Shepherding Behaviors*

Jyh-Ming Lien† neilien@cs.tamu.edu
O. Burchan Bayazit§ bayazit@cse.wustl.edu

Ross T. Sowell‡ sowelrt0@sewanee.edu
Samuel Rodríguez† sor8786@cs.tamu.edu
Nancy M. Amato† amato@cs.tamu.edu

Technical Report TR03-006
PARASOL LAB
Department of Computer Science
Texas A&M University

November 12, 2003

Abstract

Shepherding behaviors are a type of flocking behavior in which outside agents guide or control members of a flock. Shepherding behaviors can be found in various forms in nature. For example, herding, covering, patrolling and collecting are common types of shepherding behaviors. In this work, we investigate ways to simulate these types of behaviors.

*This research supported in part by NSF Grants ACI-9872126, EIA-9975018, EIA-0103742, EIA-9805823, ACR-0081510, ACR-0113971, CCR-0113974, EIA-9810937, EIA-0079874.

†Parasol Lab., CS Dept., Texas A&M University.
‡CS Dept., University of the South, Sewanee, TN.
§CSE Dept., Washington University in St. Louis.

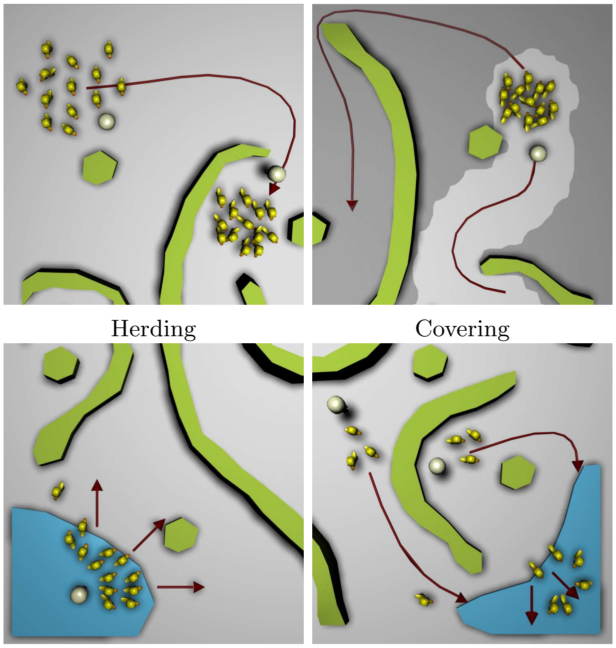

Figure 1: Shepherding behaviors.

1 Introduction

Simulating the coordinated behavior of a flock has attracted the attention of researchers in fields such as robotics, games and computer animation. Shepherding behaviors, specifically, are one class of flocking behaviors in which *one or more* external agents (called *shepherds*) attempt to control the motion of another group of agents (called a *flock*) by exerting repulsive forces from shepherds to the flock. An example found in agriculture is a sheep dog guiding a flock of sheep. Indeed, shepherding behaviors include a variety of natural behaviors and we believe that methods for simulating them will have a variety of applications.

Herding, *covering*, *patrolling* and *collecting* are common types of shepherding behaviors (see Figure 1). In the herding behavior, shepherds steer a flock from a start region to a goal region. For the covering behavior, shepherds guide a flock to visit the specified positions in the environment. In the patrolling behavior, shepherds protect a designated region and keep the flock from entering it. In the collecting behavior, shepherds gather scattered flock members into a designated region.

It is interesting to note that these various shepherding behaviors can be potentially applied in different fields of research. For example, automated methods for herding could be applied to construct robots that can function as cheaper alternatives to sheep dogs [15, 17], or and they could be used to study the neuron migration process in which a repulsive molecule pushes young cells to their permanent positions in the brain [18]. Covering methods would provide ways to accomplish tasks like mine sweeping and surveillance. Patrolling behaviors could be used to build robots to prevent birds from being sucked into airplane engines at airports [1] or to keep swimmers or children away from dangerous zones on a beach or in a school. Collecting behaviors may be used to study how predators hunt or to construct robots that gather and skim off spilt oil from oil tankers to avoid further damage to our ecosystem [8] (in this case, the skimmers are the shepherds and the floating oil is the flock).

In this paper, we study methods based on our previous work [6] for simulating simple shepherding behaviors using a *single* shepherd. In [6], a shepherd uses roadmaps to steer the flock and to re-group separated flock members. Although, the shepherd regroups flocks, [6] does not consider how movements of the shepherd will affect the flock. For example, when the shepherd approaches the flock, the shepherd does not attempt to avoid disturbing or separating the flock. Here, we focus on improving the shepherd's movements to gain better control of the flock's motion and use this improved control to demonstrate a wider variety of shepherding behaviors. The key contributions of this work are to extend our previous shepherding behaviors by:

- improving the shepherd's locomotion, and
- supporting a variety of behaviors such as herding, covering, patrolling and collecting.

2 Related Work

Reynolds' influential flocking simulation [14] established the feasibility of modeling such a system. His work showed that flocking is a dramatic example of *emergent behavior* – global behavior arising from the interaction of simple local rules. Each individual member of the flock 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 [16, 7, 9, 4] and robotics [12, 15, 13, 17, 10].

Shepherding is an interesting flocking behavior that has received comparatively little attention. Schultz et al. [15] applied a genetic algorithm to learn rules for a shepherd robot to control the movement of another robot (sheep). The sheep reacts to the shepherd by moving away from it. Vaughan et al. [17] simulate and construct a robot that shepherds a flock of geese in a simple (circular) environment. Funge et al. [9]

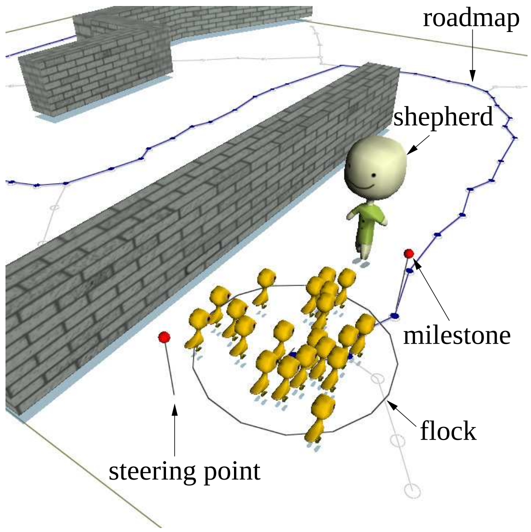

Figure 2: An environment and associated terms.

have simulated an interesting shepherding behavior in which a T-Rex chases raptors out of its territory. More recently, Bayazit et al. [6] use a roadmap to simulate shepherding in environments with obstacles.

3 Preliminaries

In this section, we define terms and concepts used in this paper. A *shepherd* is an external agent that influences the movement of the flock. A *flock* is a collection of agents that tries to keep away from the shepherd. The shepherd's task is to steer the flock to desired locations. In addition to steering, the shepherd unites separate flock *groups*. In a group, each member can *see* at least one member in that group. Usually, flock separation is caused by repulsive forces exerted from obstacles or shepherds. The *flock contour* is the smallest polygon that encloses all flock members.

A *milestone* is any position toward which the shepherd attempts to *steer* the flock, and a *steering point* is any position toward which the shepherd moves himself in order to influence the movement of the flock; see Figure 2. As in [6], a milestone is a node of a global *dynamic roadmap* close to the flock and a steering point is a point on the opposite side of the flock from the milestone. A roadmap is an abstract representation of the feasible space in a given environment. A dynamic roadmap is a roadmap storing information that changes dynamically during simulation.

Shepherd's Locomotion. We define a shepherd's locomotion as the manner in which the shepherd will move in order to control the movement of a flock. The shepherds locomotion remains invariant in different shepherding behaviors and dramatically affects the quality of simulation. We divide the shepherd's locomotion into two sub-problems: *approaching* and *steering*. In the approaching problem, we study how the shepherd goes to the steering point near the flock from its current position; see Figure 3(a). In the steering problem, we study how the shepherd steers the flock toward the milestone; see Figure 6(a).

It is important to note that we use shepherding as a broad term to describe any flocking behavior in which outside agents influence the movement of a flock, i.e., our definition of shepherding behaviors is not limited to herding behaviors as those in [15, 17, 6]. In Section 4, we propose various approaching and steering locomotions. Then, these locomotions are used as a common foundation for all the shepherding behavior simulations described in Section 5.

4 Shepherd Locomotion

All shepherding behaviors share a common set of locomotions. Here, we explore ways to have the shepherd more intelligently position himself to gain better control of its flock, e.g., reduce number or degree of flock separations. We study methods of *approaching* the flock and methods of *steering* the flock.

4.1 Approaching the Flock

In the approaching problem, we study how a shepherd should get to a steering point from its current position (Figure 3(a)). The difficulty comes from the fact that *the contour of a flock deforms dynamically due to shepherd's approaching.* In order to gain better control of a flock, the shepherd should disturb the flock as little as possible. In the following, we discuss three approaching methods that use a straight-line, a safe-zone and a dynamic roadmap, respectively.

Using a Straight Line. The simplest solution to the approaching problem is to have the shepherd move in a straight line from its current position to a steering point. This method is used in [15, 17, 6].

The problem with this method is that the straight line between shepherd's initial position and steering point often intersects the natural contour of the flock. As the shepherd traverses this line, it disturbs the flock and causes the flock to separate into two or more groups. This is illustrated in Figure 3(b).

Using a Safe Zone. A *safe zone* is a region around the flock outside of which the shepherd can move freely without disturbing the contour of the flock; see Figure 4. By not penetrating this safe zone when ap-

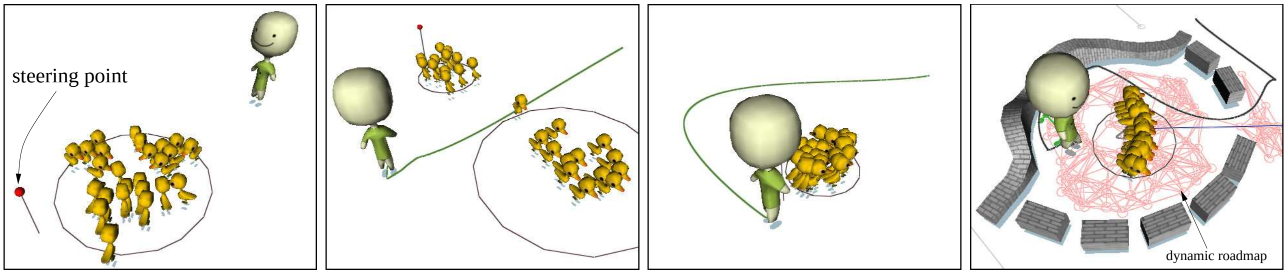

(a) (b) (c) (d)

Figure 3: (a) The approaching problem. (b) Approaching using a straight line. Shepherd's trajectories are shown in dark lines. The flock separates after the shepherd's approaching. (c) Approaching using a circular safe zone. (d) Approaching using a dynamic roadmap.

proaching the flock, the shepherd can effectively approach the flock without causing the flock to separate.

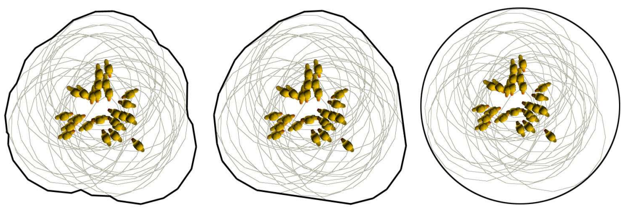

(a) (b) (c)

Figure 4: Safe zones. Small circles in the side safe zones are viewing range of the flock members.

Finding an exact safe zone (Figure 4(a)) requires expensive computation in which we need to compute the contour of the flock and then convolve the contour with the flock's viewing circle. We solve this problem by approximating the exact safe zone using a convex hull or an enclosing circle of the flock, as shown in Figure 4(b) and 4(c), respectively. Figure 3(c) shows an approaching trajectory for the shepherd based on a circular safe zone that does not disturb the flock.

Approaching using a safe zone works best in a wide open area. In an area cluttered with obstacles, the shepherd may have to access the steering point by penetrating the safe zone; see Figure 5(a).

Using a Dynamic Roadmap. A dynamic roadmap used for approaching is part of the global roadmap but constructed during the simulation in order to reflect the dynamic states of the flock. Dynamic roadmap nodes are created in the vicinity of the flock when the shepherd is approaching the flock and are removed when the flock moves away from them. Thus, a dynamic roadmap provides more detailed information in regions of current interest.

Each node v of a dynamic roadmap represents a small circular region with radius ϵ and it stores a set of flock members N_f that can see v. If $N_f = \emptyset$, v will be removed from the roadmap. Intuitively, the size of N_f provides the *safeness* of the region of v and, if the shepherd enters v, only N_f flock members will be disturbed. These dynamic nodes are sampled inside the circular safe zone of the flock using distribution $P = 1 - |P_{gauss}|$, where $|P_{gauss}|$ is the absolute value of a Gaussian distribution centered at 0. Thus, more nodes will be sampled from the boundary of the safe zone since it is better for the shepherd to travel through them than on those close to the center, see Figure 5(b) and (c). Figure 3(d) shows a trajectory of the shepherd approaching the flock using a dynamic roadmap.

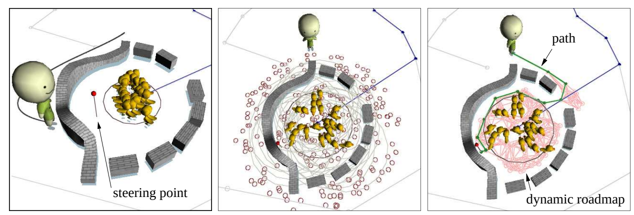

(a) (b) (c)

Figure 5: (a) The only way to access the steering point is to penetrate the safe zone. The shepherd approaching the flock using a circular safe zone fails to reach the steering point. (b) Small circles are roadmap nodes sampled in the circular safe zone. (c) A snapshot of a dynamic roadmap built from (b) and a path extracted from this map.

4.2 Steering the Flock

Once the shepherd has approached the flock, the shepherd must now determine how to steer the flock to the (next) milestone (Figure 6(a)). When the (average) heading direction of the flock is pointing to the milestone, the shepherd needs to push the flock forward. This is called *forward steering*. Otherwise, the shepherd needs to turn the heading direction of the flock toward the milestone. This is called *turn steering*. We discuss two forward-steering methods, one where the shepherd moves *straight behind* the flock and one where it moves from *side-to-side behind*, and turn-steering methods.

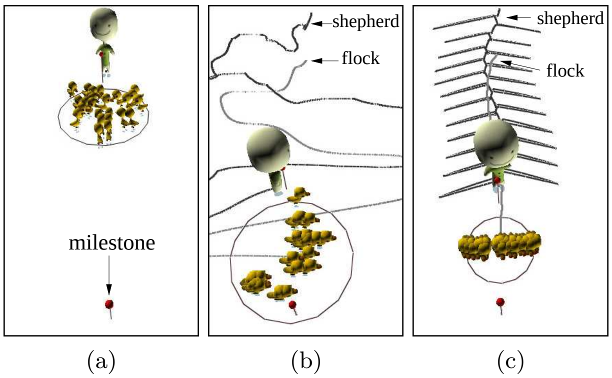

Figure 6: (a) The steering problem. (b) An example of trajectories of the flock and the shepherd when the shepherd steers straight behind the flock. Dark (light) lines are shepherd's (flock's) trajectory. (c) Trajectories from steering side-to-side behind the flock. Comparing to the trajectories in (b), the flock smoothly reaches the milestone.

Straight Behind the Flock. Assume that the flock is facing a milestone. By placing the shepherd in a steering point, the flock has been nudged towards the milestone. The simplest way for the shepherd to steer the flock farther is to compute a new steering point and then move in a straight line to that point. This process is repeated until the flock reaches the milestone.

This is a simple and effective strategy and has been used in [15, 17, 6]. However, as the shepherd moves straight behind the flock, the flock members closer to the shepherd tend to be pushed more strongly toward the milestone, while the flock members on the sides tend to be pushed further to the side, thus spreading out the flock (see Figure 8(a)). As a result, the flock becomes hard to control and some flock members may even become separated. Figure 6(b) shows steering trajectories and the difficulty of controlling the flock using this strategy.

Side-to-Side Behind the Flock. In an effort to avoid the problem encountered by moving straight behind the flock, we propose that the shepherd moves repeatedly, from one side of the flock to the other, as it advances behind the flock. Instead of simply always moving to a new steering point, the shepherd first moves to the right of this point, then to the left, and then back to the right. This side-to-side motion can be seen in Figure 6(c). In fact, this strategy is used by some dogs, such as Border Collies, as seen in Figure 7.

By following this side-to-side motion, the shepherd not only steers the flock towards the milestone but also squeezes the flock to keep them from separating by pushing the flock members on the sides towards the center, thus reducing the chance of a flock separation; see Figure 8(b).

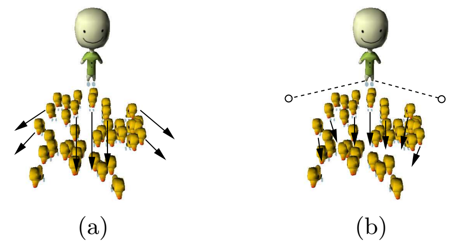

Figure 8: Reactions of the flock to steering (a) straight behind the flock and (b) side-to-side behind the flock.

Turning the Flock. Turning the flock is difficult due to the flock's momentum. Let $\vec{v}_f$ be the average heading direction of a flock F and let $\vec{v}_d$ be the desired heading direction of F. The shepherd needs to turn the flock when the angle between $\vec{v}_f$ and $\vec{v}_d$ becomes intolerable. We consider two types of turn steering: *stop-turn steering* and *pre-turn steering*.

Intuitively, in stop-turn steering, the shepherd stops the flock by moving itself in front of the flock until the flock stops and then steers it in the direction of $\vec{v}_d$ (Figure 9(b)). Thus, the stop-turn steering is used to *correct* the flock's movement. In pre-turn steering, the shepherd initiates the turning process before the turn takes place by placing itself in the direction of $\vec{v}_f - \vec{v}_d$ near the boundary of the flock (Figure 9(c)). Therefore, pre-turn steering is used to *prevent* the flock from heading to a wrong direction.

The trajectory of the flock in Figure 9(a) shows that the shepherd finally turns the flock after the flock bounces up and down for several times. The shepherd in Figure 9(b) stops the flock after the flock passes the turning point, thus reducing the traveling distance of the flock. The shepherd in Figure 9(c) prevents the flock from over shooting and therefore smoothly turns the flock.

5 Shepherding Behaviors

Shepherding behaviors can be found in various forms in our physical world and, interestingly, shepherds in these behaviors share similar locomotions. Here we show that we can simulate four types of shepherding

Figure 7: A Border Collie steers geese by moving itself side-to-side behind the flock. Images are captured from [5].

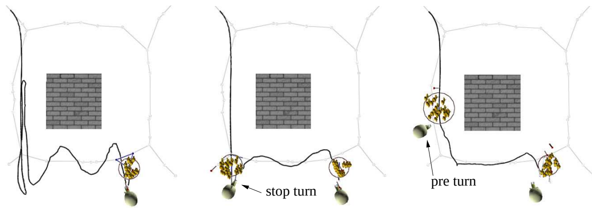

(a) (b) (c)

Figure 9: (a) Without turn steering. (b) With stop-turn steering. (c) With pre-turn and stop-turn steerings.

behaviors, i.e., herding, covering, patrolling and collection, using the previously described shepherd locomotions. We also show that our new locomotions improve the shepherd's control of the flock as compared to the methods used in [15, 17, 6]. We study these four shepherding behaviors under various conditions, e.g., open/cluttered environments or large/small flocks.

Our simulation system is implemented by *Python* scripts that use a home made simulation engine coded in C^{++} which is exposed to Python using *boost/python* library. Due to the difficulty of presenting these simulations in a paper, we encourage readers to view animations at our website [2].

5.1 Herding

Herding is a behavior in which a shepherd needs to move all flock members from a start region to a goal region. Herding is the most simple and common shepherding behavior and is studied in [15, 17, 6]. However, to the best of our knowledge, no study has been done to enhance the shepherd's ability to control the flock, even for the herding behavior. In the following, we experiment with different combinations of the shepherd's locomotions and show that our methods do improve the shepherd's control.

5.1.1 Experimental Results

Figure 10 shows two environments, one without and one with obstacles. In both cases, the shepherd needs to steer 20 flock members from their current positions until all flock members can see the goal position. Also, both environments have large open spaces. Environment herding_env1 is similar to those used in [15, 17]. In herding_env2, the flock needs to make two U-turns before reaching the goal.

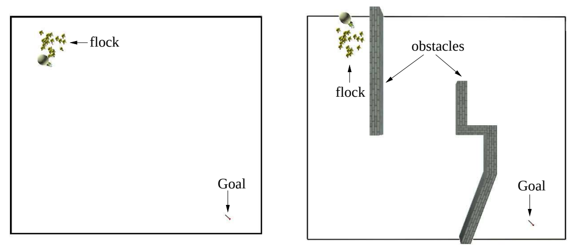

(herding_env1) (herding_env2)

Figure 10: Herding environments.

In Figure 11(a) and 11(b), we compare the total time to reach the goal and the separation time during the herding process using different shepherd locomotions. From the results, we discover that the *side-to-side steering* improves the efficiency of herding most, i.e., SS, SSS, SSSP, DSS and DSSP in Figure 11(a) and (b). We also find that *turn steering* can reduce separation time even more. In the obstacle-free environment, the total and separation time both decrease when safe-zone approaching is used (see LL and SL in Figure 11(a)). The shepherd with SS in Figure 11(a) herds the flock in the shortest time because turning is not critical in this environment and turning takes time. When obstacles are present, as seen in Figure 11(b), the shepherd with LL and SL fails to steer the flock to the goal in the maximum allowed time. Also, it is clear that turning is important in this environment as both stop-turn and pre-turn steering reduce the total time and separation time to almost half of the time without turn steering.

The traveling distance of the flock also reflects that our shepherd locomotions are better; i.e., the flock a takes shorter path in Figure 11(c) and (d). Moreover, although the shepherd seems to make more movements using safe-zone and dynamic roadmap approaching and side-to-side and turn steerings, Figure 11(c) and (d) show that the traveling distance of the shepherd using these locomotions is similar or even less than that using straight-line approaching and steering. This means that the shepherd with our locomotions is energetically more efficient.

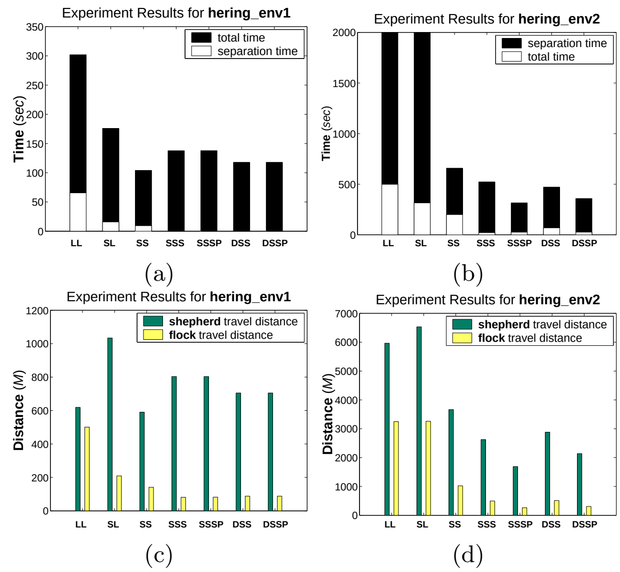

Figure 11: Experiment results. LL: Straight-line approaching and steering. SL: Safe-zone approaching and straight-line steering. SS: Safe-zone approaching and side-to-side steering. SSS: SS and stop-turn steering. SSSP: SSS and pre-turn steering. DSS: Dynamic-roadmap approaching and side-to-side and stop-turn steering. DSSP: DSS and pre-turn steering.

5.2 Covering

For the covering behavior, the shepherd guides the flock to areas of the environment that have not been visited. The shepherd finds a less visited neighboring area using a dynamic roadmap whose edge weights increase after the flock has traversed them; details are described in [6]. The covering behavior could be used, for example, for vacuuming or mowing, to guide visitors through a museum, or to guide the flock to pastures that have not been grazed yet.

5.2.1 Experimental Results

The environment that the covering behavior was tested on is shown in Figure 12(a). It consists of large open areas and cluttered areas, with the obstacles varying in size. Thus, there are many cases where the shepherd will have to steer the flock in difficult ways. This tests how well the shepherd's locomotion controls the flock.

We compare the percent of the environment covered *by the flock* and the total number of goal nodes that the shepherd was able to steer the flock toward. We study three locomotions for the covering behavior, i.e., the locomotion (LL) with straight-line approaching and steering, the locomotion (SL) with safe-zone approaching and side-to-side and both turn steerings,

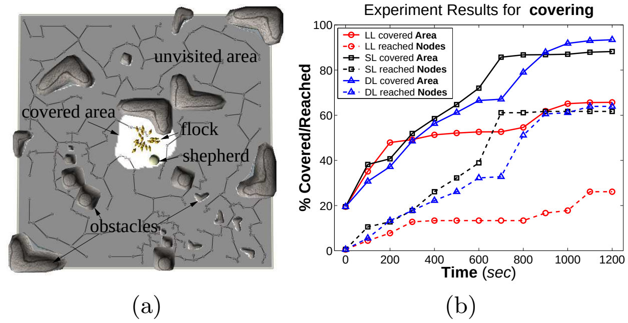

Figure 12: (a) Covering environment with 20 flock members. The roadmap is generated using both OBPRM [3] and MAPRM [11]. (b) Experiment results. LL: Straight-line locomotion. SL: Safe-zone and side-to-side locomotion. DL: Dynamic-roadmap and side-to-side locomotion.

and the locomotion (DL) with dynamic approaching and side-to-side and both turn steerings. The performance of these locomotions is shown in Figure 12(b).

It is clear from Figure 12(b) that SL and DL result in a better covering behavior. While the coverage of LL is not significantly less than SL or DL, the number of goal nodes that LL is able to reach is significantly less. This is because the shepherd with LL separates the flock and the separated flock covered more areas accidentally. The separation time using LL is roughly one-fourth of the total simulation time whereas with SL or DL the total separation time is less than one-tenth of the total simulation time. DL is able to cover more of the area and visit more goal nodes than SL although the difference is not significant.

5.3 Patrolling

The shepherd in the patrolling behavior needs to guard a designated region called the *forbidden area* or FA. Once the intrusion of the flock is found, the shepherd will chase the flock until the flock leaves the FA. If more than one intruding group is found, the shepherd will chase the larger group or the group that more deeply penetrates the FA. Patrolling behavior can be used to avoid catastrophic accidents in FA, e.g., keeping birds from interfering with aircrafts at airports [1], or preventing children from entering dangerous areas.

5.3.1 Experimental Results

In the patrolling experiments, we study how the shepherd performs with three locomotions, i.e., LL, SL and DL, using the environment shown in Figure 13(a). Performance of the shepherd is measured as the reciprocal of the average time spent in the FA for each flock member over a 1000 second period. We also examine how well these locomotions perform with different flock sizes.

Although one of the shepherd's tasks is to reunite separated flock members, grouping in the patrolling behavior does not seem to be critical. On the other hand, if the flock is kept united, then the shepherd can focus on a smaller area than when the flock is scattered. Therefore, we also compare the shepherd with or without grouping the separate members.

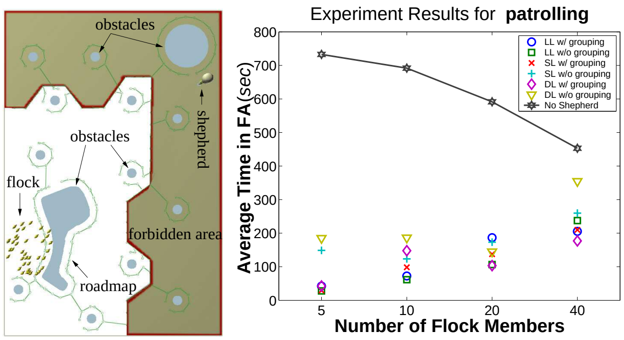

(a) (b)

Figure 13: (a) The covering environment. The dark shaded region enclosed by bold lines is the forbidden area. The roadmap is generated using OBPRM [3]. (b) Total experimental time is 1000 seconds. The x axis shows total number of flock in the environment.

Figure 13(b) shows our experimental results. The plotted line above all points indicates the time in FA when no shepherd guards the FA. We observe that, when the size of the flock is small, i.e., 5 and 10, the shepherd with SL *without grouping* performs best. When the size of the flock is large, i.e., 20 and 40, the shepherd with DL *with grouping* performs best. In fact, for SL and DL, grouping always helps the shepherd to reduce the time in the FA and grouping becomes more important for the shepherd to handle a large flock. Note that the performance of the shepherd degrades when the size of the flock increases no matter what locomotion is used.

5.4 Collecting

A shepherd in the collecting behavior gathers initially scattered flock members into a designated region, called the *home area* (HA). We assume that once a flock member enters the HA that it stays there. One of the potential applications of this study is for robots that clean oil spills, which is usually done manually using skimmers and barriers [8]. This behavior can also be found in nature, e.g., lions hunt by pushing their prey into a ambush, a dog gathers a herd of horses or cattle, etc.

5.4.1 Experimental Results

Our experiments are done in the environment shown in Figure 14 with 50 flock members. As in Section 5.2 and 5.3, we study shepherd performance using three locomotions, i.e., LL, SL and DL. The performance is measured as the time to collect all flock members. In addition, we experiment with three types of flocks: *cattle*, *ducks* and *sheep*. They differ in the way they respond to the shepherd's approach. Sheep have a tendency to stay together when the shepherd approaches and cows and ducks separate more easily [5]. Also, cows are less intimidated by the shepherd and, thus, are more difficult to steer. We examine how well the shepherd controls these types of flocks.

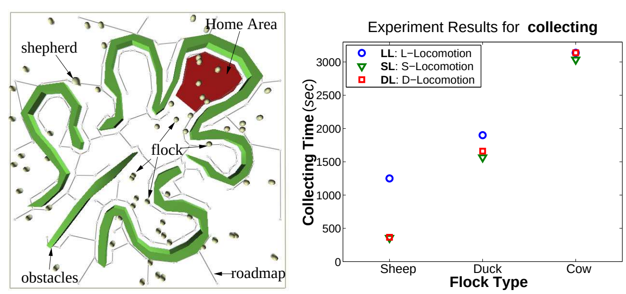

(a) (b)

Figure 14: Collecting environment with 50 flock members. The roadmap is generated using both OBPRM [3] and MAPRM [11].

Figure 14(b) shows our experimental results. The flock type used in the previous sections is sheep, so the shepherd with SL and DL performs consistently much better than the shepherd with LL for this collecting behavior. Although this remains true for duck and cow flocks, the difference between SL, DL and LL becomes smaller. Because cow and duck flocks tend to separate, it is harder to keep the flock compact and the performance degrades.

6 Discussion and Conclusion

In this paper, we have shown that a shepherd can control a flock more efficiently by using more intelligent locomotion techniques, such as the safe-zone and dynamic-roadmap approachings and side-to-side and turn steerings. In Section 5.1 and Section 5.2, flock separation and traveling time is reduced when the shepherd uses these new locomotions in both open and cluttered environments. We have also shown that our locomotions enhance the shepherd's ability to handle larger groups of the flock (Section 5.3) and to handle different types of flocks (Section 5.4).

One shepherd may not always be sufficient to accomplish a shepherding task. As we've seen in the patrolling and collecting experiments, when the size of the flock increases or the flock's fear of the shepherd and coherence decreases, the flock becomes hard to control. Multiple shepherds are necessary to handle these problems. In fact, two or three sheep dogs usually work together in fields [5]. We are currently exploring how multiple shepherds might work together to better control the movement of a flock.

References

[1] Endangered Wildlife Trust. EWT airport safety project. 2002. http://www.ewt.org.za/.

[2] http://parasol.tamu.edu/groups/amatogroup/research/shepherding/.

[3] N. M. Amato, O. B. Bayazit, L. K. Dale, C. V. Jones, and D. Vallejo. OBPRM: An obstacle-based PRM for 3D workspaces. In *Robotics: The Algorithmic Perspective*, pages 155–168, Natick, MA, 1998. A.K. Peters. Proceedings of the Third Workshop on the Algorithmic Foundations of Robotics (WAFR), Houston, TX, 1998.

[4] M. Anderson, E. McDaniel, and S. Chenney. Constrained animation of flocks. In *Proceedings of the 2003 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, pages 286–297. Eurographics Association, 2003.

[5] S. Barwig. Herding 1. Canine Training Systems Ltd.

[6] O. B. Bayazit, J.-M. Lien, and N. Amato. Better group behaviors using rule-based roadmaps. In *Proceedings of the Workshop on Algorithmic Foundations of Robotics (WAFR'02)*, 2002.

[7] D. C. Brogan and J. K. Hodgins. Group behaviors for systems with significant dynamics. In *Autonomous Robots*, pages 137–153, 1997.

[8] M. F. Fingas. *The basics of oil spill cleanup.* Lewis Publishers, 2nd edition, 2001.

[9] J. Funge, X. Tu, and D. Terzopoulos. Cognitive modeling: Knowledge, reasoning and planning for interlligent characters. In *Computer Graphics*, pages 29–38, 1999.

[10] T.-Y. Li and H.-C. Chou. Motion planning for a crowd of robots. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2003.

[11] J.-M. Lien, S. L. Thomas, and N. M. Amato. A general framework for sampling on the medial axis of the free space. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2003.

[12] M. J. Mataric. *Interaction and Intelligent Behavior.* PhD thesis, MIT EECS, 1994.

[13] S. Nishimura and T. Ikegami. Emergence of collective strategies in prey-predator game model. *Artif. Life*, 3:243–260, 1997.

[14] C. W. Reynolds. Flocks, herds, and schools: A distributed behaviroal model. In *Computer Graphics*, pages 25–34, 1987.

[15] A. C. Schultz, J. J. Grefenstette, and W. Adams. Roboshepherd: Learning complex robotic behaviors. In *Proceedings of the International Symposium on Robotics and Automation*, pages 763–768, May 1996.

[16] X. Tu and D. Terzopoulos. Artificial fishes: Physics, locomotion, perception, behavior. In *Computer Graphics*, pages 24–29, 1994.

[17] R. T. Vaughan, N. Sumpter, J. Henderson, A. Frost, and S. Cameron. Experiments in automatic flock control. *J. Robot. and Autonom. Sys.*, 31:109–117, 2000.

[18] W. Wu, K. Wong, J.-H. Chen, Z.-H. Jiang, S. Dupuis, J. Wu, and Y. Rao. Directional guidance of neuronal migration in the olfactory system by the concentration gradient of the secreted protein slit. *Nature*, 400:331–336, 1999.