

Better Flocking Behaviors in Complex Environments using Global Roadmaps

O. Burchan Bayazit Jyh-Ming Lien Nancy M. Amato
burchanb@cs.tamu.edu neilien@cs.tamu.edu amato@cs.tamu.edu

Technical Report TR02-003
PARASOL LAB
Department of Computer Science
Texas A&M University

May 6, 2002

Abstract

Flocking behavior is very common in nature, and there have been ongoing research efforts to simulate such behavior in computer animations, robotics and artificial life applications. Generally, such work considers behaviors that can be determined independently by each flock member solely by observing its local environment, e.g., the speed and direction of its neighboring flock members. Since flock members are not assumed to have global information about the environment, only very simple navigation and planning techniques have been considered for such flocks.

In this work, we investigate how the addition of global information in the form of a roadmap of the environment enables more sophisticated flocking behaviors. In particular, we study and propose new techniques for three distinct group behaviors: homing, exploring and shepherding. These behaviors exploit global knowledge of the environment and utilize knowledge gathered by all flock members. This knowledge is communicated by allowing individual flock members to dynamically update the shared roadmap to reflect (un)desirable routes or regions. We present experimental results showing how the judicious use of simple roadmaps of the environment enables more complex behaviors to be obtained at minimal cost. Animations of these behaviors can be viewed at <http://www.cs.tamu.edu/faculty/amato/research/dsmft>

*This research supported in part by NSF CAREER Award CCR-9624315, NSF Grants IIS-9619850, ACI-9872126, EIA-9975018, EIA-0103742, EIA-9805823, ACI-0113971, CCR-0113974, EIA-9810937, EIA-0079874, and by the Texas Higher Education Coordinating Board grant ARP-036327-017. Bayazit is supported in part by the Turkish Ministry of Education.

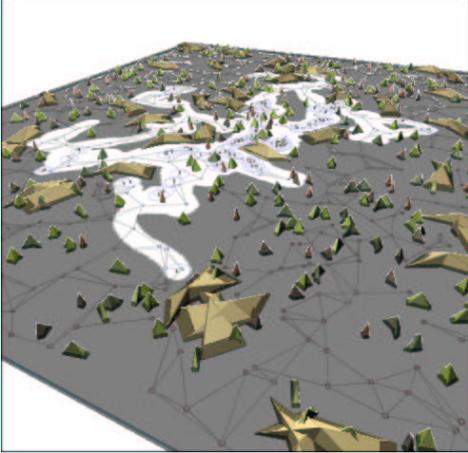


Figure 1: Complex behavior, such as covering and shepherding, can be produced using global information.

Introduction

The ability to simulate the coordinated movement of a group of creatures plays an important role in Artificial Life. For example, birds fly in flocks, fish swim in schools, and sheep move as a herd. An artificial representation of such creatures requires techniques for generating the motion of the individual entities within the flock and techniques for directing the global movement of the flock. Due to their importance, such techniques have attracted the attention of many researchers (see related work section). However, most of the research has focused on techniques for modeling individual behavior within a flock, such as Reynolds’ *boids* (Reynolds 1987). These techniques have been coupled with simple methods for guiding global flock movement, such as attractive potential fields centered at a goal location, to achieve realistic group movement in simple environments with few external obstacles, such as air or sea for birds and fish. Unfortunately existing flocking methods do not perform well if complex navigation is required, such as in cities, through crowded rooms, or over rough terrain.

In contrast, path planning algorithms developed in the robotics community are capable of navigation in complex environments (Latombe 1991). In particular we note the *roadmap methods* which can quickly answer many diverse path planning queries in the same environment using a map, typically constructed during preprocessing, containing a network of representative feasible paths in the environment. In essence, these maps function similarly to driving maps in that one plans a route by first locating their initial and final positions and then selecting a route connecting them from the roads and highways shown on the map. While many good path planning algorithms exist, they have traditionally only been used to plan paths for a single moving object (the ‘robot’). That is, they have not been customized to support coordinated group

behavior.

In this work, we explore the benefits of integrating flocking techniques with roadmap-based path planning methods. We find that the global navigation information provided by the roadmaps can also be exploited to support more sophisticated group behaviors than possible using traditional (local) flocking methods. In particular, we consider three basic behaviors: homing, exploring, and shepherding. Our new techniques can be applied to an entire flock, to individual flock members, or to an external agent that may influence the flock (e.g., a sheep dog).

In the *homing* behavior, the goal is to move the entire flock from a starting position to a goal position. The individual flock members follow Reynolds’ ‘*boid*’ dynamics, and the flock’s global motion is planned using the roadmap of the environment.

We study two types of *exploring* behavior. In the first, we consider a *covering* behavior in which individual flock members explore the environment attempting to visit all points of the scene. This behavior is useful, e.g., for mine sweeping. The second exploring behavior requires the individual members to first explore the environment searching for a goal whose position was not known *a priori*, and then, after the goal is located, all the flock members should move to the goal. This can be used, for example, to model foraging behavior in ants. This second exploring behavior illustrates another feature of our roadmap-based methods. After the goal’s location is found, it is communicated to the rest of the flock by modifying the shared roadmap, and in particular by updating the edge weights appropriately. Thus, in addition to providing a convenient data structure for storing and searching for paths, the roadmap can also be used as a communication mechanism among flock members.

The final behavior we study is *shepherding*, in which the flock is steered by an external agent (a dog). It associates a repulsive potential with the shepherd coupled with a fairly standard flocking model for the herd. This behavior has applications in robot collaboration or robotic guarding problems.

In the next section we give a summary of related work. In Section 3, we briefly describe the *probabilistic roadmap* (PRM) motion planning methods used to construct the roadmaps we use in our new flocking techniques. The three different flocking behaviors we consider are discussed in Section 4. Section 5 contains experimental results and we present our conclusions and discuss future work in Section 6.

Related Work

Reynolds’ influential flocking simulation (Reynolds 1987) established the feasibility of modeling such a system. His work showed that flocking is a dramatic example of emergent behavior: where global behavior can

arise from the interaction of simple local rules. Each individual member of the flock (boid), has a simple rule set stating that it should move with its neighbors. This concept has been used successfully by researchers both in computer graphics and robotics. Yu and Terzopoulos (1994) used flocking behaviors with intention generators to simulate a school of fish in artificial life. An intention generator is similar to a finite state machine that generates emergent behaviors based on current or past states. Later, they generalized the concept of intention generators to cognitive modeling (Funge, Yu, & Terzopoulos 1999). Nishimura and Ikegami (1997) used flocking dynamics to investigate collective strategies in a “prey-predator” game model. Ward et al. (2001) studied an evolving sensory controller for producing schooling behavior based on “boids”. Brogan and Hodgins (1997) investigated group behavior with significant dynamics, such as human-like bicycle riders. Sun et al. (2001) achieve swarm behaviors based on a biological immune system. Balch and Hybinette (200) propose a behavior-based solution to the robot formation-keeping problem. Fukuda et al. (1999) describe group behavior for a Micro Autonomous Robotics System. Mataric (1994) classifies a basic set of group behaviors which can be used to create more complex behaviors including flocking and herding. Saiwaki et al. (1997) use a chaos model to simulate a moving crowd. An interesting approach by Vaughan et al. (2000) used a robotic external agent to steer a flock of real geese.

Parker’s investigation (Parker 1993) of the use of global knowledge to model more complicated behavior supports our approach. She concluded that global knowledge should be used to provide general guidance for the longer-term actions of an agent, whereas local knowledge influences the more short-term, reactive actions. She also suggested that local information should be used to ground global knowledge in the current situation. This allows agents to remain focused on the overall goals of their group while reacting to the dynamics of their current situations.

Research on flocking behavior has been successfully applied to the entertainment area. For example, the behavior of individual characters in several animated films is based on flocking behavior. There is also an interest in flocking behavior in the game industry where the behavior of artificial characters is defined by flocking rules (Reynolds 1999). For example, real-time strategy (RTS) games usually involve moving a group of creature-like objects to accomplish tasks, like gathering resources and attacking opponents. Motion planning is required to find paths for these creatures to reach the user specified goal. The most popular motion planning method used in the game industry is A^* search (Russell & Norvig 1994) on a grid-based representation of the environment.

Although there is little research on path planning for

flocks, much work has been done on the problem of planning for multiple robots. Motion planning methods for multiple robots fall into two categories: centralized and decoupled. Centralized methods consider all robots as one entity whose degree of freedom (DOF) is the summation of the DOF of all robots. Decoupled methods first find a path for each robot separately and then resolve conflicts between these paths. Although centralized approaches can be complete and decoupled methods usually are not, decoupled methods are usually much faster and more practical. In (Li, Jeng, & Chang 2001), each group of crowds is guided by a leader and the paths of the leaders are generated using a decoupled approach. Unfortunately, followers can be trapped in local minima of their leader’s potential field.

Flocking behavior can also be used for optimization. The ant colony optimization (ACO) (M., Caro, & Gambardella 1999) meta-heuristic is a group of algorithms for discrete optimization. The method is inspired from observation of the behavior of ant colonies. Ants find shortest paths between food sources and their hill using pheromones. Ants smell pheromone and tend to select the route with higher pheromone concentration. Because shorter paths gather pheromone faster than longer paths, the ants eventually converge on the shortest path. Dorigo et al. (1999) exploit this ant-like behavior to optimize solutions for NP-Complete problems, such as the traveling salesman problem, network routing, quadratic assignment, and job scheduling. In our work, the flock’s ability to explore comes from combining a roadmap of the environment with the ACO concept. Using ACO-like heuristics to adaptively adjust roadmap edge weights enables us to represent dynamic environmental information.

Roadmap-Based Path Planning with PRMs

Given a description of the environment and a movable object (the ‘robot’), the motion planning problem is to find a feasible path that takes the movable object from a given start to a given goal configuration (Latombe 1991). Since there is strong evidence that any complete planner (one that is guaranteed to find a solution, or determine that none exists) requires time exponential in the number of degrees of freedom (DOF) of the movable object (Latombe 1991), attention has focused on randomized or probabilistic methods.

As previously mentioned, our approach utilizes a roadmap encoding representative feasible paths in the environment. While noting that our techniques could use any roadmap, our current implementation is based on the probabilistic roadmap (PRM) approach to motion planning (Kavraki *et al.* 1996). Briefly, PRMs work by sampling points ‘randomly’ from the robot’s configuration space (C-space), and retaining those that satisfy cer-

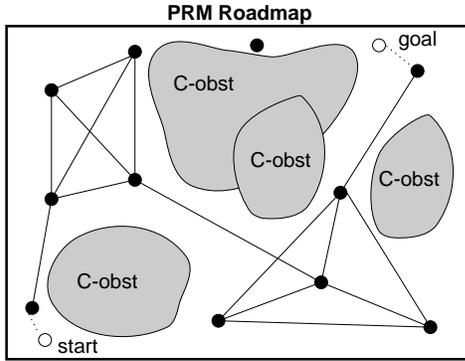


Figure 2: Querying a PRM roadmap (C-space).

tain feasibility requirements (e.g., they must correspond to collision-free configurations of the movable object). Then, these points are connected to form a graph, or roadmap, using some simple planning method to connect ‘nearby’ points. During query processing, the start and goal are connected to the roadmap and a path connecting their connection points is extracted from the roadmap using standard graph search techniques (see Figure 2).

We use a particular variant of the PRM called the Medial-Axis PRM, or MAPRM (Amato *et al.* 1998). In MAPRM, instead of generating the nodes uniformly at random in C-space, they are generated on or near the medial-axis of C-Space. MAPRM is particularly well suited to flocking behavior since roadmap nodes tend to maximize clearance from obstacles.

Group Behaviors - Homing, Exploring, Shepherding

Homing Behavior

Homing behavior consists of two sub-models, one representing the individual behavior of flock members and the other influencing the global behavior of the flock. “Boid” dynamics (Reynolds 1987) sufficiently model individual behavior in most cases. In this model, individual members should: (i) avoid collision with neighboring flockmates, (ii) match velocity with them, and (iii) stay close to their neighbors. The neighborhood is defined by a distance, and an individual member of the flock is steered by angle and directional vectors satisfying the above criteria. These individual behaviors can be seen in Figure 3.

Global behavior is usually simulated by adding another directional vector toward a goal based on a potential field method (Khatib 1986). However, this method may easily be trapped in a local minimum in an environment crowded with obstacles. A method commonly used in computer games requiring motion of a group of objects is a grid-based A^* search (Russell & Norvig 1994). In this approach, the environment is discretized to small grid cells and the search for the flock’s path is

based on expanding toward the most promising neighbor of already visited positions. Although A^* search finds shortest paths and it is usually fairly fast, it does have drawbacks. Of particular note here is the necessity of finding a completely new path for each new goal which reduces the efficiency of this approach and increases the computation time for complex environments.

In contrast, roadmap-based path planning methods do not have such drawbacks. They work on the global scale and once the roadmap is generated, finding new paths is fast and efficient. In our approach, we use MAPRM to find a path for the flock. One of the advantages of MAPRM is that the paths we find tend to have large clearances from obstacles. Once a path is found, individual flock members follow the path. The path is discretized to subgoals based on an individual flock member’s sensor range. Each member keeps track of subgoals and as soon as a subgoal comes within the sensory range the next subgoal becomes the steering direction for the global goal.

With other interacting forces from neighboring flock members and obstacles, steering toward the subgoal has the lowest priority, so individual members still move together while moving toward goal. This results in a flocking toward the goal and avoids getting trapped in local minima.

ALGORITHM FOR HOMING

```

01. for (each individual flock member)
02.   if (goal is in its view range)
03.     stay around goal
04.   else if (current subgoal is in its view range)
05.     set next subgoal as the target
06.   else
07.     steer toward the target
08.   endif
09. endfor

```

Exploring Behavior

There are several variants of exploring. For example, flocks may have good *a priori* data about the environment and they may be looking for objects whose locations are unknown, or the environment may be unknown and the objective of the exploring behavior is to gather information about the environment. For our work, we assume the environment is known. We consider two different exploring behaviors: (i) covering the environment, and (ii) searching for a goal and moving toward the goal once it is found. To achieve these behaviors, we use a roadmap graph with adaptive edge weights. In this approach, each individual member behaves independently from its flock mates and uses the roadmap to wander around. Specifically, they follow roadmap edges and there are no predefined paths. If they reach a roadmap node with several roadmap edges, they choose a roadmap edge to follow based on the weight of the edge. The edge weights represent how relevant the edge is to the current

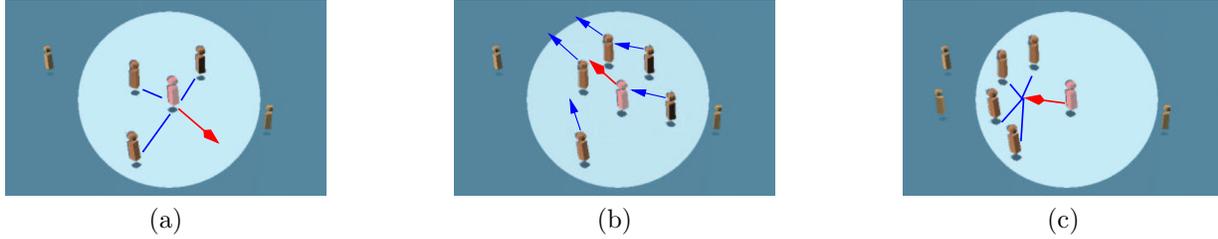


Figure 3: Individual member behavior for flocks. (a) Separation: avoid crowding neighbors. (b) Alignment: match velocity of neighbors. (c) Cohesion: stay close to neighbors. The red line represents the steering direction.

task, either covering the environment or searching for and reaching the goal.

Covering the environment In this behavior, we want some member of our flock to have visited every location in the environment. We assume we start with a roadmap covering all relevant portions of the environment. The goal is to have some flock member visit every edge and vertex of the roadmap. Initially, edges all have weight one. Flock members always move along roadmap edges. As they traverse a roadmap edge they increase its weight. This is similar to ant pheromones which increase as more ants follow the same path. Since our goal is to explore the environment, the individual flock members are biased toward relatively unexplored areas of the roadmap. This is achieved by having them select roadmap edges with smaller weights with some higher probability. The algorithm is shown below.

ALGORITHM FOR COVERING THE ENVIRONMENT

```

01. for (each flock member )
02.   while (not all nodes visited)
03.     if (not in the roadmap)
03.       move to the closes roadmap node
04.     if (current node do not have an outgoing edge)
05.       pop memory until an new branch is found
04.     else
04.       probabilistically pick a lower-weight edge
03.       increase edge weight
09.       push this node to memory.
10.   endwhile
11. endfor

```

Searching for an unknown goal Our goal searching behavior is similar to *ant colony optimization (ACO)*. Although the individual flock members know the environment, they don't know the location of the goal. If an individual reaches a location where the goal is within the sensory range, all other members try to reach the goal. Like the previous case, we implemented this behavior using adaptive roadmap edge weights. The weight of an edge shows how promising a path segment is. Again, the member chooses an edge to leave a roadmap node with some probability based on the edge's weight. As an individual traverses a path in the roadmap, it remembers the route it has taken. Then, when it reaches a goal, it

increases the weight of the edges on the route it took. If the individual reaches a roadmap node without any outgoing connections (i.e., with only one edge) or a node already contained in the current path (i.e., a cycle), the weight of the edges it followed will be decreased.

We also investigated the case where the goal is allowed to move. Searching for moving goals is similar to searching for non-moving goals. Two stages are classified in searching for moving goals: finding the goals and staying near goals. Finding moving goals is exactly the same as searching for non-moving goals. However, since goals are moving, flock members need to keep goals within sight and, also, update global information to tell other members which direction is better to go. Unlike non-moving goals, edge weights will quickly become out of date as goals move away from their original positions. To solve this problem, we interpret all "solutions" as suboptimal and use evaporation as in *ACO* to prevent flock members from being trapped in these "local minimum."

The algorithms for both approaches are summarized below.

ALGORITHM FOR NON-MOVING GOALS

```

01. for (each flock member)
02.   if (it finds goal)
03.     increase edge weight according to its memory
04.   else if (it finds a dead end)
05.     pop memory until an new branch is found
06.     decrease weight according to popped memory
07.   else
08.     select an neighboring node from current node.
09.     push this node to memory.
10.   endif
11. endfor

```

ALGORITHM FOR MOVING GOALS

```

01. for (each flock member)
02.   if (goal is not in view range)
03.     Find goal as in Alg. for Non-Moving Goals
04.   else
05.     find closest node to the goal
06.     update weight of incoming edges of this node
07.   endif
08. endfor
09. decrease every node weight

```

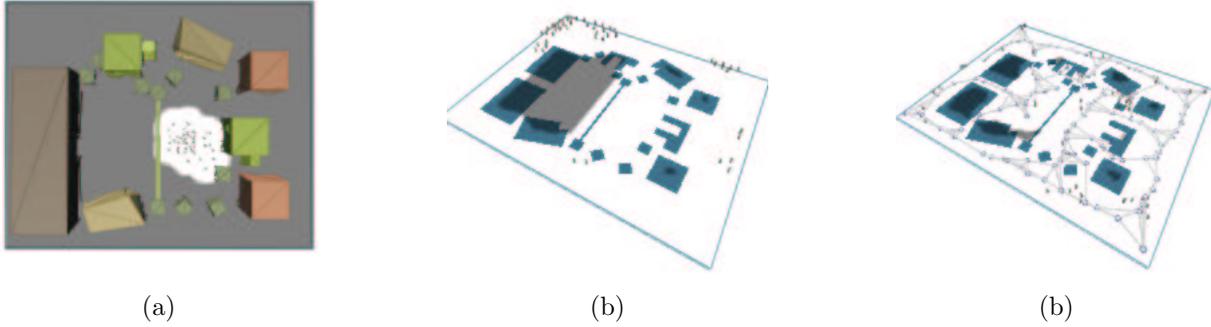


Figure 4: Covering behavior. The dark area represents unexplored regions. (a) Initial configuration of flock. (b) Normal flocking system. (c) Flocking system with roadmap-based method.

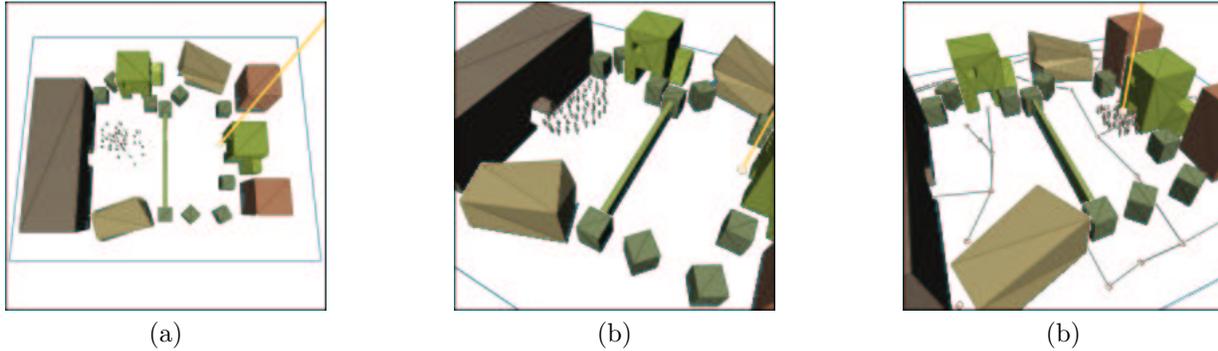


Figure 5: Searching for a goal behavior. There are 80 boids in the scene and the white tower on the right is the goal. (a) Initial configuration of flock. (b) Normal flocking system. (c) Flocking system with roadmap method.

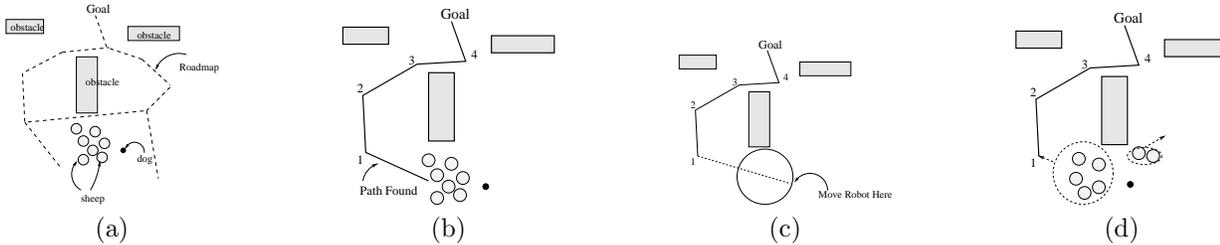


Figure 6: Shepherding: the sheep are represented by large circles and the dog is a small dark circle. (a) Roadmap, (b) dog finds a path, (c) dog's steering location, (d) flock is separated.

Shepherding Behavior

In the previous sections we have observed two distinct flocking behaviors. In the first case, the flock members were moving toward a goal together, i.e., as a flock. The motion was planned for the flock. In the second case, the flock members were exploring and planning their motions individually. In a sense, the flock had control of the motion in the first case and individual flock members had control in the second case. In our third scenario, neither the flock nor the individuals have control of the motion. Instead, an outside agent guides or shepherds them. In the simulation snapshots shown in Figure 6(a), the external agent is a dog whose objective is to move the flock of sheep toward the goal. The only motion control for the flock is to move away from the dog. A similar implementation has been done by Vaughan et al.

(2000) where a robot was programmed to move geese toward a goal position. We would like to implement a similar algorithm where a subgoal will be a roadmap node found in the path. Until the subgoal is reached, the robot will move toward that goal and then will choose the next roadmap node on the path as the next subgoal (see Figure 6(b)).

To move the flock toward the goal, the dog steers the flock from behind (Fig. 6(c)). If any subgroup separates from the flock, it is the dog's job to move the subgroup back to the flock (Fig. 6(d)).

ALGORITHM FOR SHEPHERDING (for dog)

01. Find a path on roadmap
02. **while** (not goal reached)
03. Select the next node on the path as subgoal
04. **while** (not subgoal reached)

```

05.   Move behind flock on opposite side of subgoal
06.   if flock separates
07.     Move the farthest subgroup from subgoal
        toward other subgroup
08.   endif
09. endwhile
10.endwhile

```

Experimental Results

In this section we evaluate our roadmap-based techniques for the homing, exploring, and shepherd-ing behaviors described previously. Movies illustrating these behaviors can be found on our webpage: (<http://www.cs.tamu.edu/faculty/amato/dsmft/research>).

Our experiments are designed to compare our roadmap-based techniques with more traditional approaches for simulating flocking behavior and to study the improvements possible by incorporating global information about the environment as encoded in a roadmap.

To study the efficiency of our exploring techniques, we also compare our roadmap-based techniques with ‘ideal’ variants which have complete knowledge of the environment and the current status of the search. For example, in the goal searching behavior, the location of the goal is known at all times in the ideal variant.

All of our experiments were run on a Linux system with Athlon 1.33 processor and 256MB memory.

Homing Behavior

For the homing behavior, our roadmap-based technique is compared with a basic flocking behavior using a potential field and a grid-based A^* search behavior.

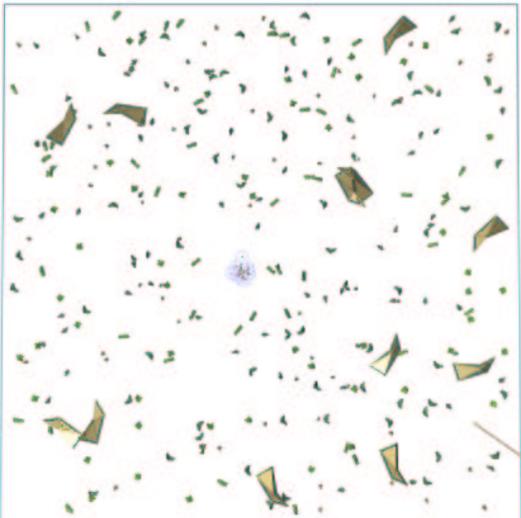


Figure 7: Environment for homing experiments.

The environment is a square with sides measuring 420 meters which contains six types of obstacles (see Figure 7). A total of 301 obstacles are randomly placed in

the environment. At any given time there is one goal, and when all flock members reach it, a new goal is randomly generated; this process continues until eight goals have been generated and reached. The experiment involves 40 flock members, which are initially placed according to a Gaussian distribution around the center of the square environment. The simulation is updated every 100 ms.

For the grid-based A^* behavior, a bitmap of the environment of 914×914 cells is constructed; the length of a side of each square cell is equal to the diameter of a flock member. Cells are classified as free cells and collision cells. Paths are found in this bitmap using A^* search. For the roadmap-based behavior, the roadmap is built using the MAPRM method to generate 400 roadmap nodes and we attempt to connect each node to its 4 nearest neighbors.

METHOD	#flockmate reaching the goal
Basic	10
grid-based A^*	40
Roadmap-based	40

Table 1: This table shows the number of the 40 flock members that reach their home within 30 seconds using the basic flocking behavior, the grid-based A^* behavior, and the roadmap-based behavior.

Table 1 shows that, without global information, only 25% of the flock members reach the goal and most of them are trapped in local minima. On the other hand, when global navigation information is utilized, either with the grid-based A^* method or our roadmap-based method, all flock members reach the goal.

BEHAVIOR METHOD	init time	find path time	local minima	
			#	escape (s)
roadmap-based	0.88	0.652	255	22.99
grid-based A^*	6.02	5.757	2005	1035.43

Table 2: This table shows the time for initialization, the average time to find a path, and the total time spent by all flockmates escaping local minima.

In Table 2 we show the time spent searching for paths, the number of local minima encountered along all paths, and the total time spent escaping from local minima. This offers some insight into the methods studied, as can be seen more clearly in Figure 8. Although the flock takes a shorter path with the grid-based A^* search than with the roadmap-based method (Figure 8(a)), the flock reaches the final goal faster with the roadmap-based method (Figure 8(b)). As A^* search is known to be fast and always finds shortest paths, this example illustrates that our roadmap-based method indeed is a competitor for grid-based A^* methods – while the paths found are a bit longer, they are found faster.

Exploring Behavior

We next consider the two types of exploring behaviors – covering the environment and searching for a goal. For both behaviors, we compare our roadmap-based methods with very basic methods (random walks) and with ‘ideal’ variants of the roadmap-based methods which have complete knowledge of the search status at all times. For example, *a priori* knowledge of the location of the goal. The former comparison enables us to establish the value of the global information stored in the roadmap, while the latter comparison enables us to determine how far our methods are from achieving optimal performance.

Covering the environment Space covering is tested on the environment shown in Figure 4, which requires flock members to pass through narrow passages to access undiscovered areas. In this experiment, we compare basic flocking behavior, roadmap-based behavior, and an ideal variant of the roadmap-based behavior which has dynamic knowledge of the undiscovered regions.

The environment (80×100) is populated with 16 obstacles (6 types of obstacles) and in total 24% of the environment is occupied by obstacles. 50 flock members are simulated and states are updated every 100ms. A bitmap is built to record discovered/undiscovered information. A bitmap cell is discovered when it is inside the sensory range of any flock member. We set the radius of the sensory circle as 5m. For the roadmap, 120 nodes are sampled and connections are attempted to each node’s 4 nearest neighbors.

The roadmap-based covering behavior is as described previously.

The basic behavior uses only local information, and is essentially a random walk through the environment. It shows that the lack of global knowledge results in some areas never being discovered, especially those surrounded by obstacles.

The behavior with perfect knowledge of the undiscovered locations uses the roadmap to find paths from a flockmate’s current position to the closest unexplored spot. Although such knowledge would not be available in this application, this variant gives us an idea of how fast an environment can be covered in the best case.

As seen in Figure 9(a), the perfect behavior rapidly covered almost 91% of the environment in the first 30 seconds. The roadmap-based behavior, using indirect communication (adaptive edge weights) takes about three times as long (90 seconds), to reach a similar coverage point of 91.6%. Nevertheless, like the perfect behavior, the roadmap-based behavior found most reachable areas. In contrast, the basic flocking behavior had difficulty covering more than 80% of the environment. However, it is interesting to note that the basic flocking behavior found more undiscovered areas than the roadmap-based approach in the first 40 seconds; this is due to the flockmates with basic behavior which tend to bounce around

and discover ‘easy’ areas very quickly.

Searching for a goal In this experiment, the roadmap-based behavior is compared with simple flocking behavior that has only local information about the environment and no knowledge of the goal position and with an ideal variant of the roadmap-based behavior that has *a priori* knowledge of the position of the goal. The environment is the same as that used in the covering experiment (See Figure 5).

We are interested in the number of flock members that reach the goal and how fast they get there. As mentioned before, the behavior with complete knowledge is used to establish a best case (lower bound) for the simulation efficiency, and the basic behavior using only local information is used to illustrate the importance of global knowledge. The results of some experiments are shown in Figure 9.

In Figure 9(b), we see that flocks using the basic behavior do not discover any goals within 35 seconds. Actually, what happens is that none of the flock members discover the narrow passage out of the confined region in which they start.

Overall, the roadmap-based behavior is competitive with the ideal roadmap-based behavior – only taking 5 seconds longer than the method in which the position of the goal is known *a priori*. In addition, it is surprising to note that two of the flock members in the roadmap-based method reach the goal earlier than any of their flockmates in the ideal roadmap-based behavior. While we expect the roadmap-based method to continue to perform well in more complex environments, we expect its efficiency relative to the ideal method to decline somewhat.

Shepherding Behavior

In shepherding, we try to investigate the following questions: (i) How does the roadmap-based method perform as compared with grid-based A^* search? (ii) How does the magnitude of the dog’s repulsive force affect the simulation. To answer these questions we analyze two sets of experiments.

In the first set of experiments, the flock consists of 30 sheep and the sheep dog is an external agent. The experiment starts the flock in a random location and the objective is to move it to a randomly selected goal. When the sheep reach the goal, the experiment is repeated again by selecting a new starting position and goal position at random.

We compare our method to a grid-based A^* method. The basic method using local information was not considered due to its observed inadequacies in the previous experiments. The environment is the same and the experiments are similar as for the homing experiments (Figure 7). In our grid-based A^* implementation, the search for the path to the goal and the dog’s path to

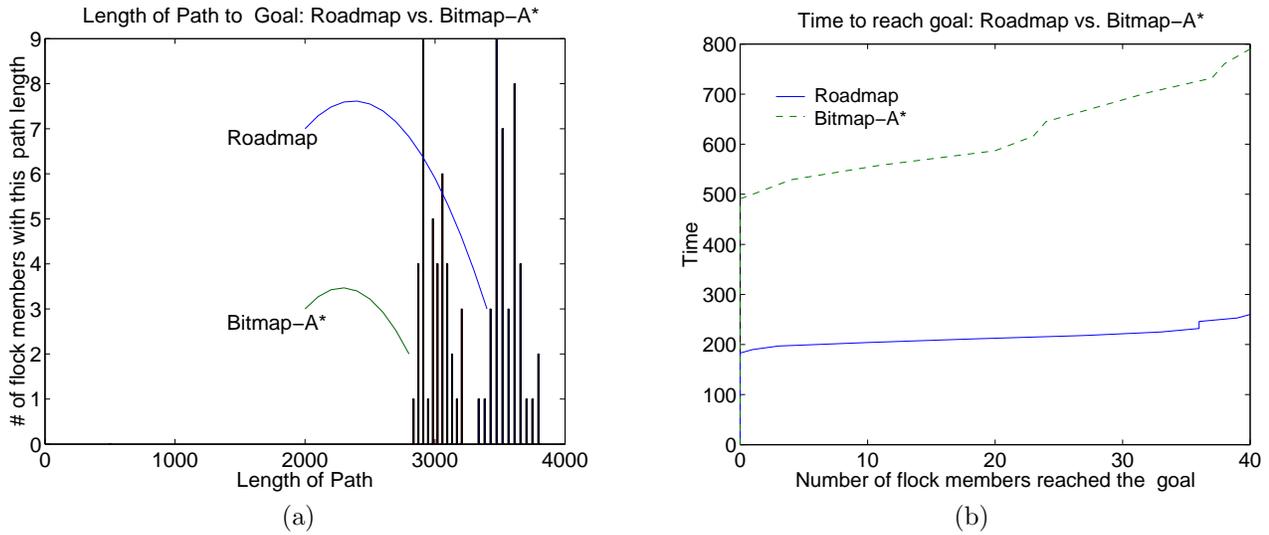


Figure 8: Homing behavior: (a) the number of flock members reaching goals with respect to the length of the paths they took, (b) the number of flock members reaching homes with the evolution of time. Although the grid-based A^* behavior finds shorter paths, the flock spends less time to reach the goals with the roadmap-based behavior.

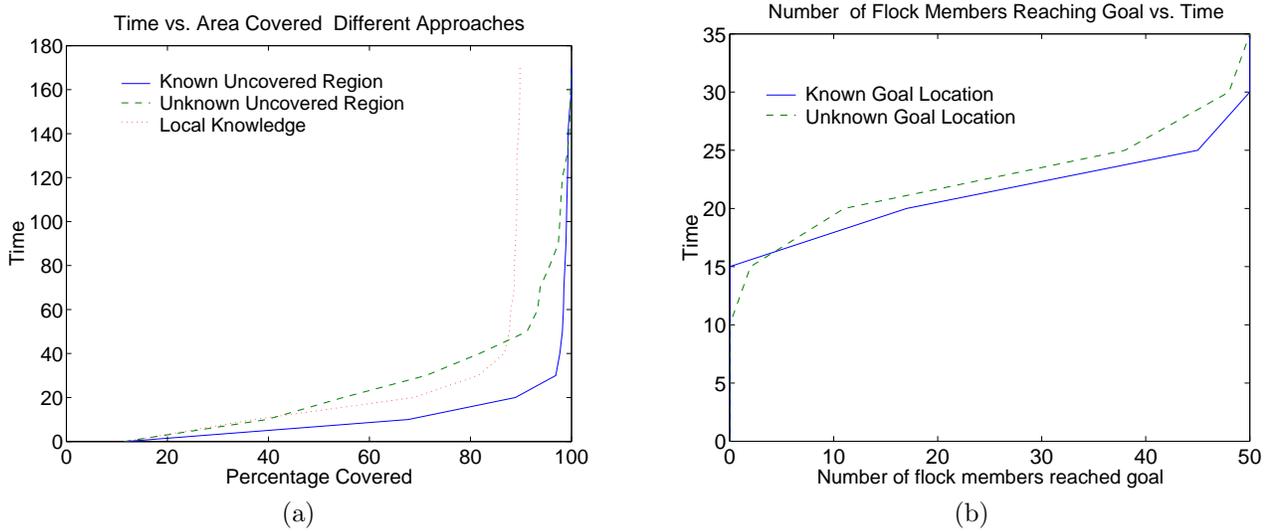


Figure 9: Covering behavior: (a) the percentage of the environmental covered in terms of time, (b) the number of flock members reaching the goal area in terms of time.

Shepherding behavior: Roadmap v.s. A^*

BEHAVIOR METHOD	Init time	Number of step to reach the Goal	Avg. Local minimum encountered
Roadmap-based	0.88	2348.17	7.8
A^* -based	6.02	10612.08	32.2

Table 3: This table shows time for initialization, the average number of simulation steps required to reach the goal, and the average number of local minima encountered.

the steering position use A^* search on a bitmap with 914×914 cells. The roadmap method used the MAPRM method with 400 roadmap nodes, with each node connected to its 4 nearest neighbors.

Our results, shown in Table 3, include initialization times, the number of simulation steps required to reach the goal, and the number of local minima encountered. All values reported are averages over 40 experiments. We see that fewer local minima are encountered here than in the previous behaviors. This is influenced by the fact that MAPRM tries to generate paths that have high clearance from obstacles and because as the dog moves to the steering position, it effects the individual members so it increases the entropy of the system, resulting in relatively more randomness. Thus, even though some members are stuck in local minima, the dog would come and retrieve them. The table shows that similar to homing behavior, our shepherding with the roadmap approach performed better than the grid-based A^* search.

In the second set of experiments, we have varied the repulsive force coefficient of the dog 0.1 and 150. The coefficient represents how much the sheep would be repulsed from the dog if it is nearby. Please note that, although the greater the coefficient the more repulsion occurs, it does not mean that the total force on the sheep is proportional to the coefficient since there are other forces on the individual sheep (to represent boid behavior). Figure 10 shows the results of this experiment. At the beginning, there is very low repulsive force so it takes longer for the dog to push the sheep to the goal. As the repulsive force increases, the dog controls the sheep better. However, after some point, the system becomes chaotic. This behavior can be explained by the fact that as the sheep become more repulsed, it becomes difficult for the dog to collect them into the herd.

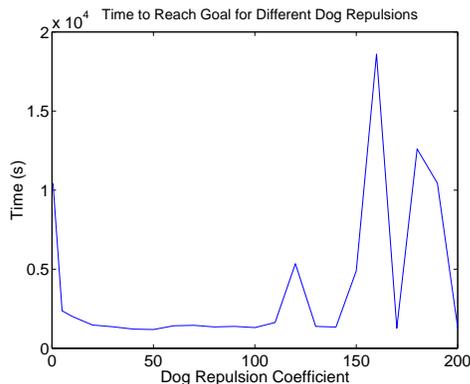


Figure 10: The time to reach the goal with different dog repulsion coefficients.

Conclusion

In this paper, we have shown that complex group behaviors can be generated if some global information of

the environment is available. The global knowledge used is a roadmap of the environment. The information it contains, such as topological information and adaptive edge weights, enables the flock to achieve behaviors that cannot be modeled with local information alone. Moreover, since in many cases global knowledge involves high communication costs between individuals, indirect communication through dynamic updates of the roadmap's edge weights provides a less expensive means of obtaining global information.

Our simulation results for the three types of behaviors studied show that the performance of the roadmap-based behavior is very close to an ideal behavior that has complete knowledge. Our future work will focus on shepherding and searching for moving goals, as in pursuit/evasion games.

References

- Amato, N. M.; Bayazit, O. B.; Dale, L. K.; Jones, C. V.; and Vallejo, D. 1998. OBPRM: An obstacle-based PRM for 3D workspaces. In *Proc. Int. Workshop on Algorithmic Foundations of Robotics (WAFR)*, 155–168.
- Balch, T., and Hybinette, M. 2000. Social potentials for scalable multirobot formations.
- Brogan, D. C., and Hodgins, J. K. 1997. Group behaviors for systems with significant dynamics. In *Autonomous Robots*, 137–153.
- Fukuda, T.; Mizoguchi, H.; Sekiyama, K.; and Arai, F. 1999. Group behavior control for MARS (micro autonomous robotic system). In *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, 1550–1555.
- Funge, J.; Yu, X.; and Terzopoulos, D. 1999. Cognitive modeling: Knowledge, reasoning and planning for interlligent characters. In *Computer Graphics*, 29–38.
- Kavraki, L.; Svestka, P.; Latombe, J. C.; and Overmars, M. 1996. Probabilistic roadmaps for path planning in high-dimensional configuration spaces. *IEEE Trans. Robot. Automat.* 12(4):566–580.
- Khatib, O. 1986. Real-time obstacle avoidance for manipulators and mobile robots. *Int. J. Robot. Res.* 5(1):90–98.
- Latombe, J. C. 1991. *Robot Motion Planning*. Boston, MA: Kluwer Academic Publishers.
- Li, T. Y.; Jeng, Y. J.; and Chang, S. I. 2001. Simulating virtual human crowds with a leader-follower model. In *Proceedings of 2001 Computer Animation Conference*.
- M., D.; Caro, G. D.; and Gambardella, L. M. 1999. Ant algorithms for discrete optimization. In *Artificial Life*, 137–172.
- Mataric, M. J. 1994. *Interaction and Intelligent Behavior*. Ph.D. Dissertation, MIT EECS.
- Nishimura, S., and Ikegami, T. 1997. Emergence of collective strategies in prey-predator game model. *Artif. Life* 3:243–260.
- Parker, L. E. 1993. Designing control laws for cooperative agent teams. In *IEEE International Conference on Robotics and Automation*, 582–587.
- Reynolds, C. W. 1987. Flocks, herds, and schools: A distributed behavioral model. In *Computer Graphics*, 25–34.
- Reynolds, C. W. 1999. Steering behaviors for autonomous characters. In *Game Developers Conference*.
- Russell, S., and Norvig, P. 1994. *Artificial Intelligence: A Modern Approach*. Prentice Hall, 1th edition.
- Saiwaki, N.; Komatsu, T.; Yoshida, T.; and Nishida, S. 1997.

- Automatic generation of moving crowd using chaos model. In *IEEE Int. Conference on System, Man and Cybernetics*, 3715–3721.
- Sun, S.-J., and Sim, D.-W. L. K.-B. 2001. Artificial immune-based swarm behaviors of distributed autonomous robotic systems. In *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, 3993–3998.
- Vaughan, R. T.; Sumpter, N.; Henderson, J.; Frost, A.; and Cameron, S. 2000. Experiments in automatic flock control. *J. Robot. and Autonom. Sys.* 31:109–117.
- Ward, C.; Gobet, F.; and Kendall, G. 2001. Evolving collective behavior in an artificial ecology. *Artif. Life* 7:191–209.
- Yu, X., and Terzopoulos, D. 1994. Artificial fishes: Physics, locomotion, perception, behavior. In *Computer Graphics*, 24–29.