation functions, a critical task is to create "good" agent evaluation functions, or evaluation functions that when maximized by the agents lead to good overall system performance. In this work we focus on "difference" evaluation functions which aim to provide an evaluation function that is both sensitive to that agent's actions and aligned with the overall system evaluation function [1, 20].

5.2.1 *Difference Evaluation Function*

Consider **difference** fitness evaluation function of the form [1, 20]:

$$D_i = G(z) - G(z - z_i) \ , \qquad (9)$$

where z_i is the action of agent i (determined by its policy) as defined in Subsection 5.1, and $z - z_i$ are the actions of all the agents with the action of agent i removed. The second term of the difference evaluation, $G(z - z_i)$, represents a counterfactual of what the performance of the system is like when agent i is removed from the system (i.e. its power level is dropped to zero). By subtracting the counterfactual from the original evaluation, this difference in some sense evaluates the agents contribution to the system.

There are two advantages to using D: First, the second term removes a significant portion of the impact of other agents in the system. This happens since the impact of actions that are irrelevant to agent i are removed by the subtraction. This benefit has been dubbed "learnability" (agents have an easier time evolving) in previous literature [1, 20]. Second, because the second term does not depend on the actions of agent i, any action taken by agent i that improves D, also improves G. This happens since any action that the agent takes can only affect the first term, because its action has been eliminated from the second term. This benefit which measures the amount of alignment between two evaluation functions has been dubbed "factoredness" in previous literature [1, 20].

Substituting Equation 7 into Equation 9, we obtain

$$D_{j'} = \frac{1}{n}\sum_{i=1}^{n}\sum_{j \in J_i} B\log_2\left(1.0 + \frac{\frac{aP_j}{r_j^2}e^{-b\frac{r_{i,j}}{r_j}}}{\sum_{j \notin J_i}\frac{aP_j}{r_j^2}e^{-b\frac{r_{i,j}}{r_j}} + k}\right) -$$

$$\frac{1}{n}\sum_{i=1}^{n}\sum_{j \in J_i, j \neq j'} B\log_2\left(1.0 + \frac{\frac{aP_j}{r_j^2}e^{-b\frac{r_{i,j}}{r_j}}}{\sum_{j \notin J_i, j \neq j'}\frac{aP_j}{r_j^2}e^{-b\frac{r_{i,j}}{r_j}} + k}\right),$$

where we are calculating the difference evaluation for agent j'. Note that the second term of the difference evaluation both removes the signal received from this agent's UAV and also removes its noise.

6. EXPERIMENTS

To test the effectiveness of evolving agents in this UAV communication domain, we perform an extensive set of experiments in simulation. In these simulations, (unless otherwise specified, such as in the scaling experiments) 100 UAVs are placed at an altitude of 20 miles (representing approximately the maximum altitude a solar powered aircraft can achieve). These UAVs are placed above a 10x10 mile square area. The task of the UAVs is to transmit data to customers within this 10x10 mile area. In all experiments the channel bandwidth is B = 1Mhz and the noise floor is $k = 0.2$. In addition, the gain radii, r_j, are distributed randomly, uniformly between 0.35 miles and 1.05 miles to represent a heterogeneous set of UAVs. For evolution, $\alpha = 0.2$ and $\epsilon = 0.25$. All experiments results are performed over 30 trials. In additional all of our major performance conclusions are statistically significant with $p < 0.05$. In this setup we test the performance of the evolving agents in cases where:

1. Agents control UAVs with no failures and full communication.
2. UAVs (i) fail; (ii) do not coordinate; or (iii) are incompatible.
3. UAVs have restrictions on their observation capabilities.
4. The number of UAVs is scaled to 1000 UAVs.

In all of our scenarios, the number of customers is the same as the number of UAVs. We do this to model the situation where a ground-based "customer" is likely to be a hotspot or signal repeater. In situations where customers were individuals, the number of customers would likely be considerably more than the number of UAVs.

For each of the cases above, we report results on four different types of agents to control the power and orientation of the UAVs:

- Static agents always choose maximum power, and strait down orientation (M).
- Random agents have random evaluation function (R).
- Evolving agents directly maximize system evaluation function (G).
- Evolving agents use difference evaluation function (D).

The first two form baselines to asses performance. The next two compare learning rates of traditional agents maximizing a common system evaluation function, and agents indirectly maximizing the system evaluation function by directly maximizing the difference evaluation function.

6.1 UAV Performance

In the first set of experiments, we have a single agent control both power level and orientation for a UAV. The results displayed in Figure 5 show how the first baseline algorithm is able to achieve a data rate of 200Kbits/s to each customer. The results also show that evolving agents maximizing the system evaluation function directly (G) perform significantly better, performing up to 300Kbits/s. This result shows that evolution can be very helpful in choosing control policies in this large system. However, agents evolved using the D evaluation function are able to perform even better, nearly doubling the performance of the system. The improved performance of the difference evaluation has to do with it being more specific to the actions of the agent. When an agent chooses a good policy that policy is likely to be evaluated well using the difference evaluation. In contrast when using

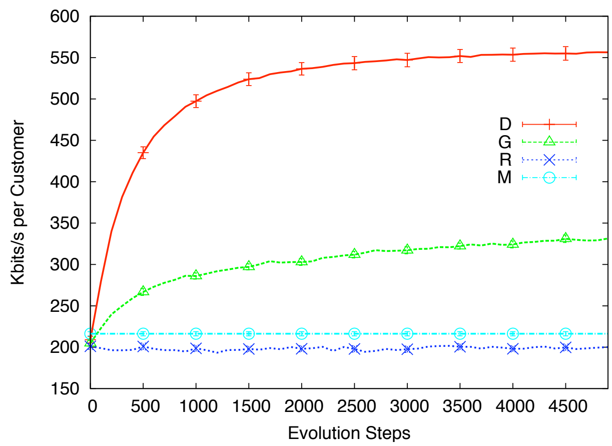

Figure 5: System of 100 UAVs where each agent optimizes both power and orientation. Agents evolved using D evaluation function outperforms all other methods by nearly 2-to-1. This is due to the D evaluation eliminating irrelevant actions.

G to evaluate that policy it may not get a good evaluation, since the evaluation depends equally on the policy choices of the 99 other agents.

6.2 Robustness to Failures, Non-coordination, and Incompatibility

For an adhoc UAV communications network to function properly, it must be robust to many types of failures. UAVs may be of different ages, be in different states of repair and may fail without notice. Even worse than failing completely, a UAV may still be transmitting at high power, but not communicating with any customer, causing them to add noise without any benefit. In this section we show how robust an evolved agent based UAV network can be against these various forms of failures.

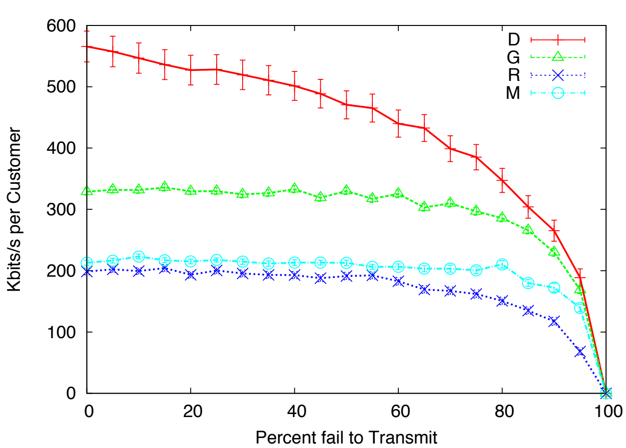

Figure 6: Performance under transmission failures. D agents outperform all other methods with up to 95% failures.

6.2.1 *Failure to Transmit*

First we consider the case where UAVs must turn off all transmissions due to failure or in order to conserve power. As agents fail, other agents must find ways to adapt to make up for the loss. As seen in Figure 6, agents using D are able to perform well, even under high failure rates. While agents using the baseline policy as well as agents using G are not hurt significantly from the failures, they still perform worse than agents evolved using D up to a 95% failure rate.

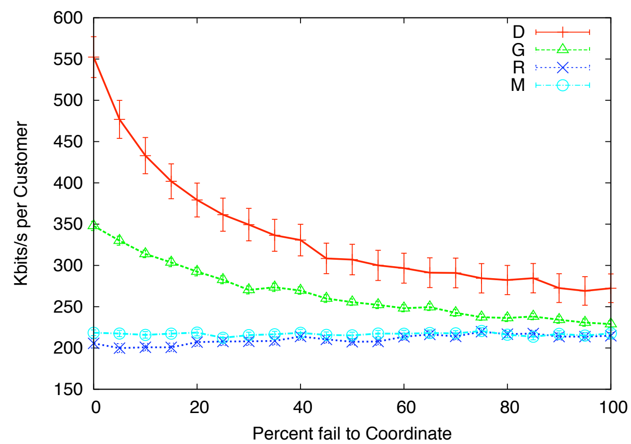

Figure 7: Performance when agents fail to coordinate. D agents outperform all other methods.

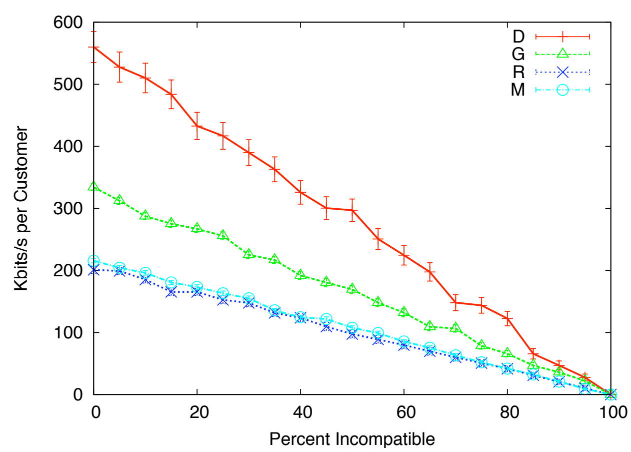

Figure 8: Performance under "incompatible" agents. Agents evolved using D evaluation function outperform all other methods with up to 90% failures.

6.2.2 *Failure to Coordinate*

Next we consider the case where some UAVs fail to coordinate with the rest of the system and end up taking uncoordinated locally greedy actions. To maximize local performance, these agents simply maximize their power level, without regard for the noise they are adding to the system. These agents still contribute to the overall system bandwidth, but they can potentially harm the system by raising the noise levels of other agents. As shown in Figure 7, agents evolved using D are able to outperform all other methods.

In fact even when 100% of agents greedily maximize their power levels, agents evolved using the difference evaluation still perform better since they are still able to efficiently choose good orientations.

6.2.3 *Failure of Compatibility*

In our final case, we consider the situation where some of the UAVs are incompatible with the current network. In this case, incompatible UAVs still send out noise, but do not actually send data to any of the customers in the system. Such cases may be common, when protocols change, software is not written to specification or there are multiple different networks communicating in the same channel. Figure 8 shows that this type of failure can be very harmful. Still, with a moderate number of incompatible UAVs, agents evolved using the difference evaluation are able to perform well.

6.3 Observation Restrictions

In order to demonstrate the concept and the suitability of the multiagent evolutionary approach to this domain, the results reported above assume that there are no restrictions in the observational capabilities of the agents (for example in computing or receiving the system objective and difference evaluation functions, or in communicating signal-to-noise ratios for all of the customers). Because such assumptions are not realistic, in this section we explore the impact of such limitation on the performance of each of the algorithms (note that "observation restrictions" here also includes inter-UAV communication, but we use the term "observation" to differentiate from the main application of UAV to ground communication).

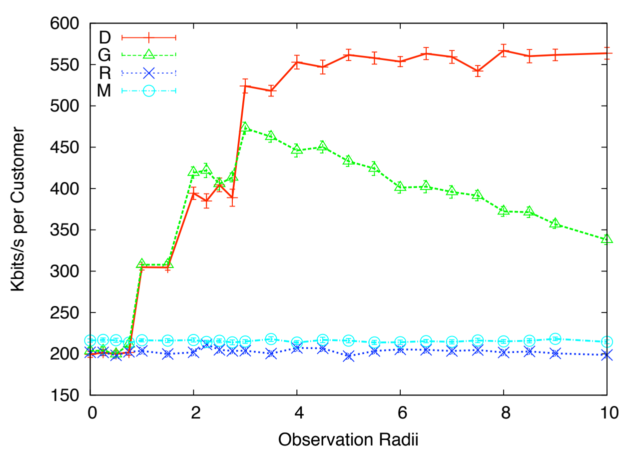

Figure 9: Performance of a 100 UAV system in the presence of observation restrictions. Agents can only observe actions of other agents that lie within a given radii from their location. Agents using D are able to effectively use additional observational information to coordinate and improve system performance, whereas agents using G can be negatively impacted by an increased amount of information.

As seen in Figure 9, agents evolved using D and G perform equally when no observations are possible (agents have access to only their own local information). As the level of observation increases and agents receive more information, an increase in performance could be expected. That is not what happens however; agents evolved using G actually experience a decrease in performance as they gain additional information. This is because the agents don't know what to do with the extra information they receive. Agents evolved using D are able to handle this information in a way that benefits the system performance, allowing them to coordinate and make decisions that positively improve system performance. Agents evolved using G begin to improve performance when the observation radius is below 2.25. This happens since the agents are receiving enough information to make more informed decisions, but not enough information to negatively impact their performance (too much information causes noise on the agents' evaluation function). After this point however, the amount of information each agent receives becomes overwhelming and they are unable to coordinate their actions.

6.4 Scalability

While all of the previous experiments have been performed with 100 UAVs, we expect future UAV systems to be much larger. In this section we test the scalability of our approach by measuring the system performance when the number of UAVs is scaled from 10 to 1000. To make scaling comparable, we also scale the number of customers and the area of the land serviced by the same amount. The results show that for more than 100 UAVs, the amount of data each customer receives is highly stable (see Figure 10). This result suggests that agents evolved using difference evaluation functions should be able to efficiently control this system, even when there are a very large number of UAVs.

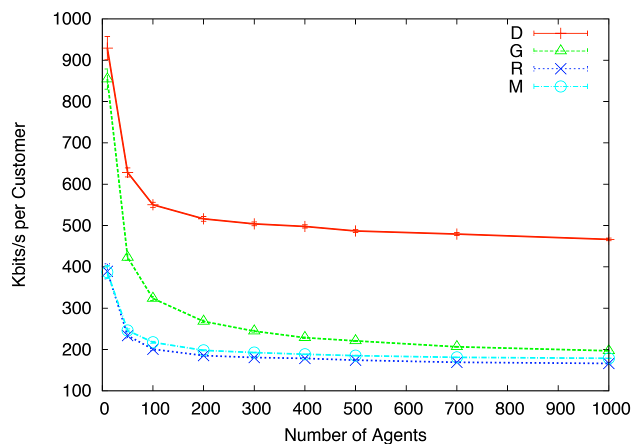

Figure 10: Performance versus scalability where the number of customers, agents, and world size are scaled proportionally. As seen, agents evolved using D are able to outperform all other methods by nearly 50% to 100% for any given system size. Agents using D with a system size of 1000 UAVs perform nearly as well as agents evolved using G when the system size is 50 UAVs. This is a twenty fold increase in system size while maintaining good performance.

7. DISCUSSION AND CONCLUSION

In this paper, we present an important application of multiagent evolution: Coordinating large numbers of UAVs to form an air-to-ground communication network. Due to advances in solar, battery and motor technology, such large UAV networks are becoming an attractive alternative to both Earth-based and satellite based communication systems. Indeed, such networks will be increasingly powerful, allowing for voice and data networks anywhere in the world without need for expensive and brittle ground-based infrastructure, or the need for expensive-to-launch and maintain satellite systems. In addition, with minor modifications, the UAV coordination algorithms can be adapted for other types of large scale UAV applications, ranging from observation systems to microwave power delivery systems.

The application chosen in this paper exhibits a number of salient features that we expect large UAV networks to have. Due to their high altitude, long-range point-to-point communication is easy. However, long-range signal congestion is prevalent. These properties contrast with terrestrial networks where point-to-point communication tends to be short-range and congestion is localized. Such topological difference make UAV-based networks much different than ground-based networks, and necessitates a higher level of coordination.

This paper shows how multiagent evolution can be used to effectively coordinate these UAV networks. In our experiments even basic evolution (i.e. evolved using G) is shown to be helpful, performing considerably better than simple baseline algorithms. However, agents evolving to maximize the "difference evaluation function" achieve twice the level of performance. In addition agents using the difference evaluation are able to scale effectively to systems with large numbers of UAVs. These results are also shown to be robust with respect to numerous types of failures, incompatibilities and observational restrictions that will be common in real-world adhoc networks. The key to these results is that they are based on large scale UAV coordination, and will extend to other domains where similar congestion exists. For example, two way communication and UAV-to-UAV communication are natural extensions.

8. ACKNOWLEDGMENTS

This research was partially supported by NSF grants CNS-0931591 and IIS-0910358.

9. REFERENCES

[1] A. K. Agogino and K. Tumer. Analyzing and visualizing multiagent rewards in dynamic and stochastic environments. *Journal of Autonomous Agents and Multi Agent Systems*, 17(2):320–338, 2008.

[2] Y. Altshuler, V. Yanovsky, I. Wagner, and A. Bruckstein. The cooperative hunters - efficient cooperative search for smart targets using UAV swarms. *Second International Conference on Informatics in Control, Automation and Robotics (ICINCO)*, 2005.

[3] N. Andre. *Design of solar powered airplanes for continuous flight*. PhD thesis, Technische Wissenschaften, ETH Zurich, 2008.

[4] G. J. Barlow, C. K. Oh, and S. F. Smith. Evolving cooperative control on sparsely distributed tasks for uav teams without global communication. In *GECCO'08*, Atlanta, Ga, July 2008.

[5] O. Barnes, editor. *Air Power: UAVs: The Wider Context*. Royal Air Force Directorate of Defence Studies, 2009.

[6] A. Bell. Learning-based control and coordination of autonomous UAVs. Master's thesis, Oregon State University, 2011.

[7] H. Choi, Y. Kim, and H. Kim. Genetic algorithm based decentralized task assignment for multiple unmanned aerial vehicles in dynamic environments. *International Journal of Aeronautical and Space Science*, 12(2):163–174, 2011.

[8] J. M. de la Cruz, E. Besada-Portas, L. Torre-Cubillo, B. Andres-Toro, and J. A. Lopez-Orozco. Evolutionary path planner for uavs in realistic environments. In *GECCO'08*, pages 1477–1484, Atlanta, Ga, July 2008.

[9] P. Gaudiano, B. Shargel, E. Bonabeau, and B. Clough. Control of UAV swarms - what the bugs can teach us. In *AIAA*, 2003.

[10] S. Herwitz, L. Johnson, J. Arvesen, R. Higgins, J. Leung, and S. Dunagan. Precision agriculture as a commercial application for solar-powered unmanned aerial vehicles. In *American Institute of Aeronautics and Astronautics (AIAA)*, 2002.

[11] S. Karaman, T. Shima, and E. Frazzoli. Task assignment for complex uav operations using genetic algorithms. In *AIAA Guidance, Navigation, and Control Conference*, August 2009.

[12] D. Miller, P. Dasgupta, and T. Judkins. Distributed task selection in multi-agent based swarms using heuristic strategies. *Lecture Notes in Computer Science*, 2007.

[13] J. Perron, J. Hogan, B. Moulin, J. Berger, and M. Belanger. A hybrid approach based on multiagent geosimulation and reinforcement learning to solve a uav patrolling problem. In *Proceedings of the 2008 Winter Simulation Conference*, 2008.

[14] N. Picon. Effects of wireless power beaming in the space industry: Modern applications and future possibilities. *Triple Helix - Explorations in Science, Society and Law*, 2011.

[15] QinetiQ. www.qinetiq.com, 2011.

[16] M. D. Richards, D. Whitley, and J. R. Beveridge. Evolving cooperative strategies for uav teams. In *GECCO '05*, pages 1721–1728, Washington, D.C., June 2005.

[17] G. Romeo, G. Frulla, and E. Cestino. Design of solar high altitude long endurance aircraft for multi payload and operations. *Aerospace Science and Technology*, 10(6):541 – 550, 2006.

[18] H. Runge, W. Rack, A. Ruiz-Leon, and M. Hepperle. A solar powered hale-uav for arctic research. In *1st CEAS European Air and Space Conference*, 2007.

[19] T. Shima, S. J. Rasmussen, and A. G. Sparks. Uav cooperative multiple task assignments using genetic algorithms. In *2005 American Control Conference*, pages 2989–2994, Portland, OR, June 2005.

[20] K. Tumer and D. Wolpert, editors. *Collectives and the Design of Complex Systems*. Springer, New York, 2004.