

# From Robots to Animals: Virtual Fences for Controlling Cattle

Zack Butler\* Peter Corke† Ron Peterson‡ Daniela Rus§

## Abstract

We consider the problem of monitoring and controlling the position of herd animals, and view animals as agents with natural mobility but not strictly controllable. By exploiting knowledge of individual and herd behaviour we would like to apply a vast body of theory in robotics and motion planning to achieving the constrained motion of a herd.

In this paper we describe the concept of a virtual fence which applies a stimulus to the animal as a function of its pose with respect to the fence line. Multiple fence lines can define a region, and the fences can be static or dynamic. The fence algorithm is implemented by a small position-aware computer device worn by the animal, which we refer to as a Smart Collar. We describe a herd-animal simulator, the Smart Collar hardware and algorithms for tracking and controlling animals as well as the results of on-farm experiments with up to 8 Smart Collars.

## 1 Introduction

Management of domestic animals is a very old human activity which has seen relatively little change over the centuries. Cattle and other livestock forage over large paddocks whose perimeters are usually bounded by conventional fences. The two essential activities are confinement by fences and mustering or herding. A typical ranch will normally have several paddocks separated by conventional fences. Animals are rotated among paddocks in an attempt to prevent overgrazing, and for husbandry practices such as during calving or weaning. Fence technology has changed from stone and hedge to electrified wires, but is still labour intensive to install and maintain and is not easily reconfigured. Mustering has changed from shepherds on foot perhaps aided by dogs, through the use of horses to motorized vehicles and even helicopters. Nevertheless herding remains a labor intensive activity when cattle and sheep graze over large paddocks and animals are rotated frequently to achieve pasture management constraints. It is physically hard work, often carried out in extreme weather conditions and remote locations. The costs of fencing and mustering remain a significant fraction of the cost of raising an animal, and has not benefited from the technical revolution in automation, computing and communication.

Within the robotics community there is much interest in human-robot interaction, but little on animal-robot interaction and this is surprising given the large number of domestic animals that we rely on, and their importance to our civilization.

Herd of animals such as cattle are complex systems. Herd societal behaviors are well known to stock farmers but their study by animal ethologists [27] has just begun. There are interesting interactions between individuals, such as friendship, kinship, group formation, leading and following. There are complex interactions with the environment, such as looking for a water source in a new paddock by perimeter tracing along the fence and random walking within the perimeter. Cattle have demonstrated spatial memory that influences habitat use [14] and also memory capable of recognizing peers and humans [30]. Our goal is to develop computational approaches for studying groups of agents with natural mobility and social interactions. Such systems differ in many ways from engineered mobile systems because their agents can move on

---

\*Computer Science Department, Rochester Institute of Technology, Rochester, NY 14623 USA, (e-mail: [zjb@cs.rit.edu](mailto:zjb@cs.rit.edu)).

†CSIRO ICT Centre, Australia, (e-mail: [peter.corke@csiro.au](mailto:peter.corke@csiro.au)).

‡Dartmouth Computer Science Department, Hanover, NH 03755 USA, (e-mail: [rapjr@cs.dartmouth.edu](mailto:rapjr@cs.dartmouth.edu)).

§Computer Science and Artificial Intelligence Laboratory, MIT, Cambridge MA 02139, USA, (e-mail: [rus@csail.mit.edu](mailto:rus@csail.mit.edu)).

their own due to complex natural behaviors as well as under the control of the environment (for example moving toward a food or water source). We wish to generate models of such systems using observed physical data and to use these models to synthesize controllers for the movement of the mobile agents. Unlike more familiar robot control problems, the animal state (stress, hunger, desire) is only partially observable and only limited control over motion can be exerted.

There are two fundamentally different approaches to controlling animal position: a physical agent such as a sheepdog or robot, and a stimulation device worn by the animal. The first approach works by applying a moderate threat which the animal will flee from, however knowledge of the herd dynamics allows a very small number of threat agents to control a very large number of animals. In this category there is the pioneering work of Vaughan [29] who demonstrated a mobile robot that was able to herd a flock of ducks to a desired location within a circular pen.

The second category relies on an active device worn by the animal that can provide an aversive stimulus as required, typically this is an odorous spray, a sound, or an electric shock. Devices to control domestic pets such as dogs are commercially available, either to prevent it from barking or to constrain it. The latter device typically works by sensing proximity to a buried perimeter wire which allows for a simple collar on the dog. Clearly such an approach, based on installed infrastructure, would be prohibitively expensive for large scale agriculture and has all the disadvantages of a conventional fence (installation, maintenance and not-configurable). However with the advent of GPS technology, and the rapid reduction in receiver cost, it is possible to implement such a fence without installed infrastructure. That is, a boundary can be defined in terms of geographic coordinates and a stimulus applied when the animal moves beyond the boundary. This functionality replaces that of a fence and is referred to as a “virtual fence”. By making the boundary dynamic we achieve the function of herding or mustering, almost for free.

In this paper we describe our work which brings together concepts from the animal behaviour and management communities, adhoc networking and robotic motion planning. In [9] we described our first experiments in controlling a herd of cows with a single static virtual fence using an approach that relies on ad-hoc networking. In [8] we extended that work on the algorithmic side, by introducing motion planning for computing dynamic virtual fences whose goal is to muster the herd to a new location. We have implemented these algorithms in simulation and deployed 10 smart collars on cows at Cobb Hill Farm in Vermont.

Our work builds on important recent studies in studying and modeling animal behavior. The Leurre project [1, 10, 11, 13, 15, 17, 19, 21] has contributed seminal theoretical modeling methodologies for individuals in animal societies, based on differential equations whose parameters are computed from exponential distribution of times in animal states from real data. This approach seems to model well ants, cockroaches, and fish. However, modeling the interaction between individuals and modeling large animals remains challenging. Vaughan [29] demonstrated a mobile robot that was able to herd a flock of ducks, and Denebourg et al [11] were able to influence the behaviour of cockroaches by the motion of a small robot. Within the networking community there is growing interest in using GPS technology to track animals with fine resolution. The Zebanet project [16] considers the tracking and data collection issues in tracking herds of zebras in their natural habitat. Related work on modeling and coordination in animal and artificial robot societies also includes [4–7, 20, 22, 24].

This paper is organized as follows. Section 2 describes the prior art regarding virtual fences, our algorithmic implementation, a detailed simulation model, and the extension to dynamic path planning. In Section 3 we describe the experimental approach in some detail including the hardware and software tools we developed to facilitate the experiments. Section 4 describes experimental results that cover: collecting data to create a grazing model for the cows; collecting connectivity data and information propagation data for the multi-hop routing method within the herd; collecting stimulus-response data for individual animals; and collecting response data for the virtual fence on a group of animals.

## 2 Virtual Fencing

The application of smart collars to control cattle is first discussed in detail by Tiedemann and Quigley [23, 28] who were concerned with controlling cattle grazing in fragile riparian environments. Their first work [23], published in 1990, describes experiments in which cattle could be kept out of a region by remote manually applied audible and electrical stimulation. They note that cattle soon learn the association and keep out




Figure 1: Notation used for the virtual fence algorithms.

of the area, though sometimes cattle may go the wrong way. Cattle learn to associate the audible stimulus with the electrical one and they speculate that the acoustic one may be sufficient after training. More comprehensive field testing in 1992 is described in [28] and also identified issues about the need to train animals to associate stimulus with spatial restrictions. This issue is also discussed by Anderson [2].

The idea of using GPS to automate the generation of stimuli was proposed by Marsh [18]. GPS technology is widely used for monitoring position of wildlife. Anderson [2, 3, 12] builds on the work of Marsh to include bilateral stimulation, different audible stimuli for each ear so that the animal can be better controlled. The actual stimulus applied consists of audible tones followed by electric shocks.

The choice of stimulus is an important component of working with animals. The animal behavior literature discusses both sound and electric shock, as well as their combination. If a sound stimulus always precedes an electric shock then the animal will become trained over time and react to the sound rather than waiting for the shock. If sound is used this raises the question about the optimum sound to use. Should the sounds be natural or unnatural, can they encode information about the direction to turn, and so on.

In our experiments we used sound stimulus only for several reasons. First was an admitted squeamishness on our part in shocking cattle — the cows we worked with are a small herd of dairy cattle owned by a cooperative and are very much like pets, with individual names such as *Linden* and *Hazel*. Second, the experimental animal protocol approved by Dartmouth College allowed only the use of gentle stimuli such as sounds (protocol assurance A3259-01). Finally, we wished to explore and characterize the scope of a minimalist and friendly approach to controlling cattle that relied on sound only.

## 2.1 The Virtual Fence Algorithm

In our algorithms, a fence is represented by a point  $F_p$  and a normal vector  $F_n$ . The notation used is shown in Figure 1. The point  $F_p$  is given in (latitude, longitude) coordinates, although higher-level interfaces such as the GUIs described in Sec. 3.3.3 use relative distances internally and convert to the absolute coordinates when sending the fences to the collars.  $F_p$  can be simply any point along the fence, and  $F_n$  points perpendicularly into the interior half-plane.

This representation allows a simple calculation to determine whether the cow is “outside” the fence. The collar first computes the cow’s position relative to the fence  $x_r$  by subtracting the  $F_p$  from the cow’s absolute position  $x_a$ . The distance  $d$  that the cow is behind (or in front of) the fence is then computed as the dot product of  $x_r$  with  $F_n$ . If  $d$  is positive, the cow is in the desired region, while if it is negative, the cow is behind the fence. To compute the exact distance correctly,  $x_r$  is first converted to meters, since the  $F_n$  is defined based on equal units in each direction. This is all done with nothing more complex than multiplication.

For most applications, several fences will be present. Since each gives an interior half-plane, setting up a number of fences with their normals pointing toward each other results in a convex polygonal paddock. To do position checking in this case, the algorithm computes a  $d$  value for each fence, and reports the fence and distance that the cow is farthest behind (the most negative  $d$ ). This is necessary to ensure the cow does not appear to be close to the interior when it is in fact far behind another fence. The algorithm, with stimulus timeout, is given in Algorithm 1.

---

**Algorithm 1** Virtual fence.  $T_{\max}$  is the maximum number of sequential stimuli that can be applied and  $T_{\text{timeout}}$  is the inhibitory period that applies after that limit is exceeded.

---

```

1:  $t_{\text{stimulus}} \leftarrow 0$ 
2:  $t_{\text{timeout}} \leftarrow 0$ 
3: loop
4:    $d_{\min} \leftarrow 0$ 
5:   for  $i = 0$  to  $N_{\text{fence}}$  do
6:      $d \leftarrow (x_a - F_p) \bullet F_n$ 
7:      $d_{\min} \leftarrow \min d_{\min}, d$ 
8:    $t_{\text{timeout}} \leftarrow t_{\text{timeout}} - 1$ 
9:   if  $d < 0 \wedge t_{\text{timeout}} < 0$  then
10:    Cow  $\leftarrow$  stimulus
11:     $t_{\text{stimulus}} \leftarrow t_{\text{stimulus}} + 1$ 
12:    if  $t_{\text{stimulus}} > T_{\max}$  then
13:       $t_{\text{timeout}} \leftarrow T_{\text{timeout}}$ 
14:    else
15:       $t_{\text{stimulus}} \leftarrow 0$ 

```

---

To present a stimulus, the simplest option is to produce a sound of a given volume when the cow goes behind the fence. This can be done by testing the sign of the cow’s fence distance. However this is quite an uninformative stimulus and the animal can only determine the desired position by moving about until the stimulus stops. Instead we encode the distance from the fence to provide graduated stimulus which gets stronger the farther the cow is behind the fence. This can be done by generating a sound with volume proportional to  $-d$ . Finally, a more complex procedure is to monitor  $d$  over time, and stop the stimulus as soon as the cow begins to move toward the desired region. This is done by keeping track of  $d_{\min}$ . If the cow’s current distance is greater, a stimulus is not produced. A further sophistication would be to reduce the stimulus when the animal turns in the right direction, providing a reward for the correct response. In the approach of Anderson [3] a directional stimulus is applied by left and right sound sources as well as electrodes on each side of the neck.

If the fence is to move, it is instantiated with a non-zero velocity  $F_v$ , in m/s. The point is then moved as a function of time along the normal

$$F_p(t) = F_p(0) + \gamma F_n F_v t$$

where  $\gamma$  is a scale factor from northing/eastings to lat/long coordinates. The fence orientation could also be a function of time, that is  $F_n(t)$ .




Figure 2: Screenshot from the simulation of the virtual fence algorithm. Twenty cows are represented by ellipses, and one fence is shown as a vertical line.

## 2.2 Modeling and Simulation

To test our ideas we developed a simulator that models the behavior of a herd of cows. We were inspired by Vaughan’s duck simulator [29], but extended the animal model to account for the differences between the species as well as their environments, and added the capability for virtual as well as physical fences, animal paths, and even wind.

### 2.2.1 Animal model

To model the behavior of a cow, we use potential fields to capture response to external stimulus along with internal state (which also affects motion). It is well known that a cow’s movement will be dependent on the location of its herd. Therefore, we model the effects of one animal’s position on another’s motion: a weak attractive herding field, and a strong local “personal space” field. We also use potential fields to model the response to non-cows in the environment, such as a robot, human, or other obstacle, which represent a threat and generate a strong repulsion. A fixed fence exerts a strong normal repulsive force over a short range, whereas the virtual fences are modeled through a different method, since they do not have remote effect.

From an internal point of view, we simplify the animals’ desires with a two-state behavior model, walking and grazing, each with an associated speed and duration distribution. Transitions between these states occur in a stochastic fashion and are used to induce motion in an approximately forward direction. We also explicitly model the stress of each animal and use this to affect the animal’s behavior. Stress is created by the virtual fence stimulus as well as the nearby presence of other fast-moving animals, isolation from the herd or threats. An animal in a low stress condition will alternate between grazing and walking as described above and exhibit very little herding instinct (as observed in the field) unless it becomes well separated. An animal that is experiencing high stress will move toward other animals, and will not resume grazing until its stress has gone down. The stress level of an animal decays exponentially over time.

In addition to inducing stress, the fence stimulus has an immediate effect on the motion of the animal. We have used two different models, each of which take inspiration from field observations. In the first model, a stimulus causes the animal to quickly turn approximately  $90^\circ$ . This behavior was also observed in [23]. In the second model, the cow walks forward for a short time when stimulated.

### 2.2.2 Simulator details

We have implemented the animal model described above within a Matlab simulation that also includes static and virtual fences. The simulator displays the position and orientation of each animal as well as a graphical




Figure 3: Screen shots from a simulation of moving virtual fences, 20 cows represented by ellipses in a  $50m \times 50m$  space, with four virtual fences forming a moving rectangle around them. (a): the beginning of a simulation. (b): after 4000 time-steps which correspond to 45 mins. See also Extension 4.

indication of the animal’s stress level. The simulator instantiates a number of animals in a moderately close grouping, giving each animal low stress and a random internal state. This also ensures that the simulator starts in a stable condition. The simulation then enters a loop in which time is advanced by small amounts and each cow’s position, velocity and stress is updated. Figure 2 shows a typical simulator display.

At each time step, the simulator determines each cow’s position relative to the rest of the herd and computes a potential-field-based force from these relative positions. If the animal is distant from all others, we add an additional large force toward the herd and increase the its stress level. Fast moving or stressed animals in close proximity add to the current animal’s stress. The simulator then checks the cow’s position relative to all fences, real and virtual. For the real fences, if the animal is in close range, an additional force normal to the fence is added to the force acting on the animal. The virtual fences induce a stimulus when the cow passes into the fence zone, and the impact of this stimulus is to provide a large force as described above. We assume that this stimulus will be strong enough to override any other behavior. Thus, for the time step when the stimulus is present, it alone is used to determine the cow’s motion. The simulations presented here are based on a bilateral stimulus, such that the cow will turn sharply away from the fence (but not based on its exact angle with respect to the fence).

After the total force acting on the cow is computed, we consider the cow’s internal state to determine its motion. At this time the simulator also determines if it is time to switch to the other state. During a state switch, the animal can make small heading changes and the simulator stochastically chooses the amount of time to spend in that state (barring external influences). If the induced force is small, this is ignored and instead a small forward force is used to generate a speed appropriate to walking or grazing.

To determine the influence of the induced force on the cow, we model the cows as non-holonomic and give them a maximum angular velocity. If the virtual force given by the potential fields is not closely aligned with the cow’s current direction, the cow will turn to align itself with the force. Otherwise, the force will affect the cow’s velocity, and in either case, the position will be updated based on the velocity and the size of the time step.

### 2.2.3 Simulation results

We have performed a variety of simulations with different numbers of animals, different parameters for the potential field calculations, and different fence configurations. Multimedia extensions 1–3 show unforced behaviour, and behaviour with slow and fast moving threat objects. In one class of experiments, presented as Fig. 3 and Extension 4, we started with a group of 20 cows within an overall  $50 \times 50$  m area and put virtual fences in a rectangle closely surrounding them. We then allowed the interior area to move without changing size, and the animals did move accordingly. We have also experimented with selective virtual fences, that is, fences that affect only certain of the animals in the herd. In simulation, this allows the cows to be separated into two groups, as seen in Fig. 4 and Extension 5. Lines drawn from a cow indicate a force vector which is the result of a virtual fence interaction.

We note here that the effectiveness of all the mustering experiments does depend on the potential field

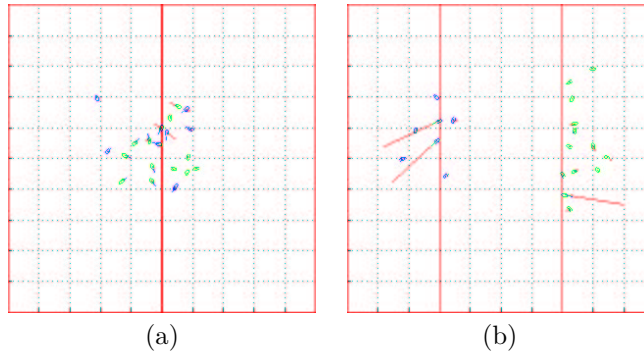


Figure 4: Simulation of selective mustering. (a): at the beginning of the simulation, the two classes of animals, represented here by dark and light ellipses, are intermingled. (b): after a short time, the two groups have been separated. See also Extension 5.

parameters and stimulus response used. We have seen good results with a variety of internal parameters provided that the stimulus is moderately effective. Our best results were obtained for the simulation of a bilateral stimulus, under which the angle between the cow and the fence is calculated and a stimulus provided that moves the cow a constant amount away from the fence. We have also tried varying the likelihood that an animal will respond to a given stimulus, and found that often 100% stimulus is not required, but that if an animal gets far beyond the initial fence location it may not be able to return to the interior.

To test the algorithms against these models, we ran virtual fences on a simulated herd with widely varying parameters. The overall goal was to move the virtual fence slowly into the herd and test how quickly the herd moved away from the encroaching fence. This was tested with different values for the grazing speed and walking speed of the cows, the level of herd-attraction and the probability that a stimulus would have the desired effect. We found that the parameters affected the overall speed of the herd in front of the fence and the number of stimuli that were applied, but in all cases the herd did move in the desired direction.

We also tested the effect of using cow orientation in controlling stimulus. Our expectation was that if the cow tends to go forward when stimulated, it would be necessary to sense the cow’s orientation and only apply stimulus when the cow is pointing in the direction we wish it to move. In the simulation, this turned out not to be necessary, since after receiving the stimulus, the cow would have increased stress and return back toward the herd even if it initially went the wrong way. However, this behavior is very dependent on the nature of the stress model.

### 2.3 Planning for Dynamic Fences

An important goal of this work is to automatically muster cows from one pasture to another. To make this possible, we have developed a path planning system for virtual fences. While this problem shares some basic features with traditional robot path planning, it has important and interesting differences as well.

The planner creates a path from one point to another using a simple occupancy grid and A\* search [25]. The object executing the plan is a virtual (polygonal) paddock, with significant extent that may change as it moves<sup>1</sup>, as long as its area remains sufficiently large for the herd. Thus, it is easier to perform planning in the workspace than the configuration space. Obstacles can overlap somewhat with the virtual fence edges, changing the effective area of the virtual paddock but not altering the plan. We use planning operators that change the dimensions of the paddock while keeping the amount of free space within the paddock sufficient for the given number of animals. Finally, we would like the motion to consist of a small number of straight-line segments. This type of optimization is necessary because changing the animals’ direction is more difficult and confusing than keeping them moving along their current vector, and to limit the number of fences that are downloaded to the collars. This optimization can be implemented by giving turns a large cost in the A\* search.

To create a plan, instead of doing an expensive search in five dimensions (paddock location, width, height

---

<sup>1</sup>Both extent and shape may change.

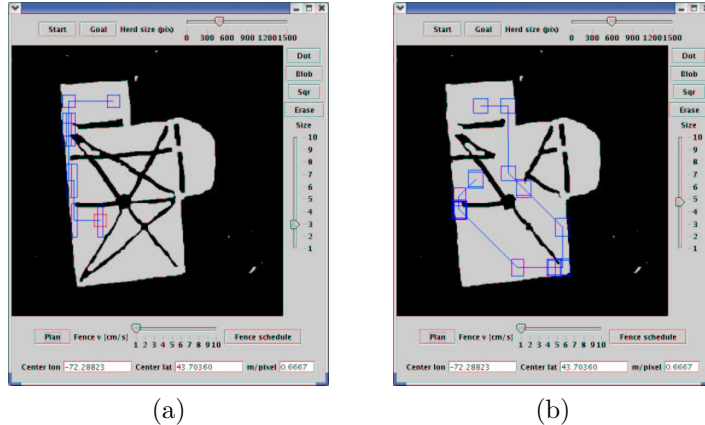


Figure 5: GUI for dynamic fence planning. Environment is based on the Dartmouth Green, with some real-world paths defined to be obstacles. In (a), finding a path from the start (in top portion of free space) to the goal (at lower left) requires a significant change in the size and shape of the virtual paddock. (b) A more complex path — note that small overlap with obstacles is considered acceptable as long as sufficient free space remains in the virtual paddock.

and motion direction) we first generate a pair of baseline plans. The first uses a point-sized virtual paddock and the second a constant-sized square paddock. These plans can be computed using the motion operators only. If the two plans are similar in length and complexity then the square paddock plan can be used as is. If not, we then gradually reduce the size of the square paddock until a plan is successful. The successful plan is then used as a base for a path with a variable-sized virtual paddock. The search for this final path is done efficiently in three dimensions (width, height and distance along the base path.) The GUI developed for this planner along with examples of resulting plans is shown in Fig. 5.

Plans are turned into schedules of fences that are executable by the collar. For each segment of the path, four fences are required to define the paddock. A velocity is set for the fences, which in turn gives each segment a time interval over which its fences are active. This list of fences, with a point, normal vector, speed and relative time interval for each, is then given to the collars. An example of this type of fence schedule is given in Sec. 4.5. The collars wait for an initialization message which tells them to make the time intervals absolute from that moment and continue to evaluate the fences and make appropriate stimuli for the course of the path.

It should be noted that the planner is only planning for the virtual fences, and assumes the desired behavior of the animals within. In the future, we plan to validate quantitative models for the animals' reaction to the stimulus. This will tell us how effective the moving fences will be, and in turn how the fences should be moved to produce the desired motion of the herd. It may be possible to make the motion adaptive to the position of the animals and only move when all animals are away from the advancing fenceline.

## 3 Experimental setup

### 3.1 The farm

Cobb Hill farm is a cooperative dairy farm in Vermont. An aerial view is shown in Figure 6. The three fields in which experiments were conducted are shown outlined in black. Field 1 is a long narrow field on the side of a steep hill, the top of which is on the left in the photo. There are two strips of bushes and trees dividing the field into three parts. The cows can walk around the dividers to reach the higher pastures. The trees in the dividers and the trees around the edge of the field provide shade. Field 2 is a larger field with much more open area. It is also on the side of a hill, the top of which is the left edge of the field in the photo. Trees around the West and South sides provide shade. The field is open with few obstructions (a stand of trees near the north end and a few steep inclines.) Field 3 is a narrow pasture that gradually widens out.






Figure 6: Aerial view of Cobb Hill farm. The fields where experiments were conducted are outlined in black. North is up. The photo displays an area approximately 1 km on a side.

Trees on the east edge provide shade. The land in field 3 is relatively flat with no obstructions to movement. A water trough is located half way up field 1 on the South side, and at the entrances to fields 2 and 3.

Collars were fitted to the cows in the barn in the morning and they were then released for the day. Their habit is to move fairly quickly to the far end of each field, and then wander around slowly till evening. The experiments were conducted in summer and fall of 2003.

### 3.2 The Smart Collar Hardware

Prototypes of a Smart Collar were constructed using commercial off-the-shelf components that are readily available. Figure 7(a) shows the components of a collar. The computer is a Zaurus PDA with a 206MHz Intel StrongArm processor, 64MB of RAM, with an additional 128MB SD memory card. It runs Embedix Linux with the Qtopia window manager. The Zaurus has a serial port and stereo sound port. A Socket brand 802.11 compact flash card provides a wireless network connection. An eTrex GPS unit is connected to the serial port of the Zaurus. A small Smokey brand guitar amplifier is used to reproduce sounds from the Zaurus audio port. A fully assembled collar is shown in Figure 7(b). Figure 7(c) shows a cow wearing an early version of the collar.

The counterweight is required to stop the collar from rotating, the GPS receiver and PDA WiFi card need to be on the top of the neck to ensure good satellite and radio reception. However in practice the collars do rotate somewhat. A very practical problem is water from long grass and early morning dew, or when the animal drinks at the trough and the lowest part of the collar goes in the water. The collar is not fully waterproof, though it is fairly water resistant since the GPS and audio amp are well sealed and the speaker has a plastic cone. The Zaurus is enclosed in a plastic case which gives it some water resistance, although the holes for the cables will allow some water in if the collar has rotated. The batteries in the Zaurus are the limiting factor in how long the collar will run, giving about two hours and forty minutes of life. The audio amplifier and speaker will produce about 90 to 100dB volume at a one foot range, depending on the nature of the sound.

The Zaurus has a custom kernel which allows running the WiFi card in ad-hoc peer-to-peer mode. This allows us to do multihop forwarding of messages for better connectivity within the herd. Ssh and scp are installed to allow remote login to the Zaurus and field upgrades of software. Each Zaurus is also configured with a shell terminal program and has a foldup keyboard for accessing and running programs directly from the console.

A laptop computer is used as a basestation for sending commands to the collars. A Cantenna brand directional WiFi antenna is used with the basestation to improve communication range to the herd.

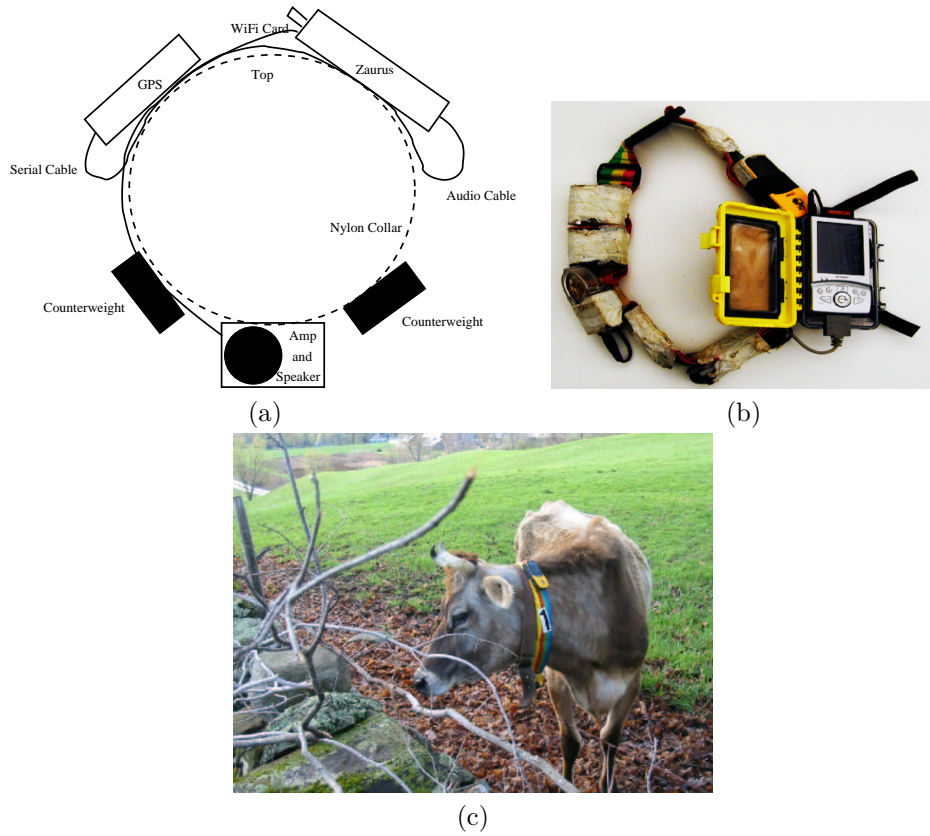


Figure 7: (a) The components of the Smart Collar include a Zaurus PDA, WiFi compact flash card, eTrex GPS, protective case for the Zaurus, an audio amplifier with speaker, and various connecting cables. (b) A fully assembled Smart Collar, with PDA case open. (c) A cow with a collar in Field 3.



Figure 8: Peter Corke and Zack Butler log into the networked cows to download code during an experiment at Cobb Hill Farms.

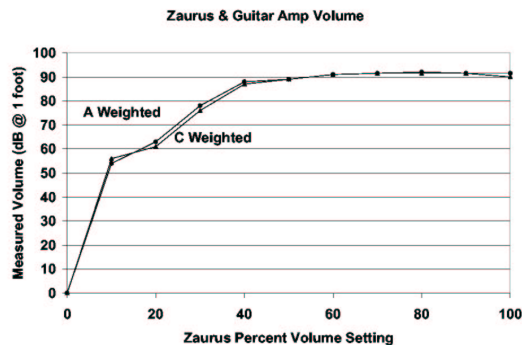


Figure 9: Volume of sound produced for Zaurus volume settings. The "C Weighted" notation indicates the sound level meter has applied a filter to adjust for the frequency response of the human ear. The A Weighted curve does not have this compensation.

### 3.3 Software Infrastructure

The components of the software used in the experiments are as follows:

#### 3.3.1 Virtual Fences

A fence is essentially defined as a point on the surface of the earth, an "interior" direction, and a velocity. The velocity can be used to move the fence over time toward or away from the specified direction. Fences can be added or removed at any time and several of them can be created at once from definitions stored in a file. Several fences can be combined to create convex polygonal shapes. When the GPS readings indicate a cow has crossed a fence a sound is triggered. The sounds are stored in WAV format files and can be selected from a list to be played on the Zaurus audio device. The sounds used in our experiments were natural such as cows and dogs; more alarming such as lions, tigers, panthers, wolves and thunder; and unnatural such as a car crash and a helicopter<sup>2</sup>.

The volume of sounds is controllable on a percentage scale from zero to 100 percent. All fences use the currently selected sound and volume, which can be changed without redefining the fences. Figure 9 shows the relationship between the Zaurus sound settings, which are expressed on a percentage scale, and the actual volume produced by the Zaurus/amplifier combination. The curve becomes flat around 40 percent due to the amplifier becoming saturated and starting to clip. This results in a progressively harsher sound as the volume is raised which is perceived as louder, but which is not actually louder.

An early version of the collar was built with two speakers to allow evaluation of stimuli with a directional effect by differential volume control. However we were unable to control the "balance" on the Zaurus sound card so we reverted to a single speaker.

The fence module also reads and interprets the GPS data which arrives every two seconds when the GPS has a good lock on the satellites. It also sends a periodic Alive message indicating the collar is functional.

#### 3.3.2 Message Handling

Wireless network and Unix pipe messages are used to control the software. The same message format is used interchangeably for both message types. This allows messages to be sent locally, which is useful for testing, and remotely via WiFi for field experiments. All WiFi messages are multihop, being forwarded once by each collar, to improve range and connectivity within the herd. There are two message channels, one outgoing

<sup>2</sup>Our experiments were done under protocol assurance A3259-01 given by the Institutional Animal Care and Use Committee (IACUC) of Dartmouth College.

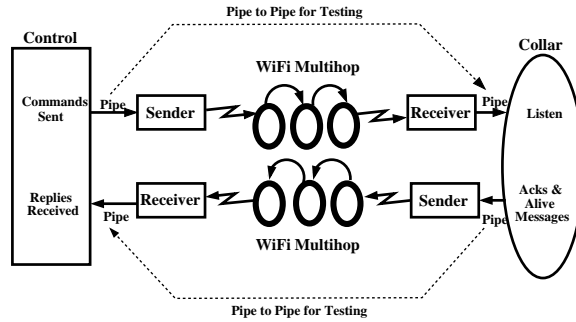


Figure 10: The message routing structure for the networked cattle herd. Messages are routed along multi-hop paths from animal to animal in the herd, i.e. a simple flooding protocol. There are two message channels, once outgoing from a basestation and one incoming to the basestation.

from a basestation and one incoming to the basestation. The outgoing channel is used for defining fences, manually triggering sounds, setting sound type and volume. The incoming channel carries “Alive” messages indicating a collar is active, and acknowledgment messages for receipt and proper interpretation of messages.

### 3.3.3 Experiment Control

Both text and GUI control programs are used to manage the collars in the field. The text control program can be run on a Zaurus or Linux laptop and allows setting and deleting fences, setting type and volume of sound, and manually triggering a sound. The GUI control programs, written in Tcl/Tk, include the functionality of the text program, and provides buttons for triggering sounds on specific cows, a map display showing current cow locations and status (i.e., relationship to fence boundary and whether a sound is playing), and a status display showing whether Alive messages have been received recently from each cow.

Figure 11 shows the control GUIs. The software programs on the collars and basestation are started and stopped with shell scripts for easy reconfiguration.

### 3.3.4 Logging and Time Synchronization

A variety of information is logged on the collars for experimental and debugging purposes including GPS location, GPS time, messages received, messages forwarded, and messages sent. All log entries are accompanied by a time and date stamp. To ensure accurate timestamps across the several programs in the collar system the Zaurus clock is initially synced to the GPS timestamp. Then a `gettimeofday()` system call is used in the various programs to log the current time. The drift in the Zaurus clocks is sufficiently low to provide good time sync for the duration of experiments which typically last two or three hours. Log data is post-processed using custom written scripts in a variety of languages.

## 3.4 Experimental Methodology

We used the methods of direct observation, video taping, and note-taking to evaluate the effectiveness of these sounds in producing desirable reactions. We used the GPS measurements of position and velocity to study the cows’ reaction to sounds. Did they avoid spaces beyond fences? Did they change direction? Did they change walking speed? Looking for correlations between sound events and changes in the GPS data was a primary analysis method, though we found it limited by the spatial accuracy of the GPS position data.

## 4 Experimental results

Our experiments addressed four issues: (1) collecting data to create a grazing model for the cows, which is used in the simulation and fence control algorithm; (2) collecting connectivity data and information propagation data, which is used to assess the performance of multi-hop routing in this environment; (3)

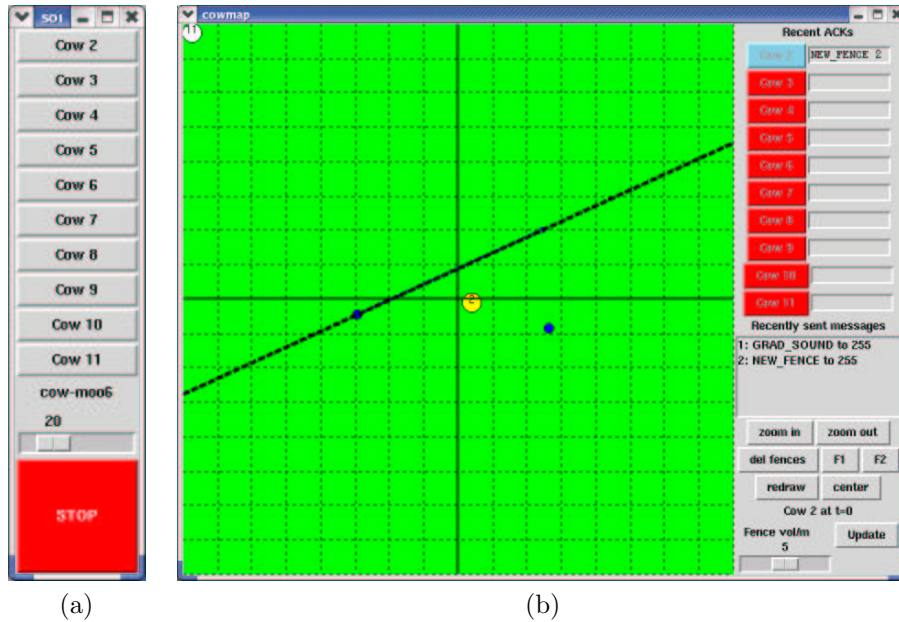


Figure 11: GUIs used on laptops to monitor field experiments. (a) Sound control GUI. Pressing a button triggers the current sound on a specific cow. Current sound and volume can also be selected. (b) Map control GUI. Shows the last reported position of each cow, whether it is currently playing a sound, and whether an Alive message has been received recently. Buttons and text boxes in upper right show recent command acknowledgments from collars.

collecting stimulus-response data for individual animals; and (4) collecting response data for the virtual fence on a group of animals.

Our preliminary results are encouraging. Animals respond to artificial potentials of sounds generated by the virtual fence by moving forward if they are on their own, or toward the group if they are in close proximity to the group. The animals responded to sounds (see Figure 13(b)) but habituation to stimuli was a problem. Others [3] have combined the sounds with shocks to avoid habituation.

#### 4.1 Acquiring a Grazing Model

**Methodology** Our first field experiments were conducted to attempt to verify the two-state grazing model used in simulation. The first experiment involved five collars placed on cows which were released into Field 1. These collars were populated only with the GPS devices and used their built-in tracking function. However, this function is designed to track human hikers who tend to move at a higher and more constant speed, and so this did not give sufficient temporal resolution to test our models. A second experiment with eight full collars on cows in Field 2 allowed for better collection of data, which we were then able to analyze. The eight collars were put on the cows in the morning, and the PDAs recorded GPS positions every two seconds until the battery ran out.

**Data** Figure 12 shows a GPS location track for eight cows moving as part of a herd of 14 cows over time. They start out in the barn, follow a path to the field, and then wander to the far side of the field and back. In addition to the position tracks, the timestamps on the graph give a rough idea of how the herd moved and how spread out they were over time. For an idea of scale, the trek from one end of the field to the other covered a distance of about 300 meters.

In order to look at an appropriate sample of the cows behavior, we present only the data from after the cows had reached the field. A histogram of the velocity for one cow is presented in Fig. 13. Each sample represents the difference in consecutively recorded positions, usually two seconds apart. Due to the resolution

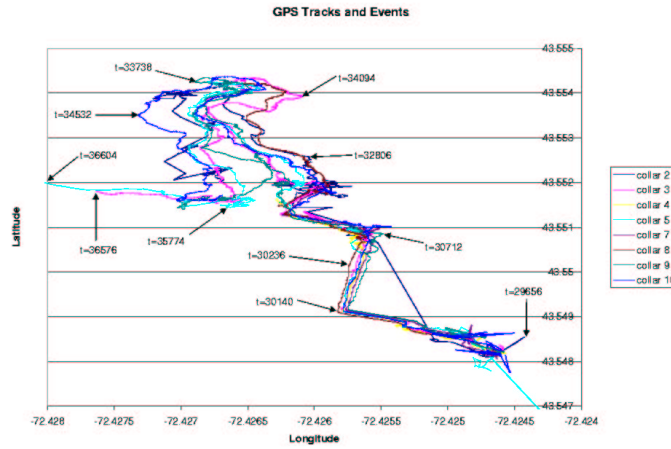


Figure 12: GPS position track of eight cows, moving as a herd, over time. A few timestamps give some sense of how the cows wandered with respect to each other.




Figure 13: Histogram of speed of one cow over a period of 40 minutes. (a) Based on raw GPS differences (b) Based on a 10-second moving average speed. Note difference in vertical scales.



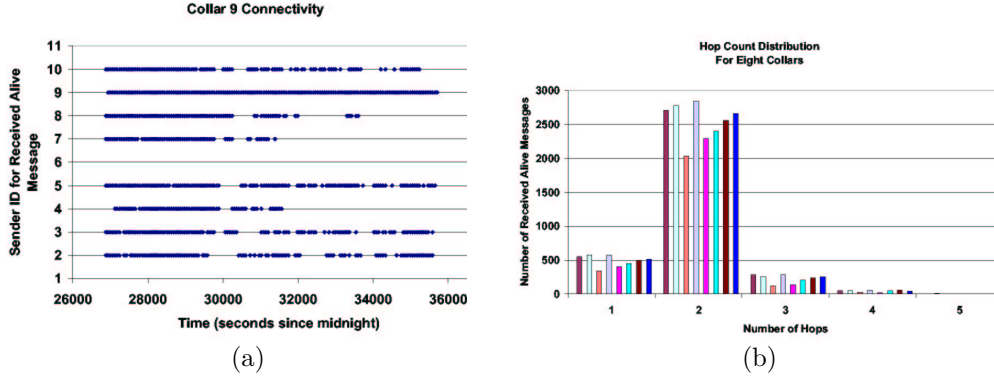


Figure 14: (a) Time history of Alive messages received by collar 9 in a typical experiment. (b) Average number of hops for an Alive message to reach the basestation.

of the GPS data, we also present a 10-second moving average of the speed data. These plots represent data collected from one cow, but other cows in the herd display very similar overall velocity profiles.

**Discussion** These data show that the cows have a range of speeds throughout the day, and the distribution is approximately bimodal: we see that the animals spend a large amount of their time moving quite slowly, and the rest of the time at higher speeds. The higher speeds vary. During our field observations we noted higher speeds when an animal moved from a grazing area to another. Even higher speeds were observed occasionally when it tried to catch up when disconnected from the rest of the group and when the animal seemed to respond to an external fearful factor such as a human or a barking dog. The speeds in the first category seemed to depend on the distance gap between the animal and the rest of the herd. The average speed for the grazing behavior is under 0.2 m/s. This can be detected reliably by consecutive GPS readings with differential GPS, and it also can be detected with lower-end GPS (like the one we used) from the smoothed speed data. For Cow No. 10, using the smoothed speed gives a smooth distribution that peaks at around 0.16 m/s. Setting a cutoff for the grazing of 0.4 m/s gives a mean grazing speed of 0.167 m/s. The cows also spent a significant time at higher speed, walking from one grazing spot to another. This behavior takes place about 15% of the time (183 raw samples or 165 smoothed samples above 0.4 m/s from 1200 total samples). However, the walking takes place at a uniform range of speeds up to 1.25 m/s, rather than a single walking speed as originally supposed.

## 4.2 Network connectivity

**Methodology** The collars periodically emit an “Alive” message which are propagated by the multi-hop message layer using a form of flood forwarding. Each collar also logs all “Alive” messages that it receives, along with the sender id and time.

**Data** Figure 14 shows an example of the connectivity achieved between collars over time. In the first half of this graph there is very good connectivity during the time the cows were all together in the barn. Around 30,000 seconds most connectivity is lost as the cows are walking end to end along a narrow path out to the field. On the right side of the graph connectivity varies as the cows wander around the field, their bodies and tall wet grass being the main causes of signal obstruction. Connectivity to collar 7 is lost before 32000 seconds because of an equipment malfunction. Collar 6 was not deployed. Each collar was connected to at least one other collar 97.5% of the time. On average, each collar was connected to 66% of the rest of the collars, via one or more hops. Thus there is room for improvement in the multihop routing algorithm which is based on simple forward-once flooding.

Figure 14 shows the number of hops required for an Alive message to reach the laptop basestation during an experiment. Most messages are relayed only once to reach their destination which indicates good

connectivity between collars.

**Discussion** Dynamic graphs of the message routing have shown us that connectivity among the herd is usually quite good since the cows tend to stay near each other. Connectivity with the base station was problematic in that there is a trade off in staying far enough away to not influence the herd (they are very curious and friendly) and staying close enough to maintain radio contact. WiFi networks are essentially line of sight and are blocked completely at times by the cows bodies. Switching to VHF transmitters to improve basestation connectivity is an option we are considering.

### 4.3 Individual Response

**Methodology** In all field experiments we visually observed the behavior of individual cows, both away from the herd and when amongst the herd. We often video taped moments of interest, see Extension 6, and have the data logs from all experiments as a basis for analyzing individual behavior. GPS velocity data was also used to correlate stimulus events with changes in velocity.

**Data** A single sound sometimes had an effect on some of the cows and had no effect when tried on other cows. Some cows never reacted to sounds, while others were more sensitive. Observed reactions to a stimulus sound varied widely and included stop eating and look up, no reaction, stop eating, look up, then walk a short distance, usually forward, etc. (see Table 1).

The velocity of an individual can be derived from the GPS tracking data to discover if a cow reacted to a sound stimulus by changing speed. Figure 15 shows a time history of the speed for a cow in our second experiment. The asterisks denote when a sound was played. In this case there seems to be a good correlation between sound events and the cow being in motion. However, some of the animals responded in a less correlated way. Two difficulties in interpreting this kind of data are, first, cows may already be in motion when stimulated, and second, the GPS data is very coarse in time (a reading every two seconds) which makes it difficult to judge if the cows motion was actually in reaction to the stimulus.

Table 1 shows some of the observed responses. We generally noted that repeated and louder sounds were more effective in eliciting a response. Cows would often react to the first instance of a sound and then not react to further instances. Waiting a half hour would sometimes result in them reacting again to an initial sound.

**Discussion** Some cows definitely reacted strongly to a sound stimulus, though they often quickly became inured to it, and stopped reacting. The orientation of the cow before the stimulus was applied played a role in determining what direction a cow moved, if it moved. Further research into effective stimulus methods and into invoking directional behavior are needed. Much louder sounds may be more effective. Sounds accompanied by something visible such as a puff of smoke may be more effective and provide some steering capability.

### 4.4 Static Virtual Fence Experiments

**Methodology** In the final field experiment, we used a total of six collars to test the effects of the virtual fence on the herd. These collars were put on the cows with one virtual fence already present, allowing us to be sure that the fence would be present even if we experienced base-station communication failures. The cows were sent into field 1, with a north-south oriented virtual fence located across the paddock about one third of the way up. We observed the cows' reactions visually to supplement the logged data. After the cows had moved through the preset fence and the fence had timed out, a second north-south fence was instantiated near the top of the paddock (approximately under the "1" label in Fig. 6). Both fences used the graduated volume algorithm, with a value of 7%/m for the first fence and 5%/m for the second.

**Data** Of the six collars, two performed very well for the duration of the experiment, two performed well but for a shorter time (perhaps due to battery failure) and two had poor to nonexistent GPS signal, probably due to rotation of the collars on the cows' necks. Figure 16 shows data from one cow's collar over the entire



Cow	Orientation	Sound/Volume	Reaction direction	Reaction magnitude	Comments
10	6	air/50	6	1 step	rapid 3 in a row, startled
10	6	air/50	-	-	2 in a row, nothing
10	6	dog/40	-	-	lifted head, looked to one side
10	6	dogx4/50	12	6 steps	walked forward while sound was on
3	-	air/40	-	-	very startled, shuddered at each sound
3	-	air/26	-	-	no reaction
3	-	air/60	-	-	no reaction
3	-	air/80	-	-	no reaction, habituated ? (we were close to the cow)
3	-	dog/50	-	-	no reaction
8	12	airx6/50	12	walked	started walking for duration of walk (initial 2s delay), actually moved toward us
8	12	airx6/50	-	-	no reaction (her back to us)
8	12	cymb/50	12		walked for duration
8	12	dog/50	10		cow and neighbour moved
8	12	dog/50	-	-	no reaction
8	12	cymb/50	-	-	no reaction
8	12	hiss/50	-	-	no reaction
8	12	crash/50	-	-	no reaction, neighbours looked up
8	12	air/100	-	-	flicked her tail
8	12	dog/100	-	-	no reaction

Table 1: Observed reaction to stimulus. Time increases from top to bottom in the table. Orientation and reaction direction are specified using “hours on the clock” notation where noon is North. Volume is percent volume setting on the Zaurus.




Figure 15: Velocity of cow 10 versus time with sound stimulus events noted by asterisks. There appears to be good correlation between sound events and cow motion.

experiment. Both fences are shown here relative to the cow's travels in the field. This figure shows the sounds being applied at the correct locations for both fences. We also show a closeup of another collar's data, showing that the fence worked correctly over multiple crossings, with the fence timeout resetting as desired. In addition, to analyze the effects of the fence, we looked at the speeds of all the cows during the times sounds were being played relative to the rest of the time.

**Discussion** Our visual observations were that in general, the cows noticed the sounds, but either ignored them or did not make the desired association with their position. For two of the cows, we observed the animal stop grazing when the sound was played, look up and walk slowly in a different direction. However, this new direction was not sufficiently different to take the cow into the interior area, and further sounds seemed to be ignored. We also observed one cow essentially ignore the sounds entirely. We were told that the cows tended to be motivated to reach the top of this paddock, especially first thing in the morning, and this motivation may have been too strong for the sounds to overcome. However, the second fence nearer the top of the hill was also not effective at keeping the animals on the desired side.

We also analyzed the logged data for the two cows that recorded good data for both fences. For both cows, the logs seem to indicate that the first fence slowed the cows' progress toward the top of the hill. This was determined by comparing for each cow (1) the cow's speed between entering the field and reaching the first fence and (2) the cow's speed while the first fence was causing sounds to be played. For cow 10, the average speeds for these two time periods were 0.380 m/s and 0.255 m/s respectively, and for cow 9, 0.590 m/s and 0.388 m/s respectively. For both, this difference is significant at the 0.01 level using a t-test, and the form of the speed distributions for these time periods looks quite similar. Later speed data is less convincing due to habituation to the stimulus. For cow 9, after the first fence stops making noise, up through and including when the second fence makes noise, its speed did not change significantly, whereas for cow 10 there was a speed increase between the fences and decrease for the second fence. For both cows, once they had reached the top of the field, their speed and range decreased significantly (again, using a t-test with a 0.01 significance level), both just under 0.2 m/s on average, similar to the grazing speeds seen in the earlier experiment.

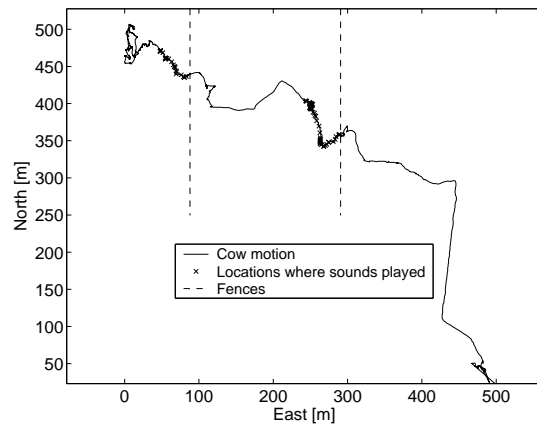


Figure 16: Trace of the position of cow #10 during an experiment with two virtual fences. The cow started in the barn, at the SE corner of the plot. Locations where sounds were played automatically are shown. The long straight line to the north is the walk from the barn to the field, which is not considered in our analysis. After the first fence (to the east) timed out, no sounds played until a second fence (to the west) was created.

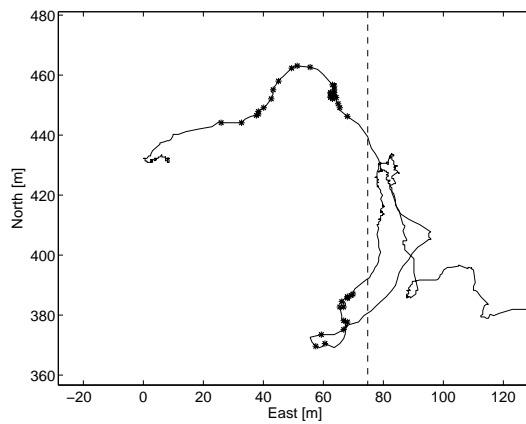


Figure 17: Trace of the position of cow #9 during the virtual fence experiment. A small portion of the experiment is shown, during which the cow walked past the fence on two separate occasions, showing correct behavior of the virtual fence algorithm.

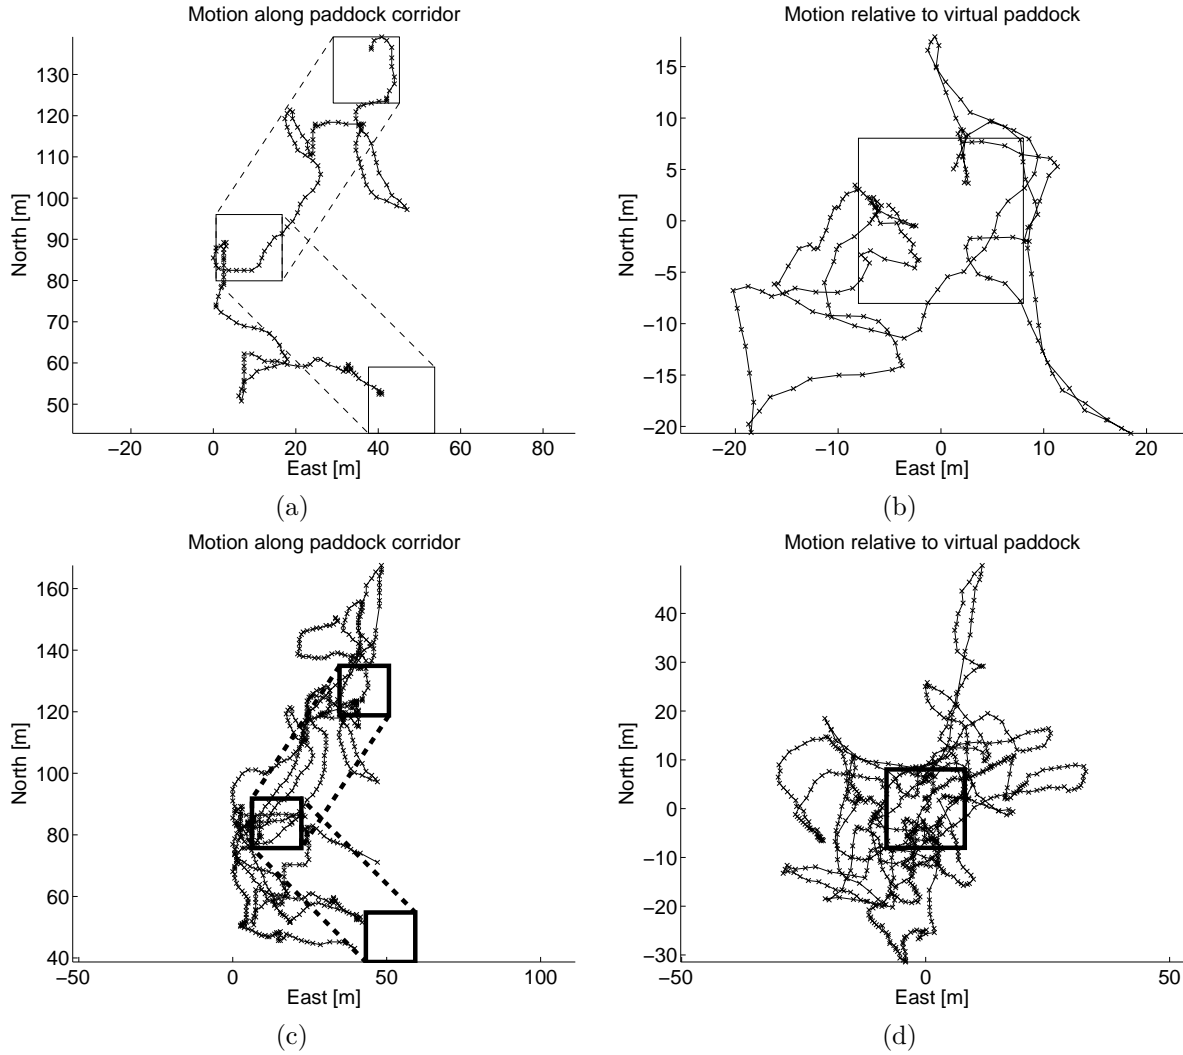


Figure 18: Dynamic fence experiment results. (a) Absolute and (b) relative motion of one person. (c-d) Absolute and relative motion data for four people.

## 4.5 Dynamic Virtual Fence Experiments

**Methodology** We have tested the dynamic fence algorithm using the smart collars on people. Plans were created in the planning GUI and exported directly to the collars. The experiment will be conducted with different hardware and different animals at the USDA<sup>3</sup> Jornada test range in New Mexico with the assistance of Dr Dean Anderson, but those results are not yet available.

**Data** Motion traces for one person and the group of four people are shown in Fig. 18. Fig. 18(a) shows the absolute coordinates of the paddock path and the person within the paddock. However if we plot the position of the person relative to the instantaneous position of the paddock the situation is more interesting, as shown in Fig. 18(b). The schedule used to generate this experiment is showed below.

<sup>3</sup>US Department of Agriculture




Figure 19: (a) The Fleck2 tracking board, and (b) incorporated into a collar.

fence-long	fence-lat	norm-x	norm-y	speed(m/s)	t-start (s)	t-end(s)
-72.28808	43.70385	1.00	0.00	-0.1302	0	215
-72.28808	43.70385	0.00	-1.00	0.2000	0	215
-72.28788	43.70370	-1.00	0.00	0.1302	0	215
-72.28788	43.70370	0.00	1.00	-0.2000	0	215
-72.28843	43.70346	1.00	0.00	0.2000	215	400
-72.28843	43.70346	0.00	-1.00	0.2000	215	400
-72.28823	43.70331	-1.00	0.00	-0.2000	215	400
-72.28823	43.70331	0.00	1.00	-0.2000	215	400

This schedule represents the fences used in the experiment portrayed in Fig. 18. The first four lines represent the four fences that form the initial safe square at the top of Fig. 18(a) and (c). Note that the normal vectors for these four fences are along each of the four cardinal directions, and that they all start at time 0 and continue for the same amount of time. After 215 seconds, these fences reach the lines shown by the square in the middle of the plots. The second set of four fences exactly overlap the first four at this moment, but have different velocities, so the safe square seamlessly moves off in a different direction, and the first set of fences is no longer in existence to impede travel.

**Discussion** The results were successful algorithmically, in that the fences appeared in the correct real-world locations at the desired times, but were less successful from a herding standpoint. The absolute position trace shows that the general motion of the people was along the desired corridor, as in Fig. 18(a). However from the paddocks perspective, the people spend very little time within its bounds, as seen in Fig. 18(b). This may be biased by the experimenters walking in the general direction of the virtual paddock motion.

One benefit of this experiment was that the subjects were able to verbalize their feelings about the system. Primarily, they felt that the resolution of the stimulus gradient was insufficient to be of assistance. Also, they tended to be inquisitive, actively exploring the acceptable boundaries of their space. Together with relatively high walking speeds (averaging over 0.5 m/s) they were able to move a large distance between sounds and were unable to relocate the virtual paddock. We believe that with a more appropriate stimulus (such as Anderson’s stimulus for cows [3]) that this can be overcome.

## 5 Conclusion

In this paper we have introduced the concept of a virtual fence, a device which applies a stimulus to an animal as a function of its pose with respect to one or more fencelines. The fence algorithm is implemented by a small position-aware computer device worn by the animal, which we refer to as a Smart Collar. We have described a simulator based on potential fields and stateful animal models whose parameters are informed by field

observations and track data obtained from the Smart Collar. We have observed the effect of sound stimuli on the animals but have had difficulties due to habituation, and the option of electric shock stimulus has not been available to us. Field implementations of static virtual fencelines have been tested and despite the limitations of sound stimuli, a change in behaviour was observable. Dynamic virtual fences were implemented and tested on human subjects which gave some interesting perspectives which we cannot get from animal subjects.

Our approach to modeling animal behavior is based on animal behavior literature, our limited empirical observations of cow behaviors, and our discussions with cattle farmers and experts from the USDA and from Heytesbury Beef, Australia. This simple model has proved to be effective. However, in the future we would like to use the animal tracking data to establish more accurate behavior models and parameter.

Our future work has a number of different directions and different locations. We will work with researchers at the USDA in New Mexico to trial the dynamic virtual fence concept. In Australia we are trialling a new collar [26] which is a variant of a Mote sensor network device, comprising an AVR microprocessor, 433MHz radio link, GPS, 3-axis acceleration, 3-axis magnetic field and temperature, programmable in TinyOS and with the ability to hold considerable amounts of data in an MMC flash memory card, see Figure 19. We also wish to take animal position data and extract information about herd social structure and dynamics. This will permit more significant grounding of this work on real data which will lead to new and improved models. These models will lead to a better understanding of cattle behavior and control at the individual and group level, which has the potential to impact not only the cattle industry, but more broadly, agriculture. By controlling grazing patterns, there will likely be an ecological impact as the land will be more carefully utilized.

## A Index to Multimedia Extensions

Extension	Type	Description
1	Video	Simulation of unforced animal motion
2	Video	Simulation of forced animal motion, slow threat object
3	Video	Simulation of forced animal motion, fast threat object
4	Video	Simulation of a virtual fence
5	Video	Simulation of two animal populations being separated
6	Video	Observed effect of stimulus on animals in the field

## Acknowledgments

There are many people we would like to thank for assisting us in this work. Simon Holmes a Court inspired us to think about this problem and has given us great technical insights into cows and herding. Steve Petty of Heytesbury Beef Australia and the staff at Pigeon Hole station showed great hospitality and provided valuable discussion and suggestions. We especially thank Steven, Kerry and Paul at Cobb Hill for their assistance and enthusiasm and the use of their cows for these experiments. We thank Lorie Loeb for facilitating our work at Cobb Hill Farms. We also received some valuable cow domain knowledge from Bud and Eunice William's Stockmanship School ([stockmanship.com](http://stockmanship.com)) who teach low-stress cow herding control to farmers. Tom Temple assembled the cow collars. Jenni Groh provided invaluable advice on the experimental procedure. The experimental part of this work has been done under protocol assurance A3259-01 given by the Institutional Animal Care and Use Committee (IACUC) of Dartmouth College. Dean Anderson of USDA Agricultural Research Service at the Jornada Test Range in New Mexico has been very helpful in discussions regarding virtual fencing. Dave Swain of CSIRO Livestock Industry has been a keen supporter of the ongoing Australian cattle tracking experiments.

## References

- [1] Jean Marc Amé, Colette Rivault, and Jean Louis Deneubourg. Cockroach aggregation based on strain odour recognition. *Animal Behaviour*, page in press, 2003.

- [2] D.M. Anderson and C.S. Hale. Animal control system using global positioning and instrumental animal conditioning. Technical Report US Patent 6,232,880, USDA, May 2001.
- [3] D.M. Anderson, C.S. Hale, R. Libeau, and B. Nolen. Managing stocking density in real-time. In N. Allsop, A.R. Palmer, S.J. Milton, K.P. Kirkman, G.L.H. Kerley, C.R. Hurt, and C.J. Brown, editors, *Proc VII International Rangelands Conf.*, pages 840–843, Durham, South Africa, August 2003.
- [4] T. Balch. Hierarchic social entropy: an information theoretic measure of robot team diversity. *Autonomous Robots*, July, 2000.
- [5] T. Balch, Z. Khan, and M. Veloso. Automatically tracking and analyzing the behavior of live insect colonies. In *Autonomous Agents*, 2001.
- [6] Tucker Balch and Ronald C. Arkin. Behavior-based formation control for multiple mobile robots. *IEEE Transactions on Robotics and Automation*, 14(6):929–39, December 1998.
- [7] Tucker Balch and Maria Hybinette. Social potentials for scalable multi-robot formations. In *Proc. of Int'l Conf. on Robotics and Automation*, pages 73–80, San Francisco, April 2000.
- [8] Z. Butler, P. Corke, R. Peterson, and D. Rus. Dynamic virtual fences for coltrolling cows. In *in the 2004 International Symposium on Experimental Robotics*, 2003.
- [9] Z. Butler, P. Corke, R. Peterson, and D. Rus. Virtual fences for controlling cows. In *ICRA*, New Orleans, LA, 2004.
- [10] G. Caprari, A. Colot, R. Siegwart, J. Halloy, and J.-L. Deneubourg. Building mixed societies of animals and robots. *IEEE Robotics and Automation Magazine*, page in press, 2004.
- [11] G. Caprari, A. Colot, R. Siegwart, J. Halloy, and J.-L. Deneubourg. Insbot: Design of an autonomous mini mobile robot able to interact with cockroaches. In *in the 2004 International Conference on Robotics and Automation*, 2004.
- [12] Don Comis. The cyber cow whisperer and his virtual fence. *Agricultural Research*, November 2000.
- [13] J. Gautrais, C. Jost, R. Jeanson, and G. Theraulaz. How individual interactions control aggregation patterns in gregarious arthropods. *Interaction Studies*, page in press, 2004.
- [14] L. Howery, D. Bailey, and E. Laca. Impact of spatial memory on habitat use. Technical report, 1999.
- [15] R. Jeanson, S. Blanco, R. Fournier, J.-L. Deneubourg, V. Fourcassié, and G. Theraulaz. A model of animal movements in a bounded space. *Journal of Theoretical Biology*, 225:443–451, 2003.
- [16] Philo Juang, Hidekazu Oki, Yong Wang, Margaret Martonosi, Li-Shiuan Peh, , and Daniel Rubenstein. Energy-efficient computing for wildlife tracking: Design tradeoffs and early experiences with zebanet. In *Proceedings of the Tenth International Conference on Architectural Support for Programming Languages and Operating Systems*, 2002.
- [17] Alfred J. Lotka. *Elements of Physical Biology*. Williams and Wilkins Co., 1924.
- [18] R.E. Marsh. Fenceless animal control system using gps location information. Technical Report US Patent 5,868,100, Agritech Electronics, February 1999.
- [19] A. Martinoli, G. Theraulaz, and J.-L. Deneubourg. Quand les robots imitent la nature. *La Recherche*, 358(november):56–62, 2002.
- [20] Maja J. Mataric. Kin recognition, similarity, and group behavior. In *Proc. of the Fifteenth Annual Cognitive Science Society Conference*, Boulder, Colorado, 1993.
- [21] P. Michelena, P. M. Bouquet, A. Dissac, V. Fourcassié, J. Lauga, J.-F. Gérard, and R. Bon. Do female and male merino sheep segregate within the same ecological conditions. *Animal Behaviour*, page in press, 2003.
- [22] Lynne E. Parker. ALLIANCE: An architecture for fault tolerant, cooperative control of heterogeneous mobile robots. In *Proc. of IEEE Int'l Conf. on Intelligent Robots and Systems*, pages 776–83, Munich, Sept. 1994.
- [23] T.M. Quigley, H.R. Sanderson, A.R. Tiedemann, and M.K. McInnis. Livestock control with electrical and audio stimulation. *Rangelands*, June 1990.
- [24] John H. Reif and Hongyan Wang. Social potential fields: A distributed behavioral control for autonomous robots. *Robotics and Autonomous Systems*, 27(3):171–94, May 1999.
- [25] S. Russell and P. Norvig. *Artificial Intelligence: A Modern Approach*. Prentics Hall series on Artificial Intelligene, 2002.

- [26] Pavan Sikka, Peter Corke, and Leslie Overs. Wireless sensor device for animal tracking and control. In *First IEEE Workshop on Embedded Networked Sensors within the 29th Conference on Local COmputer Networks*, pages 446–454, 2004.
- [27] B. Smith. *Moving 'Em: A Guide to Low Stress Animal Handling*. Graziers Hui, 1998.
- [28] A.R. Tiedemann, T.M. Quigley, L.D. White, W.S. Lauritzen, J.W. Thomas, and M.K. McInnis. Electronic (fenceless) control of livestock. Technical Report PNW-RP-510, United States Department of Agriculture, Forest Service, January 1999.
- [29] R. Vaughan, N. Sumpter, A. Frost, and S. Cameron. Robot sheepdog project achieves automatic flock control. *Proc. Fifth International Conference on the Simulation of Adaptive Behaviour.*, 1998.
- [30] I. Veissier, A. Booissy, R. Nowak, P. Orgeur, and P. Poindron. Ontogeny of social awareness in domestic herbivores. *Appl. Anim. Behav. Sci.*, (57), 1998.