Target Location among Cooperative UAVs

Roberto Galvez

CS 4960

Dr. Martin

December 8, 2010

As people strive for faster and more efficient means of accomplishing various tasks, technology in the form of computers and other tools are often used to achieve those means. More and more often, technology is not only being applied toward the goal of efficiency, but to tasks that are tedious and dangerous for humans as well. The use of aircraft known as UAVs is one such example. The acronym UAV stands for unmanned aerial vehicle. These are aircraft that are capable of flight without a crew or pilot onboard. These can be remotely piloted or be autonomous units. Several applications for UAVs exist, though the focus of UAV use tends to be military-related.

UAVs can be used for a variety of military functions, such as surveillance, target locating, target tracking, and even combat roles. Many of the benefits of UAVs come from the fact that they do not require a pilot onboard. Robert C. Nolan (1997) notes that in eliminating the presence of a pilot, the aircraft no longer requires a cockpit to hold the pilot and life support systems to keep the pilot alive in flight, which enables an aircraft to be smaller and lighter than it otherwise would be. Some important advantages of not having a pilot are that the UAV is not affected by the same physical limitations that humans are. Nolan states that due to the high endurance of a UAV, which is only limited by its fuel supply, it can perform missions of longer duration, whereas a human pilot would be limited by exhaustion and have their performance suffer over time. He also mentions that UAVs can be exposed to gravitational forces that would be unsafe for a human, as well as perform missions over hostile regions without worry of a pilot being captured or killed.

Although UAVs can prevent loss of life by removing the human element, they should also be able to perform tasks such as target location in an efficient and successful manner. For example, having a single UAV with the ability to detect targets within a one mile radius search a

2,500 square mile area would not be an effective solution, even when traveling well over 100 miles per hour. When searching a large area that is greater than the detection area of a single UAV, it is better to have several UAVs work toward accomplishing the task. Figuring out how to use several UAVs to search an area for the purpose of locating targets is a problem that has garnered attention.

As mentioned earlier, UAVs are controlled remotely or operate autonomously. The first type of UAV allow for flexibility, as they are operated by a human, and are better able to handle changes pertaining to the task at hand (Nolan, 1997). However, Nolan points out that to remotely control a UAV, a signal must be exist between the UAV and the control station. One issue with this is that to control multiple UAVs, a unique signal must exist for each UAV, but bandwidth is limited, thus inhibiting the number of UAVs able to work concurrently on the same mission. This use of a signal to send control commands to a UAV is susceptible to hostile jamming, and can result in interception of data sent by the UAV or even hostile takeover of the UAV itself.

If a group of remotely operated UAVs is not a practical solution, then the other choice is to use a group of autonomous UAVs. Patrick Vincent and Izhak Rubin (2004) categorize search plans that make use of multiple agents into four different types (Table 1). These categories depend on whether the UAVs fly predefined flight patterns or dynamically choose their paths, and whether the agents are cooperative. Vincent and Rubin note that in using a group of non-cooperative UAVs flying dynamic paths in the same target region, it is possible to have duplication of efforts. The UAVs do not share information, making it possible for several UAVs to search the same location simultaneously or immediately one after the other, resulting in a waste of resources. It is also quite possible for some subregions to go unsearched. This possibility of redundancy and leaving areas unvisited can create an inefficient and ineffective search. In order to prevent duplication of efforts and that areas go unsearched, a predefined flight path could be used. However, should a predefined flight path be used, it is possible that one or more agents will become disabled, leaving areas that are not completely searched. Vincent and Rubin state that without cooperation, the group cannot modify its search to compensate for those agents that become disabled. According to Vincent and Rubin, "The utility of using multiple agents to perform a search lies in the fact that the sensors can be used in coordination to efficiently and effectively explore a large area; allowing the UAVs to search independently of each other defeats this purpose." Vincent and Rubin argue that the aim of using multiple UAVs is to efficiently and effectively perform a search over a large area; non-cooperatively searching UAVs do not meet this goal.

Table 1.

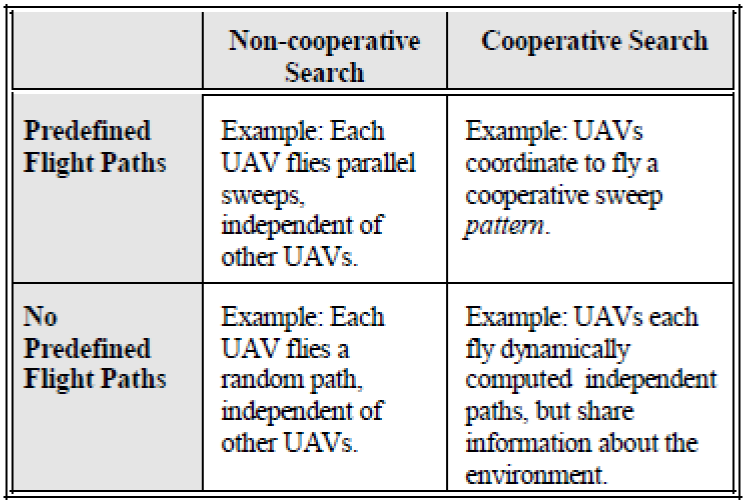

	Non-cooperative Search	**Cooperative Search**
Predefined Flight Paths	Example: Each UAV flies parallel sweeps, independent of other UAVs.	Example: UAVs coordinate to fly a cooperative sweep *pattern*.
No Predefined Flight Paths	Example: Each UAV flies a random path, independent of other UAVs.	Example: UAVs each fly dynamically computed independent paths, but share information about the environment.

(Vincent and Rubin, 2004)

If cooperation among UAVs is indeed the preferable choice, the problem now lies in choosing or devising a strategy to implement said cooperation. Various strategies for a having a group of UAVs cooperatively search for targets in a target region exist. Interestingly, some take inspiration from biological phenomena. These strategies often make use of what are sometimes referred to as “digital pheromones” (Sauter, Matthews, Parunak, & Brueckner, 2005). The strategies are inspired by the use of pheromones by animals, such as insects, to generate intelligent, complex behavior as a swarm, when as individuals they are rather simple and unintelligent. The attractive features of such a mechanism for coordination are its simplicity, scalability, and robustness (Sauter et al., 2005). In a pheromone-inspired coordination mechanism, the intelligence lies in the group of agents as a whole, and not in the individual agents themselves. They do not require the complex programming that might be required for an individually intelligent UAV. This mechanism is scalable as it requires multiple agents, and its performance generally increases with the addition of more agents. The agents are self-organizing, so the group is able to adapt to the loss of some agents. The pheromone-based strategy is seen as robust, because not only can it handle the loss of some agents, but those losses have less impact than that of more complex and possibly expensive agents.

In the strategy presented by Sauter et al. (2005), each UAV stores an artificial map in which different pheromones exist. Each of these pheromones serves the purpose of conveying particular information about the system, but all have the same basic properties. They can be deposited in a particular area, can be withdrawn from an area, are capable of accumulating, forming gradients, and they can dissipate, as well as disseminate from an area. Old information that is not refreshed becomes outdated and dissipates over time, while the dissemination of pheromones aids in spreading information. The map that each UAV stores is composed of place agents. Each of these place agents represents a geographical region. Place agents can be implemented through the use of ground sensors that communicate with local UAVs. UAV agents can start a “pheromone pump” that exists in each place agent. UAVs choose which pheromone is pumped out for some specified length of time. Place agents not only contain the pump, but also perform functions of aggregation, evaporation, and propagation of pheromones, as well as store values of each pheromone type. UAVs decide where to go by weighing the pheromone values of the place agent they are currently at, to those of each of its neighboring place agents. The strength of a pheromone at a given place agent at a given time is a function of the pheromone strength at the previous time cycle, direct deposits since then, gains through propagation of pheromone from neighbors, loss of pheromone to neighbors through propagation, and the evaporation factor. For example, if we were dealing with pheromone of type f at place agent p at time t, $s(\Phi_f,p,t\text{-}1)$, $d(\Phi_f,p,t)$, $g(\Phi_f,p,t)$, E_f, G_f, would be the strength of the pheromone at the previous time cycle, the deposits made since the last time cycle, the pheromone propagated to the place agent at time t, and the evaporation and propagation factors for that pheromone, respectively. The propagation factor determines the portion of a pheromone that is distributed from a place agent equally to all of its neighbors, while the evaporation factor determines the remaining portion of a pheromone after evaporation. The full equation for the strength of a particular pheromone at a place agent at a given time is $s(\Phi_f,p,t) = E_f, * \llcorner (1- G_f)*(s(\Phi_f,p,t\text{-}1) + d(\Phi_f,p,t)) + g(\Phi_f,p,t) \lrcorner$. This equation is calculated during each update cycle by each place agent for each pheromone type.

Using the “digital pheromones,” Sauter et al. (2005) present an algorithm for surveillance and patrolling, which can be expanded upon to create an algorithm for a target acquisitioning scenario. In the surveillance scenario, there exists an area of interest (AOI). The place agents within the specified AOI release a “Lawn” pheromone which the UAVs are attracted to. This pheromone propagates among the place agents. At some point in time a UAV encounters the attractive pheromone, and travels from place agent to place agent, each with increasing amounts of the pheromone. As the UAV visits each place agent containing the attractive Lawn pheromone, the “grass is cut,” meaning that the attractive pheromone is withdrawn, and the place agent temporarily ceases to pump out the pheromone. By altering the length of time the pump is turned off, the frequency with which the AOI is visited changes. Lengthening the time the pump is off, reduces the number of visits, while reducing the time the pump is off increases the frequency of visits. In inclusion to the removal of the Lawn pheromone, the UAV deposits a repulsive “Visited” pheromone to the area it plans to visit next. The Visited pheromone is used to prevent other UAVs from being attracted to the same location, and avoid waste of resources through duplicated efforts. In the target acquisition scenario, a UAV detects a target, but cannot identify it, and requires the help of another to do so. The UAV that detected the target deposits “NeedsID” pheromone, which is more attractive than the Lawn pheromone. Another UAV travels across place agents to increasing concentrations of the NeedsID pheromone. It eventually makes its way toward the detected target and helps identify it. Once the target is identified, the NeedsID pheromone is no longer pumped out and the second UAV deposits a large amount of Visited pheromone to counter the evaporating NeedsID pheromone. Depending on the results of the identification, the target can either be tracked or ignored. The results of an experiment using different numbers of UAVs that travel at 90 kph with random start positions on a 20 km x 20 km area that is divided into 40,000 place agents are presented in Figure 1. There are 10 targets referred to as Red units. These travel at 3 kph and occasionally rest, during which they are undetectable. During the times the targets are not at rest, Pd is the probability of detecting them when a UAV, referred to as an autonomous surveillance vehicle (ASV) in Figure 1, is in the same location.

Figure 1.

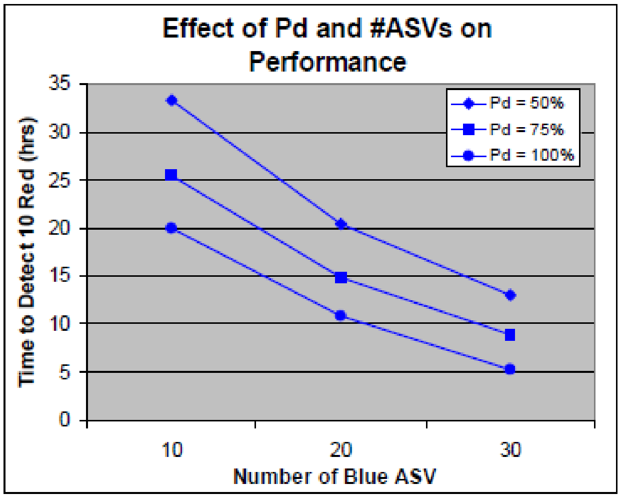

(Sauter et al., 2005)

Dasgupta, O'Hara, and Petrov (2005) also present a similar pheromone-inspired strategy. In their scenario, UAVs normally send obtained data to a central command station for identification of targets, as the computational capabilities of the UAVs is limited, while the central command has significant computational resources. However, just as the remotely operated UAVs, this situation does not scale well and is subject to malicious attacks. Therefore, a UAVs must acquire the help of *n* other UAVs, and combine resources with them to confirm the identity of a target. The number of UAVs required to help are usually determined by command according to the particular situation. The UAVs are distributed into subregions of the area to be searched at the start of the search, but are not confined to these subregions. When a UAV detects a target, it wirelessly broadcasts the pheromone value associated with the potential target via a "gossip mechanism." The UAV making the detection sends out a ping to check for nearby UAVs, who then send a ping in acknowledgement. Those nearby receiving the information, and also perform the gossip mechanism to further disseminate the information. Each UAV receiving the information updates a map similar to that of the previously presented strategy, and just as before, pheromone values decay over time. Dasgupta et al. state that to avoid having all UAVs converge onto the location of the target with the highest pheromone value, each UAV has a prioritized set of tasks corresponding to targets it is aware of. Each task consists of a target, priority value, and status, which represents whether the target is a potential or definite target. Each piece of target information is composed of a pheromone value, a location, and time the pheromone was most recently updated. Each time a UAV’s map is updated either through discovery of a target or gossiped information, the priorities of its tasks are recalculated. The priority of each task γ in a UAV’s task set is based on the destination of its current task, the current location of the UAV itself, and the distance and direction of the potential target of task γ. The function that calculates priority will only give higher priority to two kinds of tasks. One of these kinds of tasks is that corresponding to locations with significantly higher pheromone value than a UAV’s current destination. The other task is that which would reduce the distance traveled by the UAV if it traveled to that destination first, as it is on the path to the UAV’s current destination. In discovering a task of higher priority than its current task, a UAV postpones its current task for the higher priority one. As the finite storage of a UAV fills up, it starts discarding low priority tasks. When a UAV has an empty task set, such as when it is first sent to its assigned subregion, it divides its immediate vicinity into different regions and chooses the one that contains the lowest number of UAVs to move into. The UAV uses pings to determine densities of UAVs in each region. The UAV moves into the least dense region, in order to avoid possible collisions with other UAVs and to avoid duplication of efforts as well. Once a target is confirmed, the pheromone associated with it is removed from the internal map of those who identified it. No longer having pheromone information on that target, these UAVs cease to gossip information about it to other UAVs. No longer being updated, the pheromone information remaining on those that did not partake in the identification decays to zero.

If the targets are capable of movement, which is often the reality, the pheromone deposited in the area where a target is first located becomes outdated once the target moves out of that area. In order to handle this scenario, Dasgupta et al. (2005) have UAVs deposit a repelling pheromone when they arrive to the reported location of a target, but do not find a target. Pheromone is deposited as usual if the target is found in a new location.

Peng, Li, Wang, and Shen (2008) also present a biologically inspired search method for multiple UAVs. Though they use the term “digital hormone,” these are like the repulsive and attractive digital pheromones. As in the previous two examples, each UAV has a map that is continuously updated with information obtained through its sensors and from other UAVs. The environment is also broken down into a grid of different regions, with a certainty value in each cell. This value represents the certainty of the UAVs that a target exists in a particular cell region. Those areas that go unvisited have their certainty value drop during each time step. Unlike other biologically-inspired strategies, this one implements some look-ahead features. At a decision time step, a UAV takes into consideration its next q steps from its current location. At each of these steps, the UAV has the option of traveling in one of three directions, left, forward, and right. An optimal path is derived from a tree containing each possible path available in the next q steps. For example, if q were two, at time step t a UAV would consider the choices available at time step $t + 1$, as well as those available at time step $t + 2$. At $t + 1$, three options, left, right, and forward would be available, and the same would be true for each of these three possible paths (Figure 2).

Figure 2.

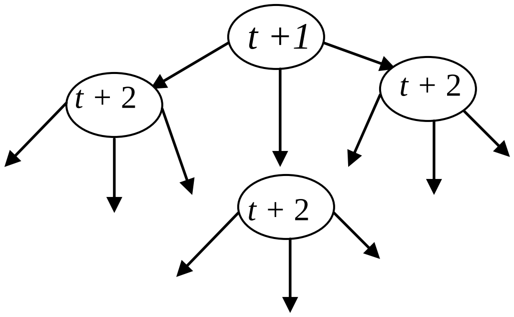

For each possible path, there exists an expected reward with an associated weight. There exist rewards for finding a target, exploring the environment, and cooperation. The environment search reward for exploration depends on the change in certainty of the visited cell. Therefore, the greatest gains for exploration would be from visiting cells with the lowest amount of certainty, which would be those that have not been visited in the longest time. In their quest to travel the paths with the highest reward values, UAVs might travel the same path. The cooperation reward exists to avoid overlapping paths. UAVs are rewarded for taking into consideration repulsive hormones. By changing the weight associated with the cooperation reward to zero, UAVs fly around disregarding pheromones. Making the value of *q* one, transforms the look-ahead into a greedy search, as at each step, the UAV's aim to move to the cell with the highest reward. Peng et al. ran several simulations comparing their method to the common search map, which lacks the use of hormones and look-ahead features. They argue that their method is superior to that of the common search map in that as more ground is covered over time, the ability of their UAVs to visit regions of higher uncertainty increases, enabling them to detect more targets (Figure 3 (b)) and obtain more information about the environment (Figure 3 (a)).

Figure 3.

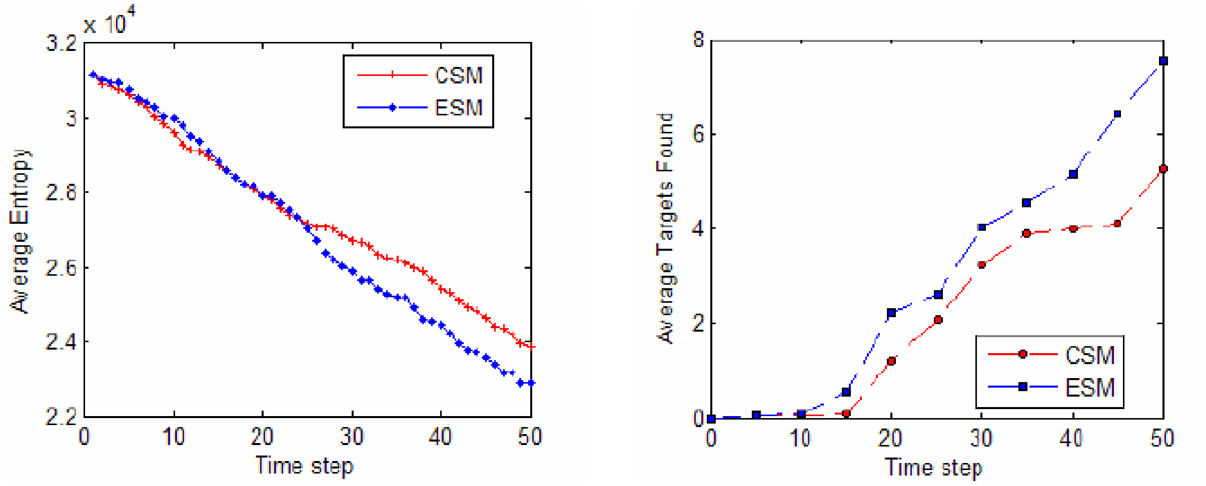

(a) (b)

(Peng et al., 2008)

As creative and effective the previous search implementations appear to be, Vincent and Rubin (2004) would probably disapprove of them. They believe that the amount of communication among a team of cooperative UAVs should be minimized, and that in the type of strategy in which UAVs periodically share the information contained in their internal maps, there is a significant amount of information being exchanged. Communication should be minimized, as bandwidth is finite and jamming can occur. They also believe that in the situation in which mobile targets exist and certainty values are used, uncertainty would continuously rise in those regions that are not being watched, because a previously empty cell might have a target move into it. In the strategy presented by Vincent and Rubin, the goal is to achieve a balance among five objectives. These are to maximize the probability of target detection, minimize the time it takes to detect the targets, minimize the number of UAVs used, be able to handle the loss of some UAVs, and minimize the amount of communication required for cooperation. In the model presented, the targets have a known max speed and have a greater range of view than the UAVs, which enables them to see a UAV and take evasive action or shoot down the UAV. However, once detected, a target cannot hide. Their search plan is one using cooperative UAVs and predefined flight paths. This kind of search can require a significant amount of communication in the situation where the pattern must be adjusted due to the loss of one or more UAVs. Each UAV must become aware of the loss and the new flight pattern. Vincent and Rubin also note that if one or more UAVs are disabled while the group is coordinating a new pattern, a surge of information exchanges would need to take place. This problem is solved by limiting the number of patterns permitted for a specific number of UAVs, as well by limiting the way they adjust their flight configurations. So, in the event that a UAV is disabled, the group knows the adjustments to be made in advance, and is not trying to choose a pattern out of an infinite number of possibilities. The UAVs exchange two kinds of messages, pattern maintenance messages and pattern update messages. The first kind of messages are used to maintain a flight pattern and check for absent UAVs, while the second kind are used to inform about the loss of one or more UAVs. The UAVs perform a parallel-path search as a single group, which is to facilitate communication and for the role of a disabled UAV to be taken over (Figure 4).

Figure 4.

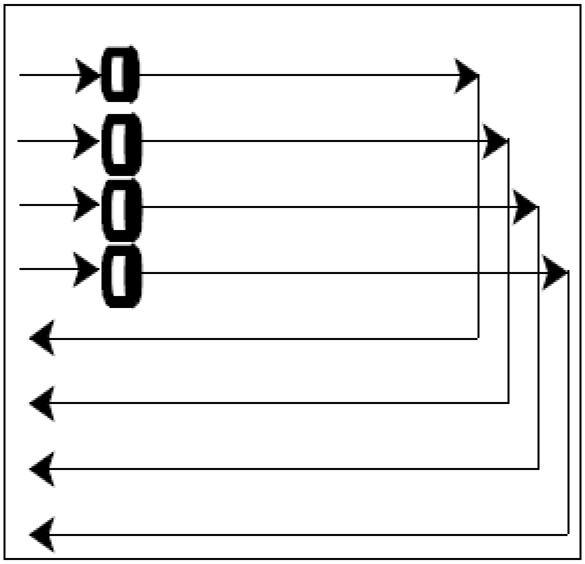

(Vincent and Rubin, 2004)

The UAVs are separated by a distance that is twice their search radius, which means that all targets within their sweep width will be detected, as no gaps exist. However, there exists the problem that a target might move into the already searched area as the UAVs travel from one edge to the other and back. In order to avoid this, there will be some overlap as the UAVs return to an edge.

As seen, different approaches to using multiple UAVs to search for targets exist. As with other tasks, there is no single correct way of doing it. Not only that, but the many ways that exist can be further improved upon. For example in the simulations run by Peng et al. (2008), it is assumed that there is no noise to interfere with communications. They mention that their strategy needs to be able to handle imperfect communication, as it would affect whether the UAVs all have the same updated map and their coordination. Dasgupta et al. (2005) are searching for a means of having UAVs determine appropriate UAVs with whom to share gossiped information with. Another concept they are interested in is finding a means in which the group of UAVs dynamically divide the AOI into subregions as they perform their search, so as to avoid dense groups of UAVs.

According to Nolan (1997), a key interest in UAVs is their ability to avoid loss of human life, and that saving human life is critical in military operations conducted by the United States. It is likely that this will continue to be the case in the future. As with other technological fields, will probably continue to expand and be improved upon.

References

Dasgupta, P., O’Hara, S., and Petrov, P. 2005. A Multi-agent UAV Swarm for Automatic Target Recognition. In *Proceedings of Defence Applications of Multi-Agent systems International Workshop* (DAMAS ’05), Thompson, G. S., Ghanea-Hercock, R. (Eds.). Springer-Verlag, Berlin, Heidelberg, 80-91. DOI= 10.1007/11683704_7 http://dx.doi.org/10.1007/11683704_7

Nolan, R. C. (1997). *The Pilotless Air Force?* (Unpublished doctoral dissertation). Air Command and Staff College. Retrieved December 01, 2010, from http://www.fas.org/irp/program/collect/docs/97-0530.htm

Peng, H., Li, Y., Wang, L., and Shen L. 2008. Hormone-Inspired Cooperative Control for Multiple UAVs Wide Area Search. In *Proceedings of the 4th international conference on Intelligent Computing: Advanced Intelligent Computing Theories and Applications - with Aspects of Theoretical and Methodological Issues* (ICIC '08), Huang, D., Wunsch, C. D., Levine, S. D, and Jo, K.(Eds.). Springer-Verlag, Berlin, Heidelberg, 808-816. DOI=10.1007/978-3-540-87442-3_99 http://dx.doi.org/10.1007/978-3-540-87442-3_99

Sauter, J. A., Matthews, R., Parunak, H., and Brueckner, S. A. 2005. Performance of digital pheromones for swarming vehicle control. In *Proceedings of the Fourth international Joint Conference on Autonomous Agents and Multiagent Systems* (AAMAS '05). ACM, New York, NY, 903-910. DOI= 10.1145/1082473.1082610 http://dx.doi.org/10.1145/1082473.1082610

Vincent, P. and Rubin, I. 2004. A framework and analysis for cooperative search using UAV swarms. In *Proceedings of the 2004 ACM Symposium on Applied Computing* (SAC '04).

ACM, New York, NY, 79-86. DOI= 10.1145/967900.967919

http://dx.doi.org/10.1145/967900.967919