

Multi-Agent Trail Making for Stigmergic Navigation

Alfred Wurr and John Anderson

Autonomous Agents Laboratory
Department of Computer Science
University of Manitoba
Winnipeg, Manitoba
Canada R3T2N2
Email: awurr, andersj@cs.umanitoba.ca

Abstract. Robotic agents in dynamic environments must sometimes navigate using only their local perceptions. In complex environments, features such as terrain undulation, geometrically complex barriers, and similar obstacles form local maxima and minima that can trap and hinder agents using reactive navigation. Moreover, agents navigating in a purely reactive fashion forget their past discoveries quickly. Preserving this knowledge usually requires that each agent construct a detailed world model as it explores or be forced to rediscover desired goals each time. Explicit communication can also be required to share discoveries and coordinate actions. The cost of explicit communication can be substantial, however, making it desirable to avoid its use in many domains. Accordingly, in this paper we present a method of cooperative trail making that allows a team of agents using reactive navigation to assist one another in their explorations through implicit (*stigmergic*) communication.

1 Introduction

Robotic agents that employ reactive navigation suffer from several problems. Agents that are only reacting may be attracted to a local stimulus, but one that is really only a waypoint on route to a goal - a local maximum [1, 2]. Box-canyons, for example, are a common local maxima problem, where agents are unable to see a goal of any type [3] and are reduced to wandering randomly before stumbling upon an exit [4]. On the other hand, when a goal is visible, these same agents may attempt to move in a direct line toward it, but be blocked by an interposing barrier - a local minimum [2]. Cyclic behaviour is another potential problem, where agents oscillate between multiple stimuli [5]. Compounding these problems, agents navigating in a purely reactive fashion, forget their past discoveries even as they are made (unless they are constructing a detailed world model as they explore) and as a result their explorations are unsystematic (they may cover the same ground repeatedly) [6].

Many researchers have attempted to remedy these and other problems using *stigmergy*, a term used in biology to describe the influence that previous changes

to the environment can have on an individual’s current actions [7]. To date, most work with stigmergy has dealt with its effects on simple homogeneous robots collectively performing foraging or sorting tasks (e.g. [8, 7, 1, 9]) based on that of the natural creatures (e.g. ants and termites) that inspire the model. More sophisticated uses of stigmergy or stigmergy-inspired approaches have been proposed by Werger and Mataric [8, 10], Balch and Arkin [4], Vaughan et al. [11] and Parunak, Sauter, et al., [9, 12] to solve problems in navigation and other common tasks. Ant and other biologically-inspired forms of computation have also been studied in numerous other tasks (e.g. [13]).

2 Stigmergic Navigation

In this paper, we employ a number of different stigmergic techniques that involve the deployment and exploitation of physical markers (as opposed to markers used only internally and communicated explicitly) in order to assist in team navigation. We refer to the process of using these marking techniques collectively to make more purposeful reactive navigation decisions in an unknown environment as *stigmergic navigation*.

The primary method employed involves a team of agents constructing marker trails leading to discovered goal locations, cooperatively and dynamically. A number of simpler methods are used as additional supports (and which are effective on their own as well). These include marking bottlenecks, marking local maxima and marking local minima [14]. Marking bottlenecks involves creating marker trails through constricted areas (i.e. hallways and doorways), which draw agents to new areas. Local maxima markers assist agents in negotiating areas devoid of sensory stimulus by first attracting agents (drawing them from walls of rooms to areas of greater visibility) and then repelling them outward toward potential exits. Local minima markers are unique in that they prescribe an action to undertake in response to a specific situation at a particular place in the environment, rather than a more basic attractive/repulsive force. In our specific implementation, these markers are dropped when an agent perceives an attractor and is blocked by an interposing barrier. Agents react to local minima markers by jumping when sufficiently close to them. This has the effect of causing agents to avoid a specific type of local minima by leaping over the problem. Alternatively, reactions such as being repelled by local minima markers are also possible, but not explored in the experiments outlined here.

More detailed explanations of these supplemental marking techniques are available in [14]. The following section provides a detailed overview of stigmergic trail-making that is the focus of this paper.

2.1 Stigmergic trail-making

Stigmergic trail-making is a process whereby a team of agents implicitly *cooperates* in the construction of a purposeful trail of markers leading to a goal or goals while exploring their environment [14]. By following these trails, agents

are then able to locate a previously discovered goal more rapidly on subsequent trips without benefit (or need) of internal world modelling, internal maps or internal path planning. Most notably, these trails are constructed, shared and used entirely without explicit communication.

The process of stigmergic trail-making is best illustrated by way of an example. Consider the situation in Figure 1. A primitive stigmergic trail is created when an agent perceives a goal and drops a marker on the ground at its current location (Figure 1a). The marker dropped is assigned a perceptible numeric value that represents its position in the trail (in this case, 1). The dropped marker identifies a vantage point from which the goal should be visible to an agent standing at the marker's location. By itself, a trail this rudimentary is of limited utility: to make the trail more sophisticated, it is necessary to extend the trail. As Figure 1 illustrates, this occurs when another agent (or the same agent at a later point in time) sees an end-point of the trail and drops a second marker (see Figure 1b). The second marker dropped is assigned a value one higher than the marker that the agent has observed (in this case, 2). As the process repeats, the trail gradually lengthens (see Figure 1c, d).

When the agent perceives the goal (or a marker already on the trail), it checks for the presence of stigmergic markers already serving as attractors to this target. If no markers are visible, the agent drops a marker as described above. This does not preclude the construction of multiple trails during this process. However, when such trails are made they will not be in obvious sight of the one another. This is illustrated in Figure 1e, where an agent sees the end of a trail, by perceiving a 3 marker but no 4, and so extends the trail in another direction. Note that under conditions such as this, extending the trail in different directions is very helpful - agents coming from different directions or from different sides of obstacles can find a useful path to a goal.

An agent only extends the trail when it perceives what appears to be its end. If the agent perceives other markers from its location, it only drops a marker if the goal or a higher valued marker is not also visible. If such a marker is visible, the agent knows it is not at the end of a trail and does not drop a marker. Such a situation is depicted in Figure 1f: the agent perceives marker 3 and two marker 4's; it knows that 3 is not the end of the trail (because it perceives a 4 marker), and that its vantage point is not far enough away to warrant dropping another marker (since existing markers already supply pertinent navigation information from the agent's current location, or the 4 marker would not be present).

These conditions allow an agent to simply follow a trail without dropping new markers when there is enough navigation information available to do so. Thus, a branching situation such as that shown in Figure 1e will only occur when an agent's perceptions indicate it is at the end of the trail when it in fact is not. In such a setting, there is not enough navigation information to conclude otherwise given the agent's current location, and so branching the trail at that point is a logical action.

Agents navigate to the goal at a trail's end by always moving toward the lowest numbered marker visible. Since the markers are dropped based on visi-




Fig. 1. Stigmergic trail making

bility, the agents are able to consistently follow one marker to the next to locate the goal. This also helps minimize the length of a path followed by an agent - if an agent perceives a 3 and a 4 marker, for example, the agent will not bother moving to the 4 marker.

The two important features to emphasize in this approach are a) that this trail-building occurs collaboratively; and b) that it occurs while agents follow existing trails and otherwise navigate through the environment.

In the case of several goals in close proximity, trails may connect. This may cause an agent to perceive more than one lowest valued marker: in this case an agent will be attracted to the marker in closest proximity, so may start following toward one goal and end up at another. Either way, a goal is reached.

3 Evaluation

Agents in this work were implemented using schema-based behaviours, which divide control into *perceptual* and *motor* schemas that are tailored to specific aspects of an agent’s environment [15, 2]. Perceptual schemas were created to detect the conditions under which markers should be dropped. In addition, motor schemas were implemented to cause agents to react appropriately to the presence of markers. Agent control was also composed of a number of other basic schemas (i.e. avoid-obstacles, wander) that were necessary to allow them to explore the chosen simulation environment; the computer game Half-Life a large-scale, 3D environment with realistic physics. The specific environment used was a standard Half-Life map, entitled Crossfire (a simplified 2-D overhead view can be found in [14]). To allow agents to mark their environment, a pre-existing game object (a small satchel) was adapted to function as a basic marker type that agents could drop on the ground as required.

Marker techniques were evaluated in a series of experiments in which the objective was for a team of six agents to find a specific goal at an unknown location in their environment as quickly and as many times as possible in a 30 minute period. Agents began each trial in one of two starting locations, so that agents would not interfere with one another excessively at the start of a trial. The goal itself was located in a room on the far side of the environment. After locating and moving to within 40 units (measured as a fraction of the total size of the environment, which is 8191 units cubed) of the goal, a point was added to the running total of goals discovered in the trial. This agent was then transported back to one of the two randomly selected starting rooms to attempt to find the same goal again. A set of 40 trials were run with a team of 6 agents that did not employ markers and combination of previously outlined marking methods. The environment, goal location and agent starting positions were consistent throughout all experimental trials for both teams.

Table 1 lists the total goals achieved across the 40 trials, the average real-world time (in seconds) and the average number of agent decision-making invocations (*AgentInv*) required to find the goal the first time in a trial. The number of agent decision-making invocations is intended as a secondary measure of time

Table 1. Summary of Results

Marking Method	Total Goals	Average Time	Average AgentInv
None	161	592.35	579822.88
BTl	278	463.67	43917.58
LocMax	224	533.48	52135.00
BTl/LocMax	305	567.02	51697.60
BTl/LocMax/LocMin	323	467.67	24068.98
StgTrl	1828	462.19	45043.30
StgTrl/BTl	2253	508.15	46896.00
StgTrl/BTl/LocMax	3817	525.3	30869.03
StgNav	3856	490.11	24814.38

None=No Markers, BTl=Bottleneck Markers, LocMax=Local Maxima Markers, LocMin=Local Minima Markers, StgTrl=Stigmergic Trails, StgNav=Stigmergic Navigation

passing in the system. While the simple forms of stigmergy all improved navigation, stigmergic trail markers resulted in a very substantial improvement over agents using no markers, or even the improved performance gained using the first two marker types. The best performances, however, occurred when all four marker types were combined. Using all marker types in combination resulted in 3856 goals being reached in 40 trials for an average of 96.4 goals per trial (with a standard deviation of 60.62 goals). This is almost 23 times more goals than agents using no markers.

Stigmergic trail building was assisted by the other marker types because the bottleneck and local maxima markers encouraged agents to travel from area to area more frequently, causing agents to perceive and therefore extend the stigmergic trails more quickly. This allowed the team of agents to benefit from the trail for longer periods of time in each trial, leading to improved performance. The reason for such a large standard deviation in the number of goals in each trial is directly connected to the variation in how quickly the agents first locate the goal and start constructing the trail. In cases where stigmergic trails are built early in the trial, the team is essentially following this trail repeatedly until time expires. If the trail is not built early or at all, performance suffers correspondingly since agents have less time to benefit from it.

4 Summary

In this paper, we have described methods that reactively navigating agents can use to increase the frequency and ease with which they discover and subsequently navigate goals, without need of high-level path-planning, world modelling or explicit communication. While the use of stigmergy in general is not new, the

development of specific techniques operating in a distributed fashion, such as that described here, encourages the application of stigmergy to real-world problems. In general, we believe the maintenance of knowledge in the world as opposed to inside an agent is strongly underappreciated in AI, and in future, techniques such as these will be applied more extensively and viewed as powerful tools.

References

1. Balch, T., Arkin, R.C.: Communication in reactive multiagent robotic systems. *Autonomous Robots* **1** (1994) 27–52
2. Pirjanian, P.: Behavior coordination mechanisms: State-of-the-art. Technical Report IRIS-99-375, Institute for Robotics and Intelligent Systems, University of Southern California, Los Angeles (1999)
3. Sgorbissa, A., Arkin, R.C.: Local navigation strategies for a team of robots. Technical report, Georgia Tech Robotics Laboratory (2001)
4. Balch, T., Arkin, R.C.: Avoiding the past: A simple but effective strategy for reactive navigation. In: *Proceedings of the 1993 IEEE International Conference on Robotics and Automation*. Volume 1. (1993) 678–685 Atlanta, Georgia.
5. Arkin, R.C., Balch, T.: Cooperative multiagent robotic systems. In Bonasso, R., Murphy, R., eds.: *Artificial Intelligence and Mobile Robots*. MIT/AAAI Press (1998) Cambridge, MA.
6. Moorman, K., Ram, A.: A case-based approach to reactive control for autonomous robots. In: *AAAI Fall Symposium on AI for Real-World Autonomous Mobile Robots*. (1992) Cambridge, MA.
7. Holland, O., Melhuish, C.: Stigmergy, self-organisation, and sorting in collective robotics. *Artificial Life* **5** (2000)
8. Werger, B.B., Mataric, M.: Exploiting embodiment in multi-robot teams. Technical Report IRIS-99-378, University of Southern California, Institute for Robotics and Intelligent Systems (1999)
9. Parunak, H.V.D., Brueckner, S., Sauter, J., Posdamer, J.: Mechanisms and military applications for synthetic pheromones. In: *Workshop on Autonomy Oriented Computation*. (2001) Montreal, Canada.
10. Werger, B.B., Mataric, M.: Robotic food chains: Externalization of state and program for minimal-agent foraging. In: *Proceedings of the 4th International Conference on Simulation of Adaptive Behavior: From Animals to Animats 4*, MIT Press (1996) 625–634
11. Vaughan, R., Stoy, K., Sukhatme, G., Mataric, M.: Lost: Localization-space trails for robot teams. *IEEE Transactions on Robotics and Automation* **18** (2002) 796–812
12. Sauter, J.A., Matthews, R., Parunak, H.V.D., Brueckner, S.: Evolving adaptive pheromone path planning mechanisms. In: *Proceedings of the First International Joint Conference on Autonomous Agents and Multi-Agent Systems*, ACM Press (2002) 434–440 Bologna, Italy.
13. Brooks, R.: *Cambrian Intelligence*. MIT Press (1999)
14. Wurr, A.: Robotic team navigation in complex environments using stigmergic clues. Master’s thesis, Department of Computer Science, University of Manitoba, Winnipeg (2003)
15. Balch, T., Arkin, R.C.: Behavior-based formation control for multi-robot teams. *IEEE Trans. On Robotics and Automation* **14** (1998)