

A probabilistic movement model for shortest path formation in virtual ant-like agents

Colin Chibaya
Department of Computer Science
Rhodes University
Grahamstown, South Africa
c.chibaya@ru.ac.za

Shaun Bangay
Department of Computer Science
Rhodes University
Grahamstown, South Africa
s.bangay@ru.ac.za

ABSTRACT

We propose a probabilistic movement model for controlling ant-like agents foraging between two points. Such agents are all identical, simple, autonomous and can only communicate indirectly through the environment. These agents secrete two types of pheromone, one to mark trails towards the goal and another to mark trails back to the starting point. Three pheromone perception strategies are proposed (Strategy A, B and C). Agents that use strategy A perceive the desirability of a neighbouring location as the difference between levels of attractive and repulsive pheromone in that location. With strategy B, agents perceive the desirability of a location as the quotient of levels of attractive and repulsive pheromone. Agents using strategy C determine the product of the levels of attractive pheromone with the complement of levels of repulsive pheromone. We conduct experiments to confirm directionality as emergent property of trails formed by agents that use each strategy. In addition, we compare path formation speed and the quality of the formed path under changes in the environment. We also investigate each strategy's robustness in environments that contain obstacles. Finally, we investigate how adaptive each strategy is when obstacles are eventually removed from the scene and find that the best strategy of these three is strategy A. Such a strategy provides useful guidelines to researchers in further applications of swarm intelligence metaphors for complex problem solving.

Categories and Subject Descriptors

1.3.2 [Computer Graphics]: Graphics systems-Distributed/network graphics; 1.3.6 [Computer Graphics]: Methodology and Techniques-interaction techniques; 1.6.5 [Simulation and Modeling]: Model development

Keywords

Emergent behaviour, search and return pheromone, dissipa-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SAICSIT 2007, 2 - 3 October 2007, Fish River Sun, SunShine Coast, South Africa

Copyright 2007 ACM 978-1-59593-775-9/07/0010 ...\$5.00.

tion, directionality, speed, quality and recruitment.

1. INTRODUCTION

Ants in nature are fascinating creatures - not so much because they are particularly intelligent on their own, but because as a colony they display compelling swarm intelligence. The collective effect of individuals wandering in search of *food* results in the formation of an established path between the *ant hole* and a *food source*. This path is an emergent behaviour of the colony.

We investigate collections of agents modelled on ants, investigating strategies for efficiently achieving shortest path formation as the resulting emergent behaviour. Such collective behaviour absorbs many individual agent failures and disabilities. Our agents are all identical, autonomous and only communicate through stigmergy (an indirect, non symbolic form of communication mediated by the environment). We support two types of pheromone with levels that co-exist independently: one to mark trails towards the goal and another to mark trails back to the starting point. Pheromone levels at each location are increased on every agent visit. Our agents have neither *a priori* knowledge of their global location nor awareness of the absolute position of the goal.

1.1 Problem statement

A number of factors can influence the rate at which path formation occurs, and how robust it is under changes in the environment. We investigate the following:

- Pheromone sensitivity: we propose three strategies for determining agent sensitivity to co-existing pheromone levels.
- Pheromone dissipation: we evaluate the effects of environmental factors such as pheromone evaporation and diffusion to model performance.
- Robustness and adaptability: we evaluate the performance of each strategy in environments containing obstacles.

1.2 Overview

The rest of the paper proceeds as follows: In Section 2, we present related work concentrating mainly on foraging algorithms and swarm control systems. Section 3 presents the algorithm and strategies that we use for controlling agent movements. Pheromone dissipation controls are also presented in this section. An in-depth description of how experiments are conducted is presented in section 4 and the

results achieved thereafter are presented and discussed in section 5. We conclude the paper in section 6 highlighting contributions and future directions of our work.

2. RELATED WORK

The basis for our agent's movement policy is the various Ant Colony Optimization algorithms that have been suggested for controlling agent mobility towards a goal [3, 13, 11, 2]. The common feature of these algorithms is the secretion of pheromone either by the agents themselves or by the search targets. Agents are very simple, unintelligent, identical, autonomous and can only communicate indirectly via the environment.

Control models where agents use a single type of pheromone [4, 3, 5, 6, 10] are common. In these, pheromone is secreted as agents forage, essentially marking trails that other agents may follow. In other cases, the search target itself secretes pheromone that can be used to create an orientation gradient [6, 5]. Return journeys, if required, make use of other *ad hoc* mechanisms. [4] proposes a model where agents recognize their return direction by the height and angle of the sun. However, this would require agents to have some memory for recording the path followed.

Multiple pheromone gradients can also be created. A good example is in the work of [1] on nanorobots communication techniques where the cancer cell targets secrete *attraction* pheromone while obstacles secrete *fear* pheromone. Work of [9] suggests two pheromones used as signals that alert agents about the availability of a goal in the environment. Such indicators would allow agents to either exhaustively search for the goal or wait until the goal is introduced again.

We investigate agent movement controls using a two pheromone laying model with agents that neither have global information nor look ahead capabilities. Existing two pheromone laying models differ in that they require global visibility of where the targets are, and determine agent's distance from the desired goal [12]. Such distance is critical when agents select the next location to visit and when deciding on the amount of desired pheromone levels to be secreted thereafter (zero to MAX_PHEROMONE) [12]. Other two pheromone laying models use look-ahead mechanisms for agents orientation [11].

The model we suggest works on a 2D square space similar to a cellular automaton. However, the desirability of each discrete location of the space is not influenced by the desirability of any of its neighbouring locations. In addition, our agents do not change their state and behavioural tendency after every update as suggested in cellular automation models [8]. Instead, each agent is autonomous and can only secrete specific pheromone that is used to update the environment.

Various strategies with which agents perceive the desirability of locations and their sensitivity to co-existing pheromone levels can be suggested. We propose, investigate and evaluate the performance of three of these.

3. MOBILE AGENT MODEL

We developed a "*foraging ants simulator*" and assume a model that is motivated by capabilities of ants in nature. Our agents are simulated as moving at a constant speed, secreting the same amount of pheromone and executing the same routine. Variation of agents speed was not an ex-

Algorithm 1 Control algorithm for ant-like agents.

```

mode ← Search
while (true)
  direction ← None
  foreach location L around agent
    if L contains no obstacles
       $R_L \leftarrow$  return pheromone at L
       $S_L \leftarrow$  search pheromone at L
      if mode = Search
         $P_L \leftarrow P(S_L, R_L)$ 
      else
         $P_L \leftarrow P(R_L, S_L)$ 
      else
         $P_L \leftarrow 0$ 
  choose direction probabilistically based on  $P_L$ 
  if mode = Search
    deposit return pheromone
  else
    deposit search pheromone
  if at goal
    mode ← Return
  if at starting point
    mode ← Search

```

periment variable in this study. No explicit coordination amongst agents occurs. All behaviour results from the interaction of agents with the pheromone in the environment. We also assume a dynamic environment where pheromone levels change in every agent movement step due to evaporation and diffusion.

3.1 Agent movement controls

The routine that we use for controlling the movements of our ant-like agents is as shown in Algorithm 1. Each agent decides its next destination based on the levels of both return (R_L) and search (S_L) pheromone in reachable neighbouring locations L . We consider a location to be reachable when it is free from obstacles (objects that prevent agents from accessing some parts of the environment). Upon deployment all agents are in the default search mode. They secrete return pheromone at their current location resulting in the formation of return pheromone trails which lead returning agents towards the starting point. They are likely to move towards locations with high levels of search pheromone while simultaneously being repelled by return pheromone. However, this tropism is reversed when the agent reaches its goal, and needs to return to the starting point. Returning agents secrete search pheromone resulting in the formation of trails which guide searching agents towards the goal.

We propose three strategies with which agents perceive the desirability of a neighbouring location based on the relative levels of co-existing pheromone. In each case, these strategies are such that agents favour high quantities of one pheromone, while penalizing high quantities of the other. The expressions we present below are of our own making although they are inspired by related work in [8, 10]. Let N be the set of reachable neighbouring locations, where the parameters τ_n and η_n respectively stand for the levels of attractive and repulsive pheromone at neighbouring location $n \in N$. The desirability P_n , of a neighbouring location $n \in N$ can be expressed as:

strategy A: The difference between levels of attractive and repulsive pheromone relative to the smallest level as shown in (2).

$$Q_{A,n} \leftarrow \left(\frac{\tau_n}{\sum_{k \in N} \tau_k} \right) - \left(\frac{\eta_n}{\sum_{k \in N} \eta_k} \right) \quad (1)$$

$$P_{A,n} \leftarrow \lambda(Q_{A,n} - P_{A,min}) + \frac{1 - \lambda}{\|N\|} \quad (2)$$

strategy B: The quotient of levels of attractive and repulsive pheromone as shown in (3).

$$P_{B,n} \leftarrow \lambda \left(\frac{\left(\frac{\tau_n}{\sum_{k \in N} \tau_k} \right)}{\left(\frac{\eta_n}{\sum_{k \in N} \eta_k} \right)} \right) + \frac{1 - \lambda}{\|N\|} \quad (3)$$

strategy C: The product of levels of attractive with the complement of levels of repulsive pheromone as shown in (4).

$$P_{C,n} \leftarrow \lambda \left(\frac{\tau_n}{\sum_{k \in N} \tau_k} \right) \left(1 - \frac{\eta_n}{\sum_{k \in N} \eta_k} \right) + \frac{1 - \lambda}{\|N\|} \quad (4)$$

In all cases, $\|N\|$ refers to the cardinality (number of elements) of N and $\lambda \in [0, 1]$ is the degree of randomness. Pheromone levels are initialized to a very small $\delta = 0.00001$ to avoid any divisions by zero.

Once the desirability measures for all the reachable neighbouring locations $P_{L,n}$ have been determined, these measures are expressed as end to end intervals. Highly desirable neighbouring locations will correspond to bigger intervals. A point is randomly selected and the interval into which it falls corresponds to the agent's next destination.

Initially agents discover the goal by random guess (strategy D), having created return trails. On the return trip they follow return trails, secreting search pheromone which guides other agents towards the goal. Both search and return pheromone trails form, creating gradients towards the goal and the starting point respectively. Shorter trails are reinforced more often than longer trails [4, 3] making them highly desirable. Once the shortest path is slightly more attractive, it is likely to get stronger still due to mass recruitment. Adding environmental features such as pheromone evaporation would dissipate the weak trails while diffusion reinforces and widens the shortest path.

3.2 Pheromone dissipation controls

Our agents forage in an environment consisting of the starting point, the goal and pheromone trails. The goal and the starting point are placed at fixed locations as depicted in Figure 1.

Pheromone dissipation is applied through evaporation and diffusion. Pheromone diffusion reinforces levels in neighbouring locations while evaporation reduce levels at the current location. We use Algorithm 2 to effect pheromone dissipation. The parameters $\rho \in [0, 1]$ and $\gamma \in [0, 1]$ respectively indicate evaporation and diffusion coefficients.

In each movement step, both return and search pheromone levels on all locations of the environment are updated. There are two reasons why pheromone levels may decrease: either because pheromone has been evaporated or it has been diffused to neighbouring locations. However pheromone levels can increase as a result of secretions from agents currently visiting that location. It can also be increased by pheromone

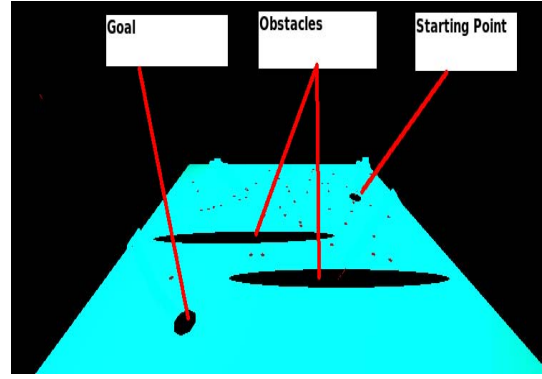


Figure 1: Simulation Environment

diffused from neighbouring locations. Our model supports diffusion towards locations with lower pheromone levels.

4. EXPERIMENTAL SETUP

We conduct three simulation experiments aimed at investigating the efficiency and robustness of each pheromone perception strategy. The first experiment investigates directionality as emergent property of laid pheromone trails. The second experiment measures the relative performance (path formation speed and the quality of the path formed) of the three pheromone perception strategies described in section 3.1. Finally we assess the robustness and adaptability of each pheromone perception strategy in environments that consist of obstacles.

4.1 Directionality

In this first experiment, we preset a path along the line segment that joins the goal and the starting point. Search pheromone in locations along this line increases linearly from the goal to the starting point. Return pheromone increases from the starting point towards the goal. Ten thousand agents in each mode are successively recruited into the formed path by placing them at positions close to the path. The frequency with which each agent orients towards the desired target when it intersects the path is recorded. An agent is classified as lost when it crosses over or retreats away from the path. In all cases, pheromone dissipation and further agent secretions are initially disabled.

4.2 Path formation

In the second experiment, we assess the performance of each pheromone perception strategy. Two metrics are of particular interest:

Speed: This measures the influence of both attractive and repulsive pheromone. We measure time in the simulation using the number of iterations executed. To measure path formation speed, we firstly determine the number of iterations executed until the first agent arrives at the goal. The model speed towards the goal is then determined using the iteration gap between the first agent's arrival time until ten more agents arrive at the goal. From observation, ten agents are sufficiently many to cause model convergence. We also record the iterations executed from the time when the goal is first discovered until the time when the first

Algorithm 2 Pheromone update rules.

```
foreach location L in environment
  foreach pheromone at L
     $P_{L,i} \leftarrow$ pheromone levels at L in iteration i
     $S_{L,i} \leftarrow$ secreted pheromone levels at L in iteration i
    foreach neighbour n around L
      if  $P_{n,i} > P_{L,i}$  then
         $D_{n,i} \leftarrow \gamma P_{n,i}$ (pheromone diffused from neighbour n in iteration i)
      else
         $D_{n,i} \leftarrow 0$ 
     $P_{L,i+1} \leftarrow (1 - \gamma)(1 - \rho) \cdot P_{L,i} + \sum_n D_{n,i} + S_{L,i}$ 
```

agent return back to the starting point. The time gap between this measure until ten more agents return to the starting point indicates the model's return speed. Smaller measures in both directions indicate few iterations before the target is found, corresponding to better performance.

Quality: This indicates the tendency of agents to follow the path, as opposed to reaching the goal by random guess. We measure trail quality using the number of successful trips of agents towards the target in a set time period, in this case 1000 iterations. The higher the quality of the path, (the greater the number of agents which found the target) the shorter it is. We also measure this metric using the intervals between successive agents arriving at the target. Shorter intervals indicate the use of a well defined path.

For each test, ten simulations are conducted. The time values (measured in iterations) are averaged over the results of all ten simulations, giving a centrally placed time value upon which the two metrics are calculated. For the parameters:

- λ , which determines the degree of randomness [7] is assigned the following values: 1 (only pheromone effects), 0.5 (a combination of both pheromone and random wandering effects) and 0 (complete random wandering)
- evaporation coefficients used are: $\rho = 0.0\%$, (no evaporation), $\rho = 0.01\%$ (decay of pheromone concentration by 0.01% in each movement step due to evaporation), $\rho = 0.05\%$, $\rho = 0.1\%$ and $\rho = 1\%$.
- diffusion coefficients used are: $\gamma = 0.0\%$ (no diffusion), $\gamma = 0.01\%$ (decay of pheromone concentration by 0.01% in each movement step due to diffusion towards neighbouring locations), $\gamma = 0.05\%$, $\gamma = 0.1\%$ and $\gamma = 1\%$.

4.3 Robustness and adaptability

The third experiment investigates each strategy's robustness and adaptability. We measure robustness by investigating path formation speed and path quality in environments that consist of obstacles. We progressively increase our scene complexity by introducing these obstacles one at a time to eventually reach the scenario depicted in Figure 1. These obstacles are placed at random positions between the goal and the starting point. Path formation speed and path quality are then measured in the same way they are measured when agents forage in environments that are free from obstacles (see section 4.2). We say a strategy is robust when a

Metric	Measure
Directionality	Agent orientation on recruitment
Path formation speed	Sensitivity to pheromone levels
Path quality	Path following behaviour
Robustness	Path formation with obstacles
adaptability	Re-organization speed

Table 1: List of measured metrics

fairly quality path is rapidly formed in an environment with obstacles.

On the other hand, adaptability refers to a situation where an even shorter path can be re-established after one or all of the obstacles are removed from an environment where a path has already formed. We measure adaptability using the frequency with which agents pass through the location where an obstacle was placed. These frequencies indicate the speed with which the model re-organizes and high frequencies are desirable.

We present in Table 1, the five metrics that we measure for each of the three strategies.

5. RESULTS AND DISCUSSION

5.1 Directionality

Table 2 presents the results achieved from the experiment testing agent orientation on recruitment (experiment 1). The values indicate the number of agents that achieved each possible outcome when ten thousand agents were recruited into an already established path.

Agents that use strategy A achieve desired orientation on intersecting with the set path. Table 2 indicates that 95% of agents in search mode successfully orient themselves toward the goal. Of these, 99% manage to follow the path until they reach the goal. Of the agents that started heading toward the starting point, 66% eventually redirected themselves towards the goal. Only 36% of lost agents wandered into the goal. Overall 97% of agents in search mode successfully found the goal. Returning agents exhibit identical behaviour.

Similarly, agents that use strategy B achieved desired orientation, although performance is lower than that for strategy A. We observe that 88% of searching agents initially achieve desired orientation, with 90% of these following the path until they reach the goal. A larger proportion of agents fail to reach the target after starting off in the correct direction. Contrary to strategy A, performance of returning agents is significantly lower than that of searching agents.

	A			B			C		
	Goal	Nest	Lost	Goal	Nest	Lost	Goal	Nest	Lost
Searching	9508 ± 110	331 ± 39	161 ± 19	8826 ± 184	923 ± 232	251 ± 32	1911 ± 843	1101 ± 428	6988 ± 771
Got to target	9441 ± 93	219 ± 18	59 ± 11	7956 ± 193	491 ± 53	79 ± 11	101 ± 41	129 ± 38	379 ± 119
Returning	302 ± 19	9577 ± 83	121 ± 17	786 ± 183	7991 ± 149	1223 ± 121	1219 ± 511	1304 ± 392	7477 ± 649
Got to target	207 ± 35	9499 ± 97	49 ± 9	547 ± 131	7033 ± 162	844 ± 101	211 ± 78	183 ± 69	411 ± 167

Table 2: Agent orientation on recruitment

P_1	attr:1.0000 rep:1.0000	P_2	attr:1.0000 rep:2.0000	P_3	attr:1.0000 rep:3.0000
P_8	attr:1.0000 rep:8.0000	Agent		P_4	attr:1.0000 rep:4.0000
P_7	attr:1.0000 rep:7.0000	P_6	attr:1.0000 rep:6.0000	P_5	attr:1.0000 rep:5.0000

Table 3: Example layout of attractive and repulsive pheromone levels at each neighbouring location.

Numbers of lost agents are much higher on the return trip, indicating agents that do not react to the path.

Unlike strategies A and B, only 19% of agents deployed in search mode and 13% of agents deployed in return mode using strategy C achieved the desired orientation on intersecting the path. Of these, 5% and 14% respectively follow the path until they reach the target. Re-orientation of initially disoriented agents occurs with only an 11% level of success for agents in search mode and 17% for agents in return mode.

Summary

Agents using strategies A and B achieve desired orientation and exhibit path following behaviour. However agent’s sensitivity to attractive and repulsive pheromone differs between the two strategies. Strategy A penalizes pheromone levels in location $P_{L,min}$ and assign to that location, the minimum possible desirability level ($\delta = 0.00001$). On the other hand, strategy B accentuates desirability levels towards locations where attractive and repulsive pheromone ratios are good in favour of attractive pheromone levels. As evidence, consider a scenario in Table 3 depicting an agent and the pheromone levels at each of its eight neighbours. The top index represents attractive pheromone while the bottom index represents repulsive pheromone at that location. Figure 2a compares the desirability of each of these locations as they are perceived by agents in search mode using each strategy. Conversely, Figure 2b compares the perception of agents in return mode to the desirability of the same locations. From this example we can see that:

- Strategy A is consistent. Levels of attractive and repulsive pheromone bear the same effects to both searching and returning agents. Consequently the absolute gradients of the curves plotted in Figures 2a and 2b are equal.
- Strategy B is inconsistent. Directionality is weak in locations where repulsive pheromone is higher than attractive pheromone as evidenced by the flatter gradient observed in Figure 2a. Agents that are recruited near the starting point (where repulsive pheromone is likely to be higher than attractive pheromone) are less

likely to correctly orient themselves, and even if they do, will be more likely to fall off the path. In addition, this strategy does not tail off to zero, allowing even unlikely cases to have some possibility of being selected - hence more chance of random wander. Thus strategy B performs worse than strategy A.

- In strategy C, agents are weakly sensitive to pheromone levels when repulsive pheromone is higher than attractive as indicated by the shallow gradient of the line segment in Figure 2a. In cases where attractive pheromone levels are high (as in the return trip example in Figure 2b) the example returns the same figures as strategy B. The simulations indicate that the directional performance of this strategy is indistinguishable from a random selection of any one of the eight neighbouring locations.

5.2 Path formation

Table 4 presents simulation results investigating path formation speed and quality. Values measured when pheromone dissipation and random wandering are ignored serve as reference levels. We assess the effects of pheromone dissipation and changes in values of λ by comparing results to these cases.

5.2.1 Path formation speed

In the reference case for strategy A, the first ten agents discover the goal in an average of 777 iterations with a standard deviation of 25 iterations. Return trips are much faster (439 iterations with standard deviation of 14 iterations) because return trails are formed when agents are still searching for the goal. These return trails give directional clues as to where the starting point is. This strategy achieves best results with parameters $\lambda = 1$, $\rho = 0.0\%$ and $\gamma = 0.05\%$.

Low levels of evaporation eliminate suboptimal trails before they influence agent’s movement. However as evaporation increases further, speed drops again. This is because high evaporation levels deplete trails before they are used. Cases where search trips are faster than return trips can be attributed to the searching agents being centrally deployed, secreting repulsive return pheromone which speeds up agents exploration away from the starting point. Increased diffusion further speeds them on their outward journey.

Returning agents do not commence their return journey at the same time. This results in an inflation in the values reported for return speed, since some agents may have returned before others have even found the goal for the first time. Additional effects of agents leaving the goal at different time intervals include the lack of search pheromone around the goal, which would repel agents, and depletion of trails due to evaporation.

	Evaporation →	0.0%			0.01%			0.05%			0.1%			1%		
Diffusion ↓	$\lambda = 1.$	A	B	C	A	B	C	A	B	C	A	B	C	A	B	C
0.0%	Speed towards goal	777	689	2568	407	489	3560	320	429	-	229	491	-	436	495	-
	Speed on return trip	439	402	-	442	857	-	458	904	-	666	962	-	1254	965	-
	Quality towards goal	35	16	4	31	21	6	30	23	-	27	21	-	24	15	-
	Quality on return trip	39	21	-	35	15	1	33	10	-	20	13	-	8	12	-
0.01%	Speed towards goal	366	557	-	362	397	-	327	338	-	359	574	-	386	491	-
	Speed on return trip	334	517	-	373	837	-	411	820	-	600	849	-	732	891	-
	Quality towards goal	40	23	2	38	27	3	35	21	-	29	19	-	26	26	-
	Quality on return trip	46	20	-	51	15	-	35	46	-	33	12	-	14	14	-
0.05%	Speed towards goal	326	545	3764	319	390	-	293	324	-	375	680	-	377	552	-
	Speed on return trip	321	804	-	523	803	-	441	744	-	762	784	-	1271	844	-
	Quality towards goal	69	17	5	51	20	6	54	24	-	42	18	-	19	13	-
	Quality on return trip	73	12	2	37	11	3	56	11	-	25	13	-	6	11	-
0.1%	Speed towards goal	321	600	2496	310	483	-	311	514	-	309	575	-	335	522	-
	Speed on return trip	438	884	-	581	867	-	50-2	733	-	473	727	-	987	623	-
	Quality towards goal	76	32	3	71	24	3	68	11	-	51	24	-	26	26	-
	Quality on return trip	53	13	2	43	10	1	52	12	-	53	17	-	10	19	-
1.0%	Speed towards goal	290	2335	1113	423	1344	-	624	740	-	656	1107	-	375	796	-
	Speed on return trip	717	3035	-	655	1881	-	613	1150	-	803	1389	-	448	1893	-
	Quality towards goal	30	2	9	23	6	6	28	11	-	25	8	-	51	14	-
	Quality on return trip	14	1	2	17	2	1	23	7	-	19	6	-	43	6	-
	$\lambda = \frac{1}{2}.$	A	B	C	A	B	C	A	B	C	A	B	C	A	B	C
0.0%	Speed towards goal	437	610	-	549	585	-	306	522	-	420	353	-	395	705	-
	Speed on return trip	494	1206	-	428	1117	-	650	778	-	580	476	-	666	1096	-
	Quality towards goal	20	16	-	25	13	-	33	19	-	25	25	-	23	11	-
	Quality on return trip	11	8	-	17	8	-	14	13	-	16	16	-	15	9	-
0.01%	Speed towards goal	561	598	-	605	704	-	324	468	-	314	710	-	594	567	-
	Speed on return trip	768	932	-	843	648	-	700	904	-	638	1291	-	695	805	-
	Quality towards goal	23	20	-	20	19	-	20	25	-	25	12	-	26	26	-
	Quality on return trip	13	12	-	11	11	-	13	12	-	14	8	-	12	15	-
0.05%	Speed towards goal	444	513	-	399	772	-	283	478	-	443	387	-	223	370	-
	Speed on return trip	752	1331	-	711	1239	-	738	580	-	752	502	-	538	454	-
	Quality towards goal	20	17	-	31	16	-	30	19	-	22	24	-	37	29	-
	Quality on return trip	11	8	-	15	5	-	17	14	-	11	16	-	17	28	-
0.1%	Speed towards goal	484	638	-	206	594	-	325	824	-	322	789	-	424	808	-
	Speed on return trip	700	624	-	561	969	-	514	993	-	697	1099	-	484	927	-
	Quality towards goal	31	19	-	37	14	-	29	18	-	27	16	-	35	18	-
	Quality on return trip	16	12	-	16	10	-	15	10	-	15	5	-	19	14	-
1.0%	Speed towards goal	306	728	-	516	625	-	757	1796	-	333	551	-	318	533	-
	Speed on return trip	863	930	-	1023	1069	-	1560	2996	-	887	1212	-	792	1434	-
	Quality towards goal	13	14	-	13	13	-	10	6	-	15	15	-	14	12	-
	Quality on return trip	10	13	-	10	9	-	7	3	-	10	7	-	11	7	-
	$\lambda = 0.$	A	B	C												
0.0%	Speed towards goal	1939	1939	1939												
	Speed on return trip	1793	1793	1793												
	Quality towards goal	5	5	5												
	Quality on return trip	5	5	5												

Table 4: Path formation speed and path quality records using each strategy.

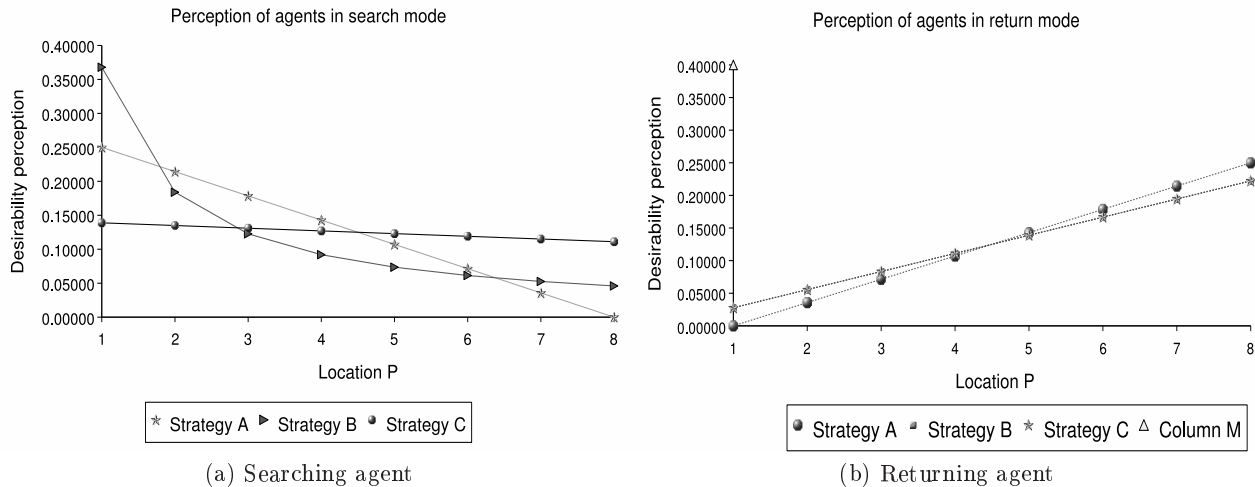


Figure 2: Agent’s perception of the desirability of locations in Table 3

We observe that diffusion improves model speed. This is because when diffused:

- repulsive pheromone speeds up exploration away from the starting point, and
- attractive pheromone creates a gradient towards the goal.

However high diffusion causes pheromone to cover the entire area evenly reverting agents to random wandering. This is because, later in simulation time, the diffusion process described in Algorithm 2 allows diffusion of more pheromone levels than an agent can secrete, resulting in paths being overwritten. Combining the effects of both evaporation and diffusion is more desirable because in addition to reinforcement through agent secretion, trails are strengthened and widened by diffusion while long and suboptimal trails get weaker still due to evaporation.

Changes in values of the parameter λ slows down the model except when evaporation coefficients are high. The reason for this is that an increase in randomness ($\lambda < 1$) in agents’ movements straightens paths when agents randomly cut corners. However, when evaporation is low, increasing randomness can derail agents.

In the reference case for strategy B, the first ten agents discover the goal in an average of 689 iterations with a standard deviation of 72 iterations. These results are consistent with data presented in Figure 2, confirming that the model is good during exploration stages. Late in simulation time, high evaporation gradually encourages the establishment of the shortest path with constant levels of both attractive and repulsive pheromone. Diffusion maintains the smoothness and widen the remaining trails while some degree of randomness ($\lambda = \frac{1}{2}$) helps to counter trail depletion.

Observation of the simulation reveals agents that use strategy C as successfully achieving exploration and goal discovery. However agents orientation is poor. We have shown in section 5.1 that agents using this strategy are less sensitive to repulsive pheromone. If they remain at a single spot for any length of time they secrete additional repulsive

pheromone which further decreases their directional sensitivity and become trapped. Any other agent that visits this location is also trapped.

Summary

Agents that use strategy A symmetrically perceive pheromone levels with equal importance regardless of their internal state or the time in simulation. However, late in simulation time, the levels of diffusing pheromone may exceed what agents can secrete causing agents to randomly wander. Agents that use strategy B perform best with high pheromone dissipation and some degree of randomness. Strategy C is worse than a random guess.

5.2.2 Quality of path

Measurements of path quality are also presented in Table 4. The values indicate the number of trips achieved within a period of 1000 iterations. Larger values correspond to better defined paths.

In the reference case for strategy A, an average of 35 ± 4 trips and 39 ± 5 trips are respectively achieved to and from the goal. As with path formation speed, return trails are of better quality because a gradient of return pheromone forms towards the starting point when agents are still searching for the goal. This strategy achieves best results with the same parameter values with which the best path formation speed is achieved. With these parameters, strategy A achieves 92% improvement in path quality compared to the reference case.

We observe similar trends in strategy B where the best path quality results are achieved with the same parameter values with which the best path formation speed is achieved. Path quality degrades late in simulation time due to drops in agents sensitivity to repulsive pheromone that reverts agents to random wandering. Figure 3a illustrates the trends in arrival times for typical sample paths produced by strategies A and B. The path formed by agents that use strategy A indicates that agents follow the path while the one formed by agents using strategy B dissipates.

When path formation is achieved through pheromone dissipation, agents using strategy C may find their targets. Agents that use this strategy have a high tendency of not

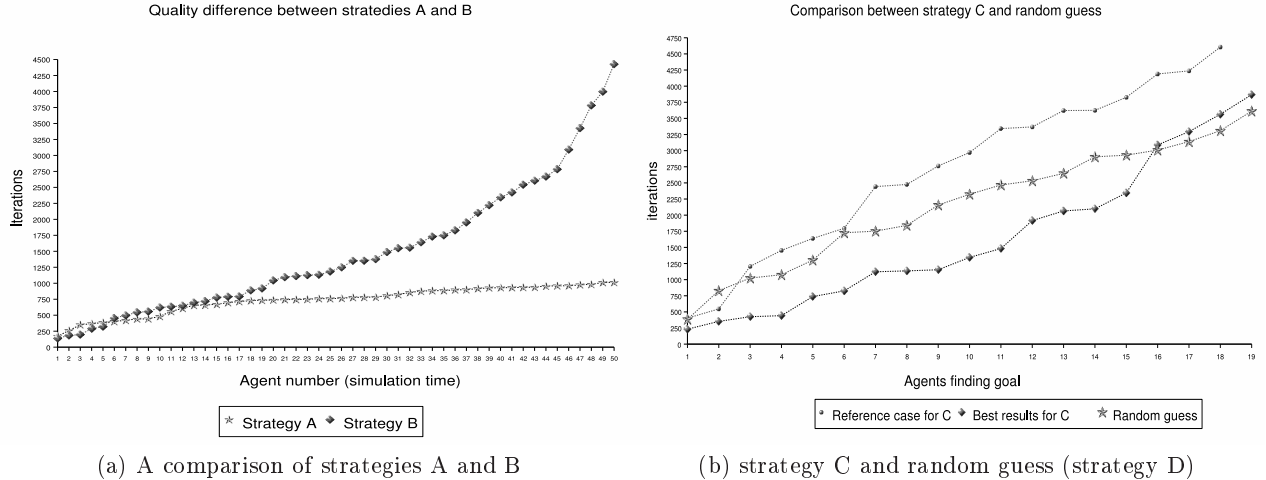


Figure 3: Quality measure using intervals of successive arrival

		A	B	C	D
No obstacles	Speed towards goal	326	370	1113	1939
	Speed on return trip	321	454	0	1793
	Quality towards goal	69	29	9	5
	Quality on return trip	73	28	2	5
One obstacle	Speed towards goal	583	2496	0	0
	Speed on return trip	802	2297	0	0
	Quality towards goal	20	7	0	3
	Quality on return trip	33	5	0	3
Two obstacles	Speed towards goal	590	3269	0	0
	Speed on return trip	960	3839	0	0
	Quality towards goal	17	4	0	1
	Quality on return trip	11	5	0	1

Table 5: Robust and fault tolerance test

getting to the desired target. For this reason, the quality of the path is low. In addition, high diffusion depletes the path when pheromone eventually covers the entire area. Figure 3b compares the trends in arrival times for typical sample paths produced by strategy C and random guess. It is indicated in Table 4 that the quality of the path formed by agents that use strategy C in the reference case is lower than the path formed by agents that use random guess. The reason is that, without pheromone dissipation, agents that use strategy C cannot follow path and are trapped on suboptimal solutions before they reach their targets.

Summary

We observe that the parameter values with which each strategy achieves best path formation speed are the same parameter values with which it achieves best path quality results.

5.3 Robustness and adaptability

The robustness measures for each pheromone perception strategy are presented in Table 5. The parameter values for which the best results were achieved in the previous experiments (see section 5.2) are used. A strategy is robust when a quality path can rapidly form in an environment consisting of obstacles.

With strategy A, path formation is achievable with the parameter values $\lambda = 1$, $\rho = 0.0\%$ and $\gamma = 0.05\%$ with which the best results were achieved in the previous experiments. Search speed drops as obstacles are introduced because agents are delayed when they are negotiating exploration paths around the obstacles. Return trips are much slower because return pheromone, which should guide agents back to the starting point, cannot be diffused from highly concentrated locations near the starting point through obstacles to locations that are on the other side of the obstacles (near the goal). When the first few agents discover the goal return trails close to the goal are still very weak because they have only been formed from pheromone secreted by these few agents. Later in simulation time, a path negotiated around the obstacle emerges depicting the scenario illustrated in Figure 4a and 4b. Pheromone dissipation through evaporation helps to eliminate suboptimal trails while diffusion straightens and widens the emergent path.

Strategy B achieves path formation slowly. The quality of the formed path is poor. The reason why the results are poor is that the repulsive effects of return pheromone drops and agents revert to random wandering before the target is found. High pheromone dissipation is detrimental as the path length increases.

We established earlier in section 5.2.1 that strategy C can not achieve reliable path formation. It is worse still in an environment containing obstacles. The strategy performs worse than a random guess (strategy D) as shown in Table 5.

Summary

Strategy A satisfies our criteria for a robust and adaptive model. Rapid formation of a high quality path is achieved in an environment consisting of obstacles. Strategies B and C do not meet the criteria of robust models. After path formation with agents that use strategy A, an even shorter path is established after obstacles are removed from the environment as evidence of the strategy's adaptability. Agents that use strategy B improve in performance immediately after obstacles are removed.

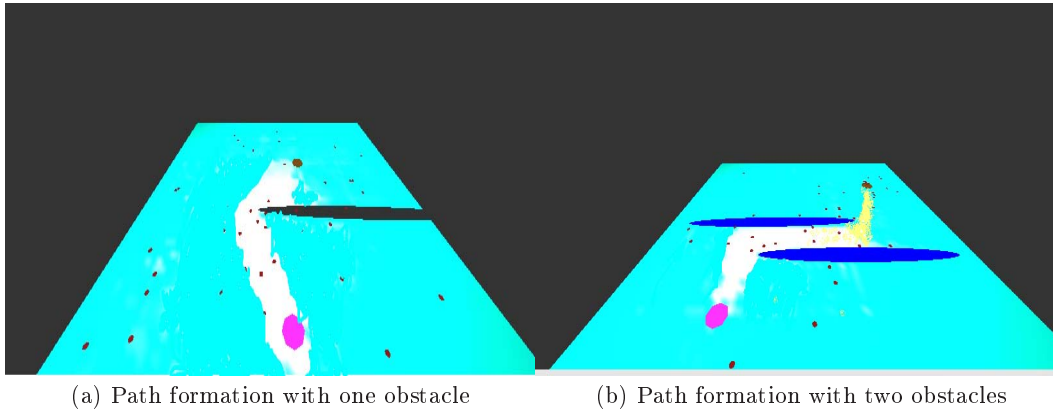


Figure 4: Path formation in environments with obstacles

6. CONCLUSION

We proposed a probabilistic movement model for controlling ant-like agents based on the use of two pheromones. Three pheromone perception strategies have been devised and are investigated. Experiments that evaluate and compare each strategy yield the following conclusions:

- Strategy A exhibits the best performance in all categories measured.
- Performance in the presence of obstacles clearly reveals weaknesses in strategies B and C.

The contributions of this work are:

- We present a novel agent control model using two pheromone types with agents that neither have global information nor look ahead capabilities.
- We devise and evaluate three plausible strategies for combining a pair of pheromone gradients into a direction vector, and assess the merits of each.
- We evaluate ways in which diffusion and evaporation effects can influence the emergence of a path.
- We devise new ways to measure and compare the performance of agent based path formation strategies.

Further research in formalizing the nature of the emergent behaviour as a function of agent density, pheromone dissipation parameters, degree of randomness and terrain complexity is underway. The work is on the initial levels of a descriptive phenomenon aimed at providing causes of emergent behaviour.

6.1 Acknowledgement

We would like to thank the entire VRSIG team at Rhodes University, without whom we would never have gotten such valuable suggestions, support and moral.

7. REFERENCES

- [1] Adriano Cavalcanti, Tad Hogg, Bijan Shirinzadeh, and Hwee C. Liaw. Nanorobot communication techniques: A comprehensive tutorial. In *IEEE ICARCV 2006 International Conference on Control, Automation, Robotics and Vision*, 2006.
- [2] Andrew Colin. Ant colony algorithms: Solving optimization problems. In *Dr. Dobb's Journal (www.ddj.com/dept/ai/191800178?pgno=1)*, 2006.
- [3] Marco Dorigo, Mauro Birattari, and Thomas Stutzle. Ant colony optimization : Artificial ants as a computational intelligence technique. Technical report, Universite Libre de Bruxelles, BELGIUM, 2006.
- [4] Marco Dorigo, Gianni Di Caro, and Luca M. Gambardella. Ant algorithms for discrete optimization. In *proceedings of Artificial Life Volume 5, MIT*, pages 137–172, 1999.
- [5] Duncan E. Jackson, Mike Holcombe, and Francis L.W. Ratnieks. Trail geometry gives polarity to ants foraging networks. In *Nature*. Nature Publishing Group, 2004.
- [6] Kurumatani Koichi and Nakamura Mari. Generating qualitative equations about macro behaviours of foraging in ant colony. In *Proceedings of the German Conference on Bioinformatics, GCB'96 (Leipzig)*, pages 142–147, 1996.
- [7] Daniel Merkle, Martin Middendorf, and Alexander Scheidler. Modeling ant brood tending patterns with cellular automata. In *Proceedings of the Workshop on Modelling of Complex Systems by Cellular Automata*, volume 2.2, pages 412–419, 2005.
- [8] Daniel Merkle, Martin Middendorf, and Alexander Scheidler. Modeling ant brood tending patterns with cellular automata. *Journal of Cellular Automata*, 2.2:183–194, 2006.
- [9] Yoshiyuki Nakamichi and Takaya Arita. An evolutionary simulation of the origin of pheromone communication. In *Proceedings of MPS Symposium 2004 on Complex System Science and its Application*, 2004.
- [10] Liviu Panait and Sean Luke. A pheromone-based utility model for collaborative foraging. In *AAMAS-2004 - Proceedings of the Third International Joint Conference on Autonomous Agents and Multi Agent Systems*, 2004.
- [11] Liviu A. Panait and Sean Luke. Ant foraging revisited. In *Proceedings of the Ninth International Conference on the Simulation and Synthesis of Living Systems (ALIFE9)*, 2004.

- [12] Liviu A. Panait and Sean Luke. Learning ant foraging behaviors. In *Proceedings of the Ninth International Conference on the Simulation and Synthesis of Living Systems (ALIFE9)*, 2004.
- [13] Steffen Priesterjahn, Andreas Goebels, and Alexander Weimer. Stigmergetic communication for cooperative agent routing in virtual environments. In *Proceedings of the International Conference on Artificial Intelligence and the Simulation of Behaviour*, 2005.