Robust Anticipatory Stigmergic Collision Avoidance in Multi-Agent Systems

(Extended Abstract)

Friedrich Burkhard von der Osten
Department of Computing & Information Systems
The University of Melbourne
Victoria 3010, Australia
fvon@student.unimelb.edu.au

Michael Kirley
Department of Computing & Information Systems
The University of Melbourne
Victoria 3010, Australia
mkirley@unimelb.edu.au

Tim Miller
Department of Computing & Information Systems
The University of Melbourne
Victoria 3010, Australia
tmiller@unimelb.edu.au

ABSTRACT

Anticipatory Stigmergic Collision Avoidance (ASCA) is a reciprocal collision avoidance model in which agents interact by indicating intended future paths to other agents in their local environment. Other agents can then incorporate that information to guide their path planning decisions. Additionally, constraints such as restricted movements and noise are incorporated. Our evaluation showed that ASCA is consistently more robust in noisy environments in which transmitted information can be lost or degraded.

Categories and Subject Descriptors

Computing methodologies [**Control methods**]: Motion path planning

General Terms

Algorithms, Performance, Reliability, Experimentation

Keywords

collision avoidance, stigmergy, simulation, cooperation

1. INTRODUCTION

Many reactive planning solutions assume that moving objects are *passive*, it is assumed that the objects will continue moving in a somewhat predictable manner. In multi-agent environments, such an assumption may not be valid, because other agents will also be reactively planning to avoid collisions. For example, unmanned aerial vehicles (UAVs) or unmanned underwater vehicles (UUVs) have stricter movement, localisation and communication constraints than the agents usually used in existing implementations of collision avoidance models.

In this paper, we present an approach to collision avoidance, *Anticipatory Stigmergic Collision Avoidance* (ASCA) to create a robust multi-agent collision avoidance method that is less reliant on accurate sensor information about the environment, and is scalable, as it is decentralised. Existing collision avoidance models deal with either noise [2] or movement restrictions [1], however these models do not combine these characteristics into a single solution. An important feature of our model is the ability to deal with movement restrictions and noise in a combined solution.

Unlike standard stigmergy-based models, in which agents deposit pheromones to inform other agents what actions they have taken [4], anticipatory stigmergy deposits pheromones to inform others of intended future actions [3].

We evaluate our model by comparing it to a state-of-the-art collision avoidance model for multi-agent systems, Reciprocal Velocity Obstacles (RVO) [5] and its extensions. The results show that ASCA outperforms RVO in the presence of noise. In environments with perfect information RVO outperforms ASCA.

2. ANTICIPATORY STIGMERGY COLLISION AVOIDANCE

The Anticipatory Stigmergy Collision Avoidance (ASCA) model makes two assumptions not typically made in existing collision avoidance work: a) agents have restricted movement (change of direction and speed) and b) perception of the local environment is noisy. Unlike standard stigmergy-based models, in which agents leave interpretable information on the trail passed, anticipatory stigmergy agents deposit information (i.e. pheromones) on *planned* trails to indicate their intended paths. A high level overview is given in Algorithm 1.

In each time step an ASCA agent has to decide on a path. If the optimal path p towards the goal is free up to a certain lookahead distance, an agent takes it immediately. If the path is not free, the agent evaluates (samples) a number of alternative paths C towards its goal.

A path is evaluated by the sum of pheromones other agents have deposited on it. One of the evaluated paths is selected probabilistically, with the probability of a given path being determined by the inversely proportionate sum of pheromones. This is due to the fact that higher sums of pheromones indicate congested areas and thus a lower likelihood of choosing that path.

As a last step, pheromones are deployed on the selected path j. The amount of pheromones deposited decreases the further the cell is from the agent, thus modelling the fact

Appears in: *Alessio Lomuscio, Paul Scerri, Ana Bazzan, and Michael Huhns (eds.), Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014), May 5-9, 2014, Paris, France.*

Algorithm 1 The ASCA algorithm for choosing a path

```
while agent has not reached goal do
    calculate optimal path p
    if p is free (no pheromones) then
        choose p
        update current position, velocity
        deposit pheromones on p
    else
        sample paths C around p
        weigh pheromones on all paths in C
        choose path j ∈ C probabilistically
        update current position, velocity
        deposit pheromones on j
    end if
    remove pheromones from previous path
end while
```

that closer cells will be arrived at sooner, and are more likely to be actually used by the pheromone-depositing agents.

If the selected path is not the original path, the pheromones are removed from the original path. This mechanism replaces the typical pheromone decay used in stigmergy models.

As an example in a real-world application, agents could broadcast their virtual pheromones on GPS coordinates (or by means of alternative localisation methods) in a limited radius. Such a broadcast would not require any message exchange and would be limited locally. An agent would only need to broadcast its own intended path, which requires very little communication bandwidth. If an agent receives a broadcast, it integrates the received intended path into its worldview (local data structure). Thus there is no need for a global data structure (e.g. maps) to be stored or shared.

3. EXPERIMENTAL RESULTS

In experiments, ASCA has been evaluated with the goal of minimising collisions, while also efficiently arriving at its target. Additionally, a comparison to the state-of-the-art method in collision avoidance, Reciprocal Velocity Obstacles (RVO) [5], is conducted to demonstrate ASCA's abilities in application. For a fair comparison the RVO model is subject to the same constraints as the ASCA model.

Experiments have been conducted over different numbers of agents on the same environment size (effectively increasing agent density) and various *noise levels*, which aim to simulate unreliable communication channels or sensoring. The two main measures of interest are the number of collisions occurring and the efficiency of an agent reaching its goal, defined as the percentage of the optimal distance over the actual distance travelled. Those measures have been evaluated on a number of different scenarios, i.e. geometrical setups of agents and their goals.

Figure 1 shows 3D plots for one exemplary scenario, outlining the number of collisions under increasing agent density and noise (plots for other scenarios are similar).

The results from the evaluation demonstrate that ASCA performs robustly in a noisy environment when compared to RVO. The results support our hypothesis for applications such as multi-UUV environments, where perception is difficult and the vehicles have large turning circles and cannot stop quickly, existing collision avoidance models cannot be used directly.

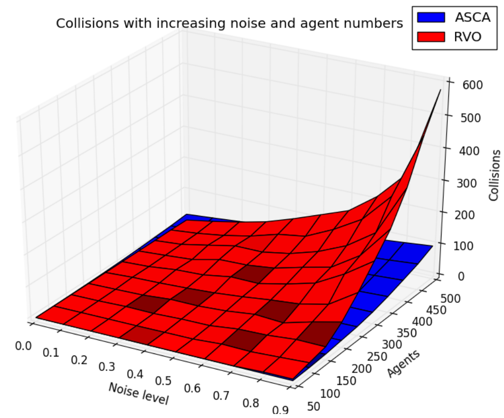

Figure 1: Collisions with increasing agent numbers and noise

4. CONCLUSIONS AND FUTURE WORK

In this paper, we have presented a simple and effective robust collision avoidance model. In ASCA, an agents' decision process is independent of other agents' decisions, while still taking other agents' behaviour into account. We assume stricter constraints on how agents can move, making it suitable for applications such as aerial or aquatic vehicles. Compared with RVO, ASCA results in less collisions in noisy environments, while offering similar efficiency.

In future work, we will extend the model to a continuous 3-dimensional space, and will include sampling of different velocities, which we believe can additionally improve the performance in terms of reducing collisions. Additionally, it would be interesting to investigate different scenarios, such as following a moving target.

5. REFERENCES

[1] Javier Alonso-Mora, Andreas Breitenmoser, Martin Rufli, Paul Beardsley, and Roland Siegwart. Optimal reciprocal collision avoidance for multiple non-holonomic robots. In *Distributed Autonomous Robotic Systems*, pages 203–216. Springer, 2013.

[2] D. Hennes, D. Claes, W. Meeussen, and K. Tuyls. Multi-robot collision avoidance with localization uncertainty. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, volume 1, pages 147–154. International Foundation for Autonomous Agents and Multiagent Systems.

[3] T. Ito, R. Kanamori, J. Takahashi, I. Maestre, and E. De La Hoz. The comparison of stigmergy strategies for decentralized traffic congestion control: Preliminary results. In *PRICAI 2012: Trends in Artificial Intelligence*, pages 146–156. Springer, 2012.

[4] G. Theraulaz and E. Bonabeau. A brief history of stigmergy. *Artificial life*, 5(2):97–116, 1999.

[5] J. Van den Berg, M. Lin, and D. Manocha. Reciprocal velocity obstacles for real-time multi-agent navigation. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 1928–1935. IEEE.