NTNU – Trondheim
Norwegian University of
Science and Technology

Use of Repellent Pheromones In Collective Exploration

Filip Fossum

Master of Science in Computer Science
Submission date: June 2014
Supervisor: Pauline Haddow, IDI

Norwegian University of Science and Technology
Department of Computer and Information Science

Abstract

This report explores the possibilities of using repellent pheromones in order to achieve efficient dispersion of a swarm of robots. The goal is to have a set of autonomous robots cover an unknown environment by collectively visiting each area once.

Unlike previous research in collective exploration this report aims at achieving intelligent dispersion by relying solely on local perception of repellent pheromones, without any centralized control mechanisms. Using local sensory input and decentralized controllers promotes better scalability of the swarm, and is well suited for application on physical robots.

A small number of homogeneous autonomous robots are tasked to explore randomly generated environments. The robots possess no previous knowledge of the environments nor the precise whereabouts of other robots. Repellent virtual pheromones serve as a collective memory of previously visited areas, aimed at preventing several robots from searching the same locations twice, thus increasing the exploration efficiency.

The usefulness of repellent pheromones in exploration is addressed through agent-based simulations, with respect to population density, environmental layout, and pheromone evaporation rate.

Experimental results show that repellent pheromones can assist robots in exploring unknown environments more efficiently and consistently. The benefits of using pheromones is shown to be greatest in low population swarms and when environments are large or contain few obstacles. Moreover, pheromones evaporating at a slow rate is shown to provide consistent improvements over rapidly evaporating pheromones.

Sammendrag

Denne rapporten utforsker mulighetene for å bruke avstøtende feromoner til å oppnå effektiv spredning av en sverm av roboter. Målet er få en mengde selvstendige roboter til å dekke et ukjent område ved å kollektivt besøke alle steder én gang.

I motsetning til tidligere arbeid innen kollektiv utforskning forsøker denne rapporten å oppnå intelligent spredning ved å kun bruke lokale sanseinntrykk fra avstøtende feromoner, uten noen sentralisert kontrollmekanisme. Bruk av lokale sensordata og desentraliserte kontrollere fremmer skalerbarhet av svermen, og er godt egnet for bruk på fysiske roboter.

Et fåtall selvstendige roboter er gitt i oppgave å utforske tilfeldig genererte omgivelser. Robotene har ingen tidligere kunnskap om hverken omgivelsene eller plasseringen til andre roboter. Virtuelle avstøtende feromoner fungerer som en felles hukommelse over tidligere besøkte områder, som forsøker å hindre flere robot fra å lete i de samme stedene flere ganger, og dermed øke effektiviten av utforskningen.

Nytteverdien av avstøtende feromoner i utforskning er vurdert gjennom agentbaserte simuleringer i forhold til populasjonstetthet, omgivelsenes utforming og feromonenes fordampningsrate.

Forsøksresultater viser at avstøtende feromoner kan hjelpe roboter med å utforske ukjente omgivelser mer effektivt og mer konsistent. Fordelene ved bruk av feromoner er vist å være størst i svermer med lav populasjon, og i omgivelser som er store eller inneholder få hindringer. Videre er det vist at feromoner som fordamper sakte gir konsistent forbedringer sammenlignet med feromoner som fordamper raskt.

Preface

This master's thesis is written at the Department of Computer and Information Science at NTNU during the spring semester of 2014.

I would like to thank Jean-Marc Montanier and Pauline Haddow for their assistance and advice throughout the project. I would also like to thank Christian Berg Skjetne for his technical assistance.

Filip Fossum

Trondheim, June 11, 2014

Contents

List of Figures

List of Tables

Chapter 1

Introduction

This chapter presents the background and goals of this project. A summary of previous work and the motivation for this study is presented in section 1.1. Section 1.2 then explains the goal and research question for this study. An overview of how the research will be conducted is described in section 1.3, before finally the structure of the thesis is summarized in section 1.4.

1.1 Background and Motivation

Deploying autonomous mobile robots in unknown environments is a challenging problem in a number of tasks. In a military setting robots can be useful in scenarios such as reconnaissance, patrolling, surveillance, target tracking and target acquisition [22]. Robots has also seen real world application in both Afghanistan and Iraq where they have been used to search for chemical and nuclear weapons [29]. Another area where robot deployment can be useful is during urban search and rescue operations. In disaster areas, such as collapsed structures, robots can search for survivors in environments that are either inhospitable or inaccessible to rescuers [5]. When deploying robots in unknown environments to search for points of interest, it becomes unnecessary to have the robots make the same discovery twice. Moreover, redundant coverage of the environment directly affects the efficiency of the search.

Coverage of an environment can be considered either a *static* or a *dynamic* problem [2]. The static coverage problem aims at converging to a robot configuration such that every point of the environment is continuously within a robot's sensor range. Obtaining complete static coverage of an environment requires a critical number of robots. Without any previous knowledge of the environment, determining this number can be difficult or impossible. This report is concerned with the dynamic coverage problem. When there are not enough robots to achieve complete static coverage, robots must continually move to new locations in order to uncover all areas of the environment. Rather than measuring coverage as the areas that are currently observable, dynamic coverage is measured by the locations that have been visited at some point in the past.

In swarm robotics, much of the work is inspired by the self-organizing behaviour of social insects, such as honeybees and ants. A well documented form of communication in social insects is the release of pheromones, a chemical compound that can be perceived by other individuals. When ants discover a food source they leave a pheromone trail on the ground on their way home. Other ants can follow this trail and amplify it with their own pheromone to attract even more ants [14].

Attractive pheromones

The use of virtual pheromones for guiding robots have been researched extensively. Yet most attention is directed at emulating the type of pheromones most commonly seen in nature – pheromones that attract other individuals. In [13] Garnier et al. used attractive pheromones to make robots favour a single common path to a food source. This experiment was later reproduced in [25] supporting the positive effect of pheromones for collective behaviour.

Payton et al. proposed in [18] an approach inspired by the diffusion of pheromones. By passing a decrementing counter between neighbouring robots they simulated pheromone diffusion without any centralized control mechanism nor external dependencies. In a similar fashion [23] simulated the exchange of nutriments through direct mouth-to-mouth contact – a behaviour seen in many social insects, known as trophallaxis. By evenly sharing resources between robots upon physical contact, a gradient emerged by the amount of resources each robot was carrying. Not unlike pheromone behaviour, this gradient could be traced to the location at which the resource originated. Both of these approaches bears resemblance to the response of actual pheromones. However, rather than distributing navigational cues indirectly through the environment, the information used to guide the swarm is embodied in the robots themselves. Therefore, in order to produce detailed navigational cues, both of these approaches require a large number of robots.

Several other techniques for distributing pheromones have been documented. The use of chemicals to represent pheromones has been studied in for instance [20], intended for environments where inter-robot radio communication is either infeasible or undesirable. However, since these chemicals are invisible, it becomes difficult study exactly how they influences the robots. In [11] a centralized pheromone server was used to store the pheromone data. While the robots performed a path-following task, they could query the server for pheromone information at their current location. The robustness of this approach is however questionable – the robots are highly dependent on the reliability of continuously receiving and transmitting data from and to the pheromone server.

Repellent Pheromones

While virtual pheromones has predominantly seen use in aggregation and foraging tasks, it has also been applied in the context of exploration and dispersion of robots. Pheromone with repellent properties has been studied in for instance [19] where it was used to disperse robots in an unknown environment in the context of search and rescue operations. In this study robots were mounted with a virtual pheromone emitter. When deployed the robots could perceive the distance to other robots and attempt to maximize it. However, maximizing dispersion does not necessarily entail efficient coverage of the environment. Maximum dispersion is achieved when all robots are evenly distributed throughout the environment. Yet, when environments are complex as a result of obstructions and secluded pathways, more efforts may have to be directed towards particular areas. When robots are dispersed evenly throughout the environment, a large number of robots may be required to ensure that the coverage is accordingly even.

The use of repellent pheromones was also reported in [24], where it was secreted by robots searching a maze for a single point of interest. Repellent pheromones was reported as a successful technique for making robots favour unexplored areas of the environment. The pheromone data was however stored in internal maps that was distributed globally between all robots. By using this global data, exploration was performed according to predetermined paths calculated by an `A`* algorithm. While the approach was reported as effective, distributing large amounts of data between robots can limit the scalability of the swarm. The number of inter-robot connections required to distribute data globally grows exponentially with the size of the swarm. This can make the approach computationally infeasible outside of simulated environments.

In [22] multiple different applications of virtual pheromone was studied for use in military scenarios. Several types of pheromones, with both repellent and attractive properties, were utilized to complete tasks in surveillance, patrolling, target tracking and target acquisition. Promise was reported on the versatility of pheromones to coordinate autonomous robots in a variety of tasks.

Motivation

Virtual pheromones with attractive properties have been studied extensively. Yet little attention has been received by its repellent counter part. Despite its scarcity in the literature, repellent pheromones is reported to show promise in collaborative exploration. Previous work has however taken advantage of aspects such as global distribution of data, which can make the transition to physical robots difficult. This report examines the possibility of retaining the positive effects of repellent pheromones when the information required by the robots is distributed indirectly through the environment. Under this condition robots must react solely based on local sensory input.

Moreover, there is little mention in the literature of precisely when repellent pheromones might be useful, nor under which conditions it ceases to be advantageous. This report aims to address these questions.

1.2 Goals and Research Questions

The goal of this study is to investigate the possibilities of exploring unknown environments with a swarm of robots that possess no way of communicating directly with each other. Collaboration between individuals often involves gathering at specific locations in order to overcome complex tasks. However, collaboration can also involve delegating tasks to different individuals in order to achieve a common goal. This study proposes an exploration task where a swarm of robots must work individually in order to collectively cover all locations of an unknown environment.

Goal *Achieve efficient exploration of an unknown environment by having a swarm of robots visit each area once*

When robots are able to complete tasks individually their proximity could potentially lead to unnecessary redundancy. It becomes desirable for them to disperse in order to explore as much of the environment in the shortest possible time.

Research question: *Can the use of repellent virtual pheromones limit redundant exploration when robots are deployed in unknown environments?*

While attractive pheromones have been researched extensively, there is a shortfall of studies on pheromones with repellent properties. There is also little mention of when repellent pheromones might be useful in exploration. If the number of agents changes, will it influence the benefits of using pheromones? Can the pheromones be used in the same manner irrespective of the appearance of the environment? How is the performance of the agents affected by the longevity of the pheromones? These questions will be assessed in this report.

The ultimate goal of the work in this report is future exportation to physical robots. Previous work on repellent pheromones in exploration has taken advantage of global distribution of pheromone data. While this may improve the performance, it can be computationally infeasible for physical robots. In this report, pheromones will be distributed indirectly through the environment. Robots must therefore rely solely on local sensory input. Given this constraint, this report assesses the possibility of achieving efficient collective coverage of an unknown environment.

1.3 Research Method

The experiments in this report are performed with an agent-based simulator. This enables the robots to be examined carefully, in order to better understand their emergent behaviour. Robots used in the simulations are modelled to resemble the ChIRP robots [26], which are designed and manufactured at NTNU. Hopefully this will promote future study on the subject. While the presented technique is intended for future exportation to physical robots, this report will focus solely on the performance in simulated environments.

When working with a swarm of robots there are many parameters that can affect the overall performance. In this report these parameters are concerned with the number of agents used, diffusion and evaporation of pheromones, and the appearance of the environment. By performing the experiments in simulation, these parameters can be modified effortlessly. Furthermore, simulated environments eases the collection of data required to measure the performance of the swarm.

To assess the viability of the pheromone technique a series of simulations is performed. Many of the experiments are performed in pairs to see the effect when the only difference is the presence of pheromones. While there are many random factors present in the experiments, pairs of experiments are always initialized with the exact same randomly generated data. This concept is further elaborated in the experimental setup in section 4.3.

1.4 Thesis Structure

This report is structured as follows. Chapter 2 gives a brief overview of the field of swarm intelligence including some of the biological concepts it is inspired by. Chapter 3 summarizes the development of the simulator with respect to the representation of virtual pheromones and the control architecture of the robots. In chapter 4 the experiments are presented along with their setup and final results. Finally, chapter 5 presents an overall conclusion of the work, as well as suggestions for future work.

Chapter 2

Background Theory

This chapter introduces some fundamental theory in the field of swarm intelligence and swarm robotics. Some basic concepts of biology relating to indirect coordination of a swarm are presented, along with a few examples of their application in the real world. Finally, in section 2.3 the literature review protocol is summarized, including search engines, search terms and inclusion- and quality criteria.

2.1 Swarm Robotics

Swarm intelligence is concerned with the emergence of collective intelligent behaviour in a large number of individuals. It is inspired by the biological studies of the self-organizing behaviour seen in social insects [14]. Honeybees and ants are examples of social insects which by themselves are very simplistic creatures, but through collaboration with others constitute a complex system. Substituting the social insects with simple autonomous mobile robots incorporating their behavioural traits is the essence of swarm robotics [3].

2.1.1 Properties of a Swarm

A large collection of robots does not necessarily constitute what we refer to as *swarm robotics*. Natural swarms have been observed to posses certain properties, namely robustness, scalability and flexibility [30]. In order to correctly emulate the behaviour of social insects these serve as guidelines of how to create an efficient swarm of robots.

In terms of fulfilling these properties, an important factor is to decentralize the control mechanism. This entails that no single entity should dictate the behaviour of the individuals in the swarm. Studies have revealed that there is no centralized coordination mechanism governing the synchronized behaviour of social insects [9]. Yet, social insects show impressive collective problem-solving capabilities.

Decentralizing the control mechanism promotes *robustness* of the system. Robustness requires that the swarm is able to continue to operate despite failures in the individuals. If the coordination of the swarm is dependent on a single entity, its failure would affect the entire system. When the control mechanism is self-contained, the loss of a single entity should merely reduce the efficiency of the swarm as a whole. The robustness of a swarm is naturally contributed by the redundancy present in the swarm. This accounts for both the redundancy in the amount of individuals, as well as redundancy in the individuals' sensing and control mechanisms.

The number of individuals in the swarm directly relates to *scalability*. This involves that the inclusion of additional entities in the swarm should not negatively impact the system as a whole. Quite contrary, such an addition should make the swarm more proficient at completing its task. With centralized control mechanisms the addition of entities could lead to infeasible computational overhead.

The final property observed in social insects is their *flexibility*. This involves the swarm's ability to dynamically utilize different coordination strategies in response to environmental changes. Flexibility is also concerned with the swarm's capability of generating solutions to new tasks, as well as dynamically coordinating the execution of multiple tasks at once.

2.2 Pheromones

Social insects communicate in many different ways. The honeybee waggle dance is an example of direct coordination. When honeybees return from a food source they move in an eight-figure pattern waggling their body to indicate the direction and distance to the discovered food source.

Indirect coordination within the swarm by leaving traces in the environment is known as *stigmergy*. These traces can be perceived by other individuals and trigger different behavioural patterns.

One well known form of stigmergy is the release of pheromones. Foraging ants are known to release a chemical substance when they discover a food source. Other ants will follow this trail to find the same location. When these ants return they strengthen the scent by releasing more pheromones. In this way ants are not only able to find the paths to food sources, but also deduce based on the pheromone intensities, which food source is closer [14]. Most species of ants release pheromones in such a manner – between two points. However [10] shows that certain types of ants release pheromones continuously during exploration in the absence of food source discoveries. These pheromones indicate the direction of previous explorers, thus guiding new recruits to the frontier, avoiding searching the same area twice.

Pheromones are most commonly used to attract other individuals in the process of recruiting workers. However, social insects have also been proven to secrete pheromones with repellent effects. Honeybees leave scent-markers to signify the rejection of exhausted or recently visited flowers, such that they do not return to the same flower at a later stage. These scent-markers have been shown to also affect other individuals of the swarm [15].

Certain types of ants also release repellent pheromones, albeit in a different manner. When paths branch in different directions, Pharaoh's ants have been proven to leave repellent pheromones at the entrance to the unrewarding path, thus guiding the swarm in the direction of a food source [21].

While there is an abundance of different types of pheromones that elicit different behaviours, they can be generalized into two categories – releaser- and primer pheromones [27]. *Releaser pheromones* evokes immediate behaviour reaction upon reception. *Primer pheromones* on the other hand alter long-term behavioural traits. This report is concerned with the secretion of releaser pheromones.

2.3 Structured Literature Review Protocol

For the literature review the following research questions were posed:

RQ1 What are the existing solutions to controlling a swarm of robots with the aid of virtual pheromones

RQ2 Can the use of repellent pheromones increase the efficiency of a swarm performing a foraging task.

	Swarm	Pheromone	Repellent	Foraging
Term 1	Robotics	Virtual	Repulsive	Dispersing
Term 2	Intelligence	Digital	Attractive	Path Finding
Term 3		Artificial		Aggregation
Term 4		Synthetic		Recruiting

Table 2.1: Search words and synonyms/relevant terms

RQ3 Can the combination of both attractive and repellent pheromones trigger new possibilities in a foraging task.

RQ4 How does the use of repellent (and attractive) pheromones compare to solutions with pure attractants.

The literature search was performed using the terms in table 2.1. The top row lists some relevant terms while each corresponding column accounts for either synonyms or directly relevant terms. These terms where used to create the boolean query in (2.1). This query was used to search the following search engines:

- IEEE Xplore
- ScienceDirect
- SpringerLink
- Wiley Inter Science
- Google Scholar

For each search engine the 20 most relevant papers were extracted, for a total of 100. These were put through an iterative elimination process. Duplicates and papers unavailable through NTNU's license were immediately discarded.

Initially papers were eliminated from the corpus based on the *title* alone, following the criterias of 2.2. The second iteration eliminated papers based on the *abstract* and their inability to fulfill the inclusion criteras in table 2.3. Papers remaining after screening the *content* based on the same inclusion criteras, were subject for full reading, and kept based on the quality critera of table 2.3.

$$
\begin{aligned}
swarm &\wedge (robotics \vee intelligence) \\
pheromone &\wedge (virtual \vee digital \vee artificial \vee synthetic) \\
&\wedge (repellent \vee repulsive \vee attractive) \\
&\wedge (foraging \vee path\,finding \vee aggregation \vee recruiting)
\end{aligned}
\tag{2.1}
$$

	Title criteria
1	The study is concerning swarm robotics or is closely related (i.e. biology)
2	Main focus appears to be on techniques for swarm behaviour.

Table 2.2: Title screening guidelines

ID	**Criteria**
IC 1	The study relates to swarm robotics
IC 2	The study is a primary study
IC 3	The study study incorporates use of virtual pheromones
IC 4	The study describes a specific method to complete a specific task
QC 1	There is a clear statement of the purpose of the study.
QC 2	The study is put into context of other studies and research.
QC 3	Design decisions are justified
QC 4	Concepts are applicable to physical robots
QC 5	Results are thoroughly and unbiasedly analysed.
QC 6	Performance metrics are explained and justified.
QC 7	The experimental procedure is thoroughly explained
QC 8	The study reproducible
QC 9	The study incorporates repellent pheromones
QC 10	The study incorporates both repellent and attractive pheromones.

Table 2.3: Inclusion and quality criteria

Chapter 3

Architecture/Model

This chapter introduces the simulation tools used to conduct the experiments. Section 3.1 introduces the simulator Roborobo!, and how its functionality has been extended for this study. In section 3.2 the development of the pheromone model is documented, followed by descriptions of the diffusion and evaporation of pheromones. The complete architecture of the robot controller is presented in section 3.3, before finally section 3.4 presents an algorithm created to generate random environments.

3.1 Simulator

A simulator called Roborobo! [7] was utilized to conduct the experiments. Roborobo! is an open source swarm robot simulator written in `C++`. It uses the multi-platform SDL graphics library[1] to display robots in a 2-dimensional world. Figure 3.1 shows a graphical view Roborobo! before any modifications have been made.

For this project Roborobo! has been extended with functionality to release and detect pheromones. The robot model of Roborobo! is modified to more closely resemble the sensor placement of the ChIRP robots. This includes eight proximity sensors that are distributed evenly around the robot. Figure 3.2a depicts the location and orientation of these proximity sensors. Each sensor has a range

[1] `http://www.libsdl.org/`

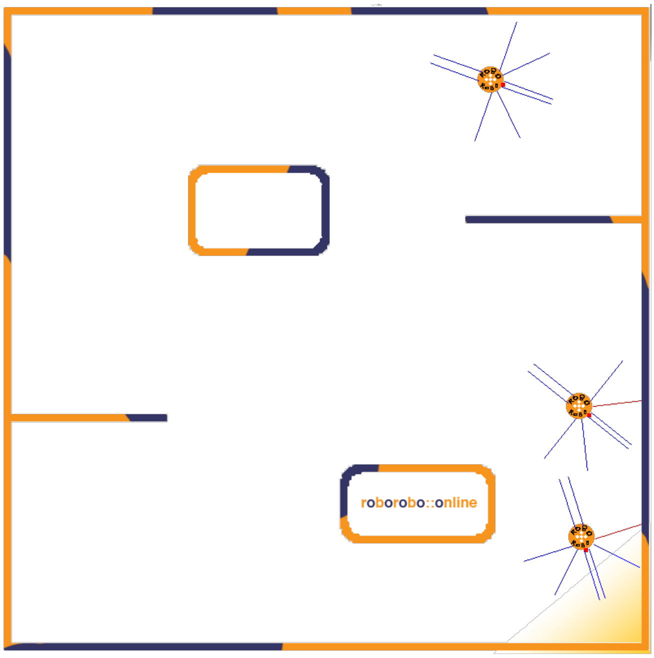

Figure 3.1: Roborobo! – an agent-based simulator

of 30 pixels in its respective direction and is able to measure the distance to obstacles within this range.

The robots in Roborobo! are also equipped with 8 pheromone sensors. Pheromone sensors are currently not available on the ChIRP robot. However, due to its modular design, the ChIRP robot can easily be extended to include additional sensors [26]. The pheromone sensors of the robots in Roborobo! are distributed evenly around the robot, as depicted in figure 3.2b. Each sensor is positioned 20 pixels from the robot center and measures the pheromone intensity level at a single pixel.

The robot model emulates the movement of ChIRP robots by assuming that it has two *actuators* – or wheels. When both actuators are moving at identical speeds, the robot will move in a straight line. By slowing down either actuator the robot will turn in the corresponding direction at an angle proportional to the difference in actuator speeds.

Each robot is 5 pixels in diameter and can move at a maximum velocity of 1 *pixel/second*.

(a) Proximity Sensors (b) Pheromone Sensors

Figure 3.2: Sensor placement of robots

3.2 Virtual Pheromone

Several different models for representing the pheromones were developed. These are documented in section 3.2.1, along with positive and negative properties of each approach. The final model, using a cellular automata, is presented at the end of this section. Details on the evaporation and diffusion of the chosen model is elaborated in section 3.2.2 and 3.2.3, respectively.

3.2.1 Pheromone Models

The presented models aim to incorporate multiple key factors of pheromones. The most important of these are *evaporation* and *diffusion*. Evaporation determines the rate at which the pheromones decrease in intensity, while diffusion is concerned with how quickly the pheromones propagate throughout the environment. Distinct pheromone trails should also possess the ability to blend, such that two weak trails constitute one potent. This can yield a stronger repellent effect in areas that have been visited multiple times. Another important property of pheromones is their intensity variations – recently secreted pheromones should have a high intensity value that gradually declines. By moving in the direction of lower intensity pheromones, robots can reduce the likelihood of encountering other robots. Moreover, since the robots are relying on local sensory input, they are dependent on receiving different readings across their various sensors – if all sensors detect the same intensity value it renders it impossible for the robot to determine its best course of action.

Seed Points

Early versions of the model simulated pheromone through a series of expanding circles. While navigating, *seed points* were placed at the robots' present locations. This occurred at fixed intervals. The seed points were the center of gradually expanding circles. All seed points were independent entities that together constituted a pheromone trail. Since every seed point was independent, both diffusion and evaporation could be controlled irrespectively. Diffusion was simulated by drawing circles around the seed points with increasingly large radii. Evaporation was handled by changing the translucency each of the circles were drawn with. Figure 3.3a depicts the resulting pheromone trail of this model. The circles around more recently activated seed points can be seen to overlap neighbouring circles. This creates a seemingly smooth transition in intensity values. However, since the color of each distinct circle is uniform, the variations do not offer enough precision. The pheromone sensors of the robots are positioned relatively close together. Therefore, robots often detected intensity values from the same circle across all sensors, rendering it impossible to make informed movement decisions.

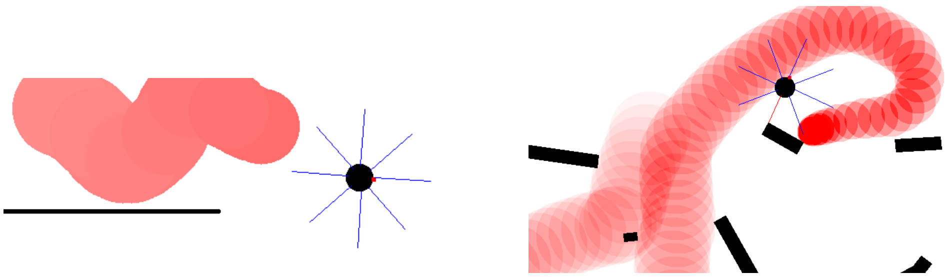

(a) Overlapping expanding circles (b) Blending expanding circles

Figure 3.3: Pheromone model using expanding circles

The model presented above lacked the ability for pheromones to blend. In an attempt to incorporate this property, the same technique of using seed points was employed. However, rather than making adjacent circles overlap, their intensity levels were instead added together. Figure 3.3b shows the result of this change. While intersecting pheromone trails can be seen to blend, the same applies to adjacent circles. This causes the intensity variations in the trail to be inconsistent, which lead to erratic sensor readings for the approaching robots.

Lattice-gas Automaton

In order to achieve a more realistic representation of the pheromone trails, a *lattice-gas automaton* was implemented. Lattice-gas automata are types of cellular automata used to model systems of moving particles. In such a system, cells are defined by the presence of particles. A set of transition rules determine under which conditions these particles move to adjacent cells.

This pheromone model was developed by incorporating the transition rules of the HPP gas model [16]. The HPP gas model is based on a two-dimensional grid where every cell has a set of four velocity vectors – one for each cardinal direction. Every iteration particles translate to neighbouring cells according to these vectors. When two directly opposing particles collide their direction is shifted by 90 degrees. In order to represent the pheromones an additional intensity value was stored in each cell. Figure 3.4 depicts the pheromone representation of this model. Despite the simplicity of the model, simulating particles is computationally expensive. Furthermore, because of the particle collision, the resulting pheromone trails were often erratic. The lacking control of the pheromones' diffusion ultimately rendered this model unsuitable to reliably assist the other robots.

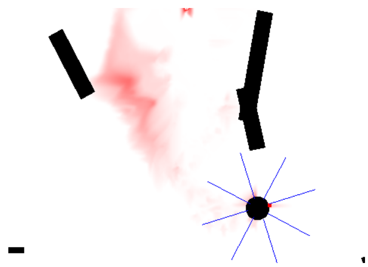

Figure 3.4: Pheromone model using HPP lattice gas automaton

Cellular Automaton

To accommodate the lacking predictability of lattice-gas automata, a more traditional cellular automaton was developed. Cellular automata are discrete dynamic systems that models complex behaviour through simple local rules. A cellular automaton is generally represented using a d-dimensional lattice of cells, which is referred to as the *cellular space*. The information contained in a cell represents its *state*, which is one of a finite number of acceptable values. Often one of the states is known as the *quiescent state*. This represents an inactive or resting state for the given cell. *Transition rules* determine how the state of cells change over time depending on the cell's *neighbourhood*. The neighbourhood of a cell can be arbitrary, however it is typically comprised by a small number of adjacent cells [12]. Figure 3.5 depicts two common types of neighbourhoods.

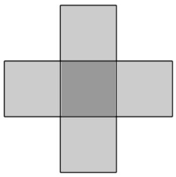

(a) Von Neumann

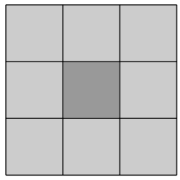

(b) Moore

Figure 3.5: Two common neighbourhoods in cellular automata

The cellular space of the developed pheromone model is defined as a two-dimensional grid where every cell corresponds to one pixel of the environment background. Each cell may take on a value in the range $[0 - 255]$ which represents the pheromone intensity level of that cell. A cell with a value of 0 is in its quiescent state. Robots secrete pheromones by setting the value of a single cell, directly beneath it, to 255. The transition rule of the model is based on a *Von Neumann neighbourhood*, where every cell assumes the greatest value among its four neighbours. This transition rule is depicted below.

```
buffer_cell[x][y] = max(  cell[x+1][y],
                          cell[x-1][y],
                          cell[x][y+1],
                          cell[x][y-1] ) * Evaporation Factor
```

Figure 3.6: Transition rule of cellular automata pheromone model

This transition rule is applied to each cell in the cellular space at every iteration. In order to prevent the cell values from cascading throughout the grid, values must be placed in a temporary buffer before they are applied to the grid. This ensures that every cell is assessed based on the values of the previous iteration. Every cell is also multiplied by an evaporation factor < 1 to simulate the evaporation. The details of the evaporation factor is elaborated in section 3.2.2.

Lastly when the grid is updated, an *averaging filter* with a *Moore neighbourhood* is applied to the grid. This is done to smooth out the rough edges produced by the transition rule, thus creating greater variations in neighbouring values.

The averaging filter functions in a similar fashion to the above rule, except every cell receives the average value of its eight surrounding neighbours. Figure 3.7 depicts the results of this pheromone model. This is the model used to conduct the subsequent experiments.

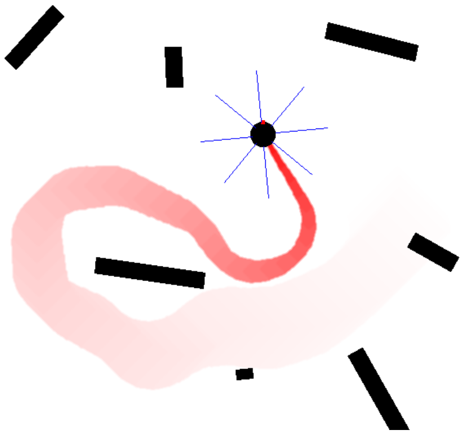

Figure 3.7: Pheromone model using cellular automata

3.2.2 Evaporation

Evaporation is the rate at which the pheromones' potency diminish over time. This potency is simulated by continually reducing the alpha value of pixels subjected to pheromones, thus increasing the pixels' level of transparency. The alpha channel of each pixel is expressed by 1 byte, such that the intensity of the pheromone is represented by a value in the range $[0, 255]$. A pixel with an alpha value of 255 is completely opaque, while a pixel with an alpha value of 0 is fully transparent.

The evaporation rate function is inspired by the work in [13]. While the original formula is intended for attractive pheromones the underlying functionality is the same for repellent pheromones. Pheromone intensities are updated at a given timestep interval Δt. Existing intensity values are then reduced at every pixel by an *evaporation factor*, given in (3.1). The evaporation factor is dependent on two parameters – The number of timesteps between every update, Δt, and a constant t_c which determines how long it should take for the pheromone to evaporate completely. In order to preserve consistency between internal data and the visual representation, the resulting values are rounded down to the nearest integer. This formula is given in (3.2).

$$Evaporation\ Factor = \exp\left(\frac{\log(\frac{1}{2})}{t_c}\Delta t\right) \tag{3.1}$$

$$alpha_{\,t}(x,y) = \lfloor alpha_{\,t-\Delta t}(x,y) \times Evaporation\ Factor \rfloor \tag{3.2}$$

3.2.3 Diffusion

Diffusion is the rate at which the pheromones expand and propagate throughout the environment. Since the pheromones are modelled using a cellular automaton this diffusion is determined by the neighbourhood of the transition rule. As the pheromone model uses a Von Neumann neighbourhood, the pheromones will diffuse by one pixel in each of the four cardinal directions every time it is updated. The diffusion can however be controlled by modifying the update interval Δt. Lower values of Δt will cause the pheromone to be updated more frequently and thus diffuse at a faster rate.

3.3 Robot Controller

The robot controller dictates how the robots will navigate throughout the environment. Every robot is homogeneous, meaning that they all contain the exact same controller. Moreover, the robots are all independent entities – there is no centralized control system affecting the robots in any way.

The navigation of the robots is determined solely based on local sensory input. Steering mechanisms often related to this property is presented in section 3.3.1. These mechanisms constitute elements of the complete controller, which is based on the subsumption architecture. This type of architecture is described in section 3.3.2, followed by a detailed description of how its various elements are implemented in the robots' controller.

3.3.1 Braitenberg Vehicles

The behaviour of the robots in this report is based entirely on local sensory input. These sensors are able to detect both pheromones and obstacles in the environment. Each of the active sensors on an agent relate to one or both of its actuators. Whenever a sensor receives input, the appropriate actuator is signalled, resulting in a change in movement.

Robot steering architectures utilizing a direct sensor-actuator coupling is often referred to as Braitenberg Vehicles [6]. In his book Valentino Braitenberg describes a series of robot models that have very simple internal structures, yet are able to produce complex and seemingly intelligent behaviour. The robot models he presents differ in the way sensors are connected to actuators, and also whether the connections are *inhibitory* (negative) or *excitatory* (positive). Excitatory connections causes an increase in movement upon sensor input, while inhibitory connections results in a reduction in movement upon sensor input.

Figure 3.8 depicts two Braitenberg vehicles that both have inhibitory sensor-actuator connections. The only way the two vehicles differ is in the way the sensors are connected to the actuators. Both vehicles maintain a steady forward translation until the depicted light source is detected. When the light source is detected by a sensor, the speed of its connected actuator is lessened. When vehicle **A** detects the light source on its left sensor, the speed of its left actuator is slowed down. This causes the vehicle to turn left, towards the light source. Conversely, since vehicle **B** has cross-coupled connections, sensor input on the left side will result in the vehicle turning right – away from the light source.

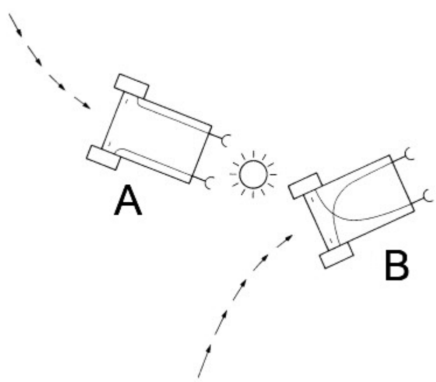

Figure 3.8: Braitenberg vehicles with inhibitory sensor-actuator connections

The control mechanisms of the robots in this report are concerned with cross-coupled inhibitory sensor-actuator connections, such as vehicle **B**. This type of control mechanism was successfully applied to physical robots in for instance [17]. Details on how the pheromone- and proximity sensors are connected to the actuators are elaborated further in section 3.3.2.

3.3.2 Subsumption Architecture

The controller of the robots is designed using Brooks' subsumption architecture [8]. In the process of designing a robot control system for physical robots, [4] evaluated a series of architectures. Brooks' subsumption architecture was favoured for its simplicity while still being able to produce complex behaviour. Since the ultimate goal of this project is future exportation to physical robots, this type of architecture is appropriate also for this controller.

The subsumption architecture decomposes the complete behaviour of an agent into sub-behaviours. These sub-behaviours are organized in a vertical hierarchy of layers, as depicted in figure 3.9. Each layer of the architecture introduces a new behavioural competence to the agent. All layers of the architecture run in parallel. Therefore, behaviours found on the higher tiers still utilize the competences of the layers below. However, higher levels *can* suppress the output of the lower levels.

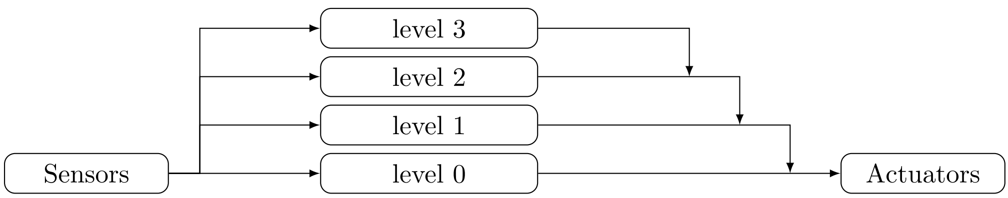

Figure 3.9: Brooks' subsumption architecture

The subsumption architecture enables different control systems to be developed independently. This property is desirable for the controller of this project, as behaviour related to the release and perception of pheromones can be modified irrespective of the agent's core functionality. Certain layers may also be omitted from the controller altogether. This property is utilized during the experiments in section 4.4.1 and 4.4.3, where the layer concerned with secretion of pheromone is electively omitted, in order to demonstrate the performance with and without pheromones.

The complete architecture of the robot controller is depicted in figure 3.10. Each layer of the architecture is elaborated below.

Level 0: Move

This is the most basic building block of architecture. All this layer is concerned with is setting both actuator (or wheel) speeds to maximum, causing the agent

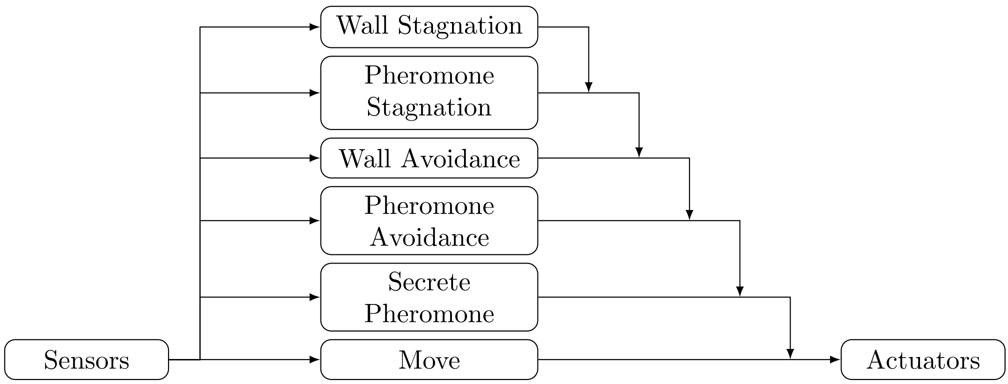

Figure 3.10: Robot controller architecture

to move forward.

Level 1: Secrete Pheromone

This layer handles the release of pheromone. If not suppressed, pheromone is continuously secreted directly beneath the agent.

Level 2: Pheromone Avoidance

This layer remains latent until pheromones are detected on either of the agent's active pheromone sensors. Four of the agent's eight sensors are being used to detect pheromones – two on each side of the robots. Each of these sensors are coupled to the actuator on the opposing side of the robot, as depicted in figure 3.11. When pheromones are detected on either of these four sensors, all active pheromone sensors are queried. However, only the sensor reading the highest level of pheromone concentration will send a signal to the actuator it is coupled with. This signal will dampen the speed of that actuator, effectively causing the robot to steer away from the direction with the highest concentration of pheromone. The speed reduction of the actuator is proportional to the intensity of the detected pheromone. Thus, higher concentration of pheromones will cause the agents to make a more abrupt turn.

There is one exception to the behaviour presented above. If the maximum sensor reading is equal on both a left and a right sensor, the agent will be unable to make a fair assessment as to which direction to move towards. In such situations the robot will simply continue on its current path and await variations and the subsequent sensor readings.

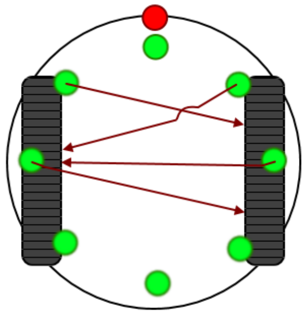

Figure 3.11: Pheromone sensors: sensor-actuator connections

Level 3: Wall Avoidance

Similarly to pheromone avoidance this layer remains latent until obstacles are detected across either of the agent's active proximity sensors. Obstacles include walls in the environment, the outer bounds of the environment, as well as other agents. Since these obstacles are physical constraints to the agent, avoiding them becomes a higher priority task than avoiding pheromones. Therefore, this layer suppresses the output of pheromone avoidance.

Five of the agent's proximity sensors are used to avoid obstacles. These sensors are directly coupled to the agent's actuators as illustrated in figure 3.12. These connections are inhibitory, such that sensor input is directly translated to a reduction in speed of the connected actuator. Since the sensors and actuators are cross-coupled the agent will steer away from obstacles. For instance, if an obstacles is detected on the agent's right side, the left actuator will slow down, causing the agent to steer left. The speed reduction on the actuators are proportional to the distance input. Thus, if the agent is close to an obstacle it will perform a sharper turn to avoid it. Notice that the front sensor of the agent is connected to both actuators. Therefore, when agent an approach obstacles head-on, it will slow down. This is necessary to prevent the agent from colliding with obstacles, as it won't be able to make an abrupt enough turn at full forward velocity.

This type of internal robot structure is similar to Braitenberg's third robot model, namely *the explorer* [6].

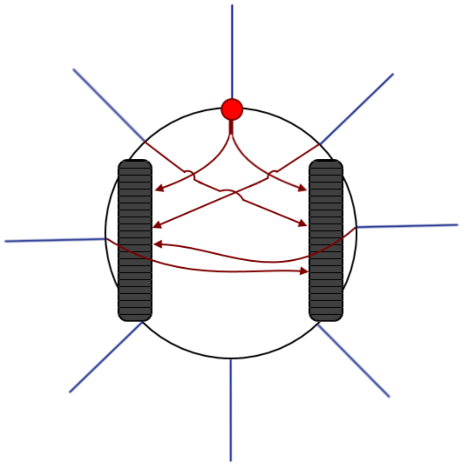

Figure 3.12: Proximity sensors: sensor-actuator connections

Level 4: Pheromone Stagnation

The purpose of the pheromones is to improve the robots' ability to disperse. Yet it quickly became apparent that the pheromones could inadvertently create situations that temporarily halted the robots' progression. Since the robots continuously secrete pheromones they become highly susceptible to their own trail. Figure 3.13 depicts a situation where two agents are caught in spirals of their own pheromone trail.

Figure 3.13: Agents can become temporarily trapped by their own pheromone trail

In an attempt to disrupt monotonous behaviour that emerges when agents intersect with their own pheromone trail, a stagnation mechanism was developed. This mechanism utilizes two counter – one for left movement, and one for right movement. Whenever the agent's reaction to pheromones result in a leftward movement, the left counter increments. Similarly, when pheromones causes the agent to turn right, the right counter increments. Conversely, when the pheromones do not cause cause the agent to turn in a specific direction, the corresponding counter decrements. Whenever one of the counters reach a specified threshold of 50, pheromone avoidance is suppressed for 10 timesteps.

Despite the inclusion of this mechanism, pheromones can still inadvertently cause delays in the swarm. This problem is further discussed in section 5.1.3.

Level 5: Wall Stagnation

Wall stagnation is the top level control scheme and assumes the highest priority. It is directly concerned with the survival of the robots. Wall stagnation is detected based on the agent's lack of movement in the environment. Throughout the simulation Cartesian coordinates of the robot's 100 latest positions are recorded. Standard deviation of the x and y values are calculated respectively. If the sum of these deviations is less than a specified threshold of 10, the robot is considered stuck. When this occurs the layer assumes full control of the robot and suppresses the output of all other layers.

Stagnation resolution is carried out in two steps. The first step is to reverse for a short duration. This allows the robot to get some clearance from the obstacle which is stagnating its movement. Next, the robot will stop and rotate to face the direction with the least obstruction. This direction is decided by querying each distance sensor along with its two adjacent sensors. The collection of sensors that yields the lowest accumulated distance value determines the most desirable orientation. When the robot has rotated to face the correct direction, control is restored to the other layers.

3.4 Random Environment Generator

Simple wall avoidance is one of the key factors of the robots' controller. Therefore the performance of the system is greatly influenced by the layout of the environment the robots are navigating. Ideally the system should function well in any environment. However when both the environment and the controller are made by the developer it becomes difficult to neglect bias in the construction process. On a subconscious level the developer may know which obstacles might prove a challenge for the system and which may not.

In order to reduce the potential of any bias a random environment generator was created. The developed generator is inspired by an algorithm for modelling street patterns, presented in [1]. The original algorithm uses local optimization to create a network of roads such that a number of cities are interconnected. Rather than creating an interconnected network of roads, similar principles are employed to generate separable line segments based on randomly spawned points. A description of the environment generator's implementation is presented below.

Algorithm

The first step of the algorithm is to initiate a set amount of *points* at random positions throughout the environment. A given number of these points will be marked as *seed points*. For every seed point, vectors are drawn to all regular points within a specified *radius*, as depicted in figure 3.14a. These points are removed and the average of these vectors constitute a line that acts as an obstacle in the environment. This is illustrated in figure 3.14b. At the end position of the newly created line an additional seed point is spawned. This seed point is promptly treated similarly to the one before, thus extending the line. The maximum number of times the line will be extended is determined by a random number in the range [1, $maxJumps$], where $maxJumps$ is a user specified variable. If this threshold is exceeded, or no points are detected within range of the specified search radius, the next seed point among those initially spawned is assessed. This process continues until no seed points remain.

The complete source code for the random environment generator can be found in appendix A.1.

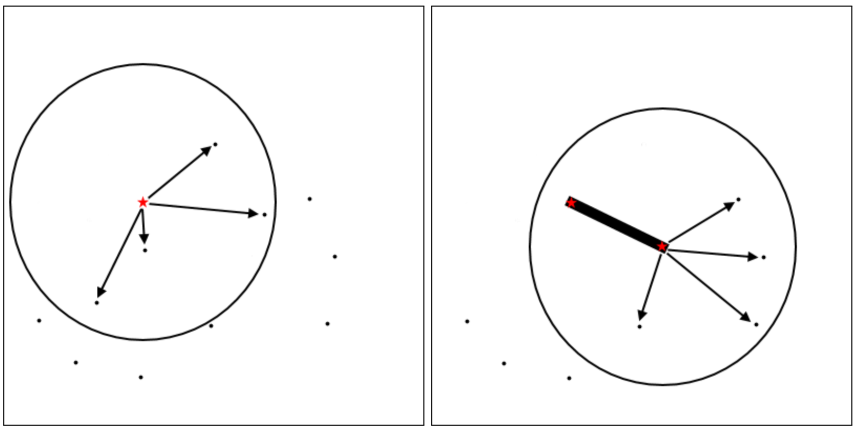

(a) Vectors are drawn from a seed point to all points within the threshold radius

(b) These points are deleted and the average vector form an obstacle as well as a new seed point. This repeats until no points are in range.

Figure 3.14: Random environment generator

Chapter 4

Experiments and Results

This chapter presents the experiments performed as well as their results. The experiments aim to assess the viability of repellent pheromones in an exploration task. Furthermore, experiments will assess how the agents' exploration abilities are affected by increasing number of agents, changes in environmental complexity, as well as differences in the pheromone evaporation rate. The hypotheses for these various situations are presented in section 4.1. Section 4.2 addresses briefly how the experiments will aim to answer these questions. Parameters and further details of the experiments are introduced in section 4.3. Finally, each question is assessed in turn in section 4.4 through the results of the experiments.

4.1 Hypotheses

The main hypothesis is that repellent pheromones can assist a swarm of robots in more efficiently covering all locations of an unknown environment. Pheromone trails left in the environment can yield information about recently visited areas which is believed to help the agents disperse and avoid exploring the same locations. Agents releasing pheromones are therefore expected to cover all areas of an environment more efficiently than agents whose movement are entirely dictated by the avoidance of obstacles.

Since the agents act exclusively on local sensory input, pheromones will naturally have no effect until other agents discover them. However, increasing the likelihood of pheromone being discovered by excessively increasing the longevity of pheromone trails, might not be a viable solution. Quite contrary, it is believed that excessive amounts of pheromones may impede the agents' exploration abilities. Since the agents react to pheromones by assessing intensity variations. An abundance of pheromones might seclude these variations, disabling the agents from making informed decisions.

The amount of pheromone present in the environment is influenced by several factors. Naturally, the evaporation rate as well as the number of agents secreting pheromones, will both have a direct impact on the presence of pheromones. Additionally, since the propagation of pheromones is bounded by obstacles, environments with open areas will be more susceptible to the diffusion of pheromones.

Pheromones are expected to still yield positive effects in exploration as the number of agents is increased. Since the increase in population density reduces the distance between agents, the likelihood of agents encountering pheromone trails become higher. Thus, it is believed that in order to prevent excessive amounts of pheromone, while still retaining its ability to disperse the agents, the rate of evaporation will need to be increased accordingly as the population grows.

Similarly, when environments contain many open areas agents become more likely to discover the pheromone trails of other agents. Also in this case it is expected that the pheromones need to evaporate quickly, as not to diffuse excessively and impede the progression of the agents. On the other hand, if there are many obstacles in the environment, it is expected that the evaporation rate must be slow. The complexity of the environment may cause large time gaps between each time a pheromone trail is discovered. The surrounding obstacles should also help contain the pheromones to the subjected area. This should enable the pheromones to safely evaporate at a slow rate.

4.2 Experimental Plan

The hypotheses of the previous section will be assessed through multiple simulations using Roborobo!. Details on which experiments will be run to address the effect of various properties is described below.

Repellent pheromones

To address the viability of repellent pheromones in an exploration task, all experiments will be performed both with and without pheromones. Despite the many random factors at play during simulations, efforts will be made to ensure that the initial conditions of all corresponding runs are identical, with the exception of pheromones. This will enable the precise effect of the pheromones to be reflected in the results.

Population
: It is yet unclear how well the pheromone technique will perform with increasing agent populations. All experiments will therefore also be performed with 1, 3, 6, 12 and 24 agents. This will help in assessing how higher densities of agents will influence the performance of the exploration.

Environment
: Another question to be addressed is how the appearance of the environment will influence the agents. To better understand this, multiple environments will be randomly generated for the agents to explore. Each environment aims to portray certain general properties in terms of complexity and size.

Evaporation rate
: The final question to be addressed is how the evaporation rate of the pheromones will influence the agents' performance with respect to the population size and the environment. Experiments will therefore also be performed with a selection of different evaporation rates.

4.3 Experimental Setup

For any given set of parameters, 40 independent simulations are performed. The performance of this parameter set is measured as the average percentage of the environment that has been collectively covered by the agents at coinciding timesteps of these 40 simulations.

In each simulation agents are deployed throughout the environment at random locations. However, in order to mitigate the effect of randomness between similar experiments, a collection of spawn locations is only generated once for a particular environment. This entails that all experiments with varying parameters that are performed in the same environment, use the 40 exact same randomly generated spawn locations for its agents. Because of this property, experiments can be performed with different parameters with respect to pheromones, without the results being influenced by any other factors.

Further details of the experimental setup are elaborated in the upcoming sections.

4.3.1 Measure of Coverage

Coverage measures how much of the environment has been visited by the agents. The environment is divided into a lattice of square tiles. Whenever a robot is within the bounds of a tile, that tile is marked as visited. The coverage at any given timestep of the simulation is the total number of tiles that are marked as visited. Figure 4.1 depicts a visual representation of how coverage is measured.

$$Coverage = \frac{Tiles\ Visited}{Total\ Tiles} \tag{4.1}$$

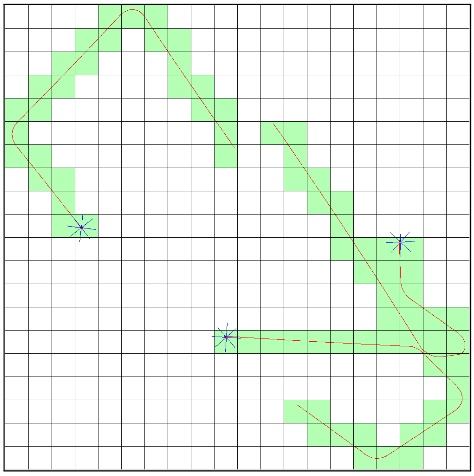

Figure 4.1: Graphical representation of coverage. Highlighted tiles are covered

All the following experiments are performed using a tile size of 50×50 pixels. Multiple different tile sizes were however attempted during preliminary experiments. Figure 4.2 depicts the movement trace of a single agent over the same scenario, and how its coverage of the environment is affected by the tile size.

Figure 4.2a shows the effect of tiles that are too small. The agent can be seen to enter and exit a confined area, yet only about half of tiles within this location are measured as covered. Furthermore, tiles in close proximity to obstacles may be unreachable to the agent. Tiles are registered as covered based on the location of the agent's center. Since the proximity sensors of the agent extends 30 pixels beyond this point, the agent's wall avoidance mechanism may cause the agent to never reach certain tiles.

Figure 4.2b depicts the scenario with a tile size that is too large. Unlike before, the tiles within the confined area are now properly covered. However, since the tiles are so large, they extend far beyond certain obstacles. This causes areas that have never been visited by the agent to be registered as covered.

Figure 4.2c depicts the scenario with the opted tile size of 50×50. Areas visited by the robot are now properly registered as covered without the tiles extended too far beyond obstacles. Additionally, given a sensor range of 30 pixels, all tiles are now reachable by the agent.

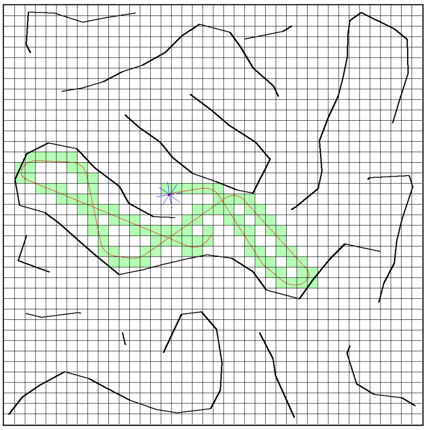

(a) 25×25

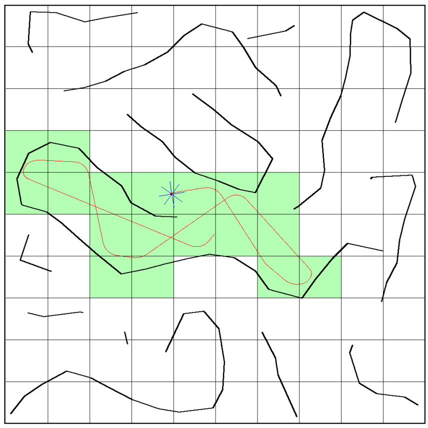

(b) 100×100

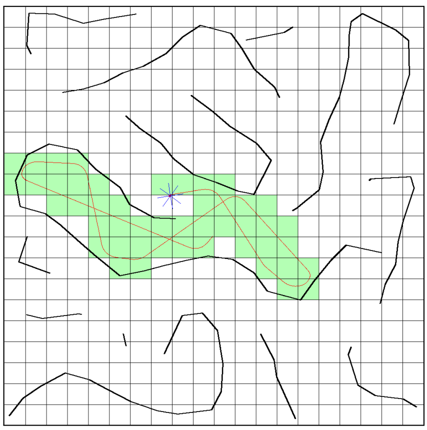

(c) 50×50

Figure 4.2: Different tile size greatly influence how coverage is measured

The number of tiles that need to be visited in order to achieve full coverage, naturally depends on the size of the environment. Section 4.3.3 presents the environments used in the experiments. These vary in size between 1000×1000 and 2000×2000. Since all experiments are performed with a tile size of 50×50, maximum coverage of these environments will be achieved at 400 and 1600 tiles covered, respectively.

4.3.2 Pheromone Evaporation Rates

Four different rates of pheromone evaporation was used for the experiments. The evaporation rate is determined by the constant t_c of equation (3.1) presented in section 3.2.2. These values are summarized in table 4.1.

Evaporation Rate	t_c
Slow	2000
Semi-slow	1000
Semi-fast	500
Fast	250

Table 4.1: Virtual pheromone parameters

The update interval Δt which affects the diffusion of pheromones is kept at a constant value of 25. Too high values of Δt causes the pheromone trails to appear as very thin lines. This can limit the pheromones' ability to influence the agents. Agents may simply pass straight through individual trails, without enough time to react to them. Furthermore, when multiple thin trails intersect, the agents' sensor readings can become erratic. Conversely, too low values of Δt limits the pheromones' ability to produce distinguishable trails. By keeping the diffusion rate constant, it also enables the focus to be directed solely towards the evaporation rate.

Figure 4.3 depicts a visual representation of the difference between slow and fast evaporation rates. The enclosed frames are both 1000×1000 pixels of size.

4.3.3 Generated Environments

Simulations are performed in three different environments which are created using the algorithm presented in section 3.4. Each environment is intended to accentuate one general property. One environment has few obstacles and is relatively open. Another is more complex and has many continuous obstacles. The final environment also contains multiple continuous obstacles, however it is four times larger than the two former.

All the environments, as well as the parameters used to generate them, are presented in figure 4.4.

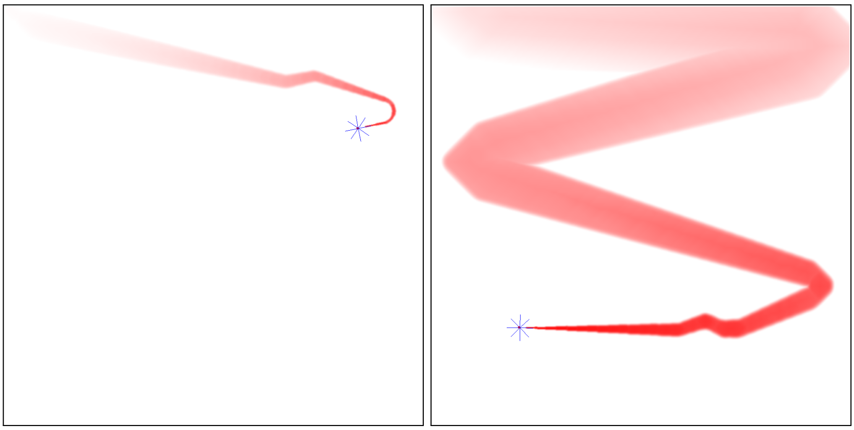

(a) Fast evaporation (b) Slow evaporation

Figure 4.3: Visual representation of pheromone evaporation rates

4.4 Experimental Results

The results of the experiments will be presented in the following sections. The experiments aim to address four main questions:

- Is repellent pheromones a viable technique in collective exploration?
- What will be the influence of increasing agent density?
- How does the appearance of the environment affect the agents?
- How does the evaporation rate of the pheromone influence the performance of the agents?

4.4.1 Repellent Pheromones in Exploration

To assess the viability of using repellent pheromones in exploration, simulations are performed where the only difference is the presence of pheromones. The average coverage is measured at every timestep over 40 simulations. Figure 4.5 depicts the difference in coverage when exploring the environment with few obstacles. The center line of each plot depicts the average amount of covered tiles at the corresponding timestep, while the surrounding fill represents the standard deviation. The pheromone in this simulation evaporates at a slow rate ($t_c = 2000$).

The initial coverage exhibited by both kinds of agents coincide perfectly. This is

Parameter Name	*Value*
Height	1000
Width	1000
Points	2000
Seeds	50
Radius	100
maxJumps	40

(a) Many Obstacles

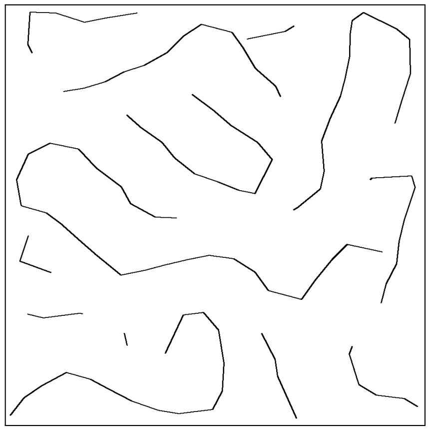

Parameter Name	*Value*
Height	1000
Width	1000
Points	500
Seeds	20
SearchRadius	300
maxJumps	5

(b) Few Obstacles

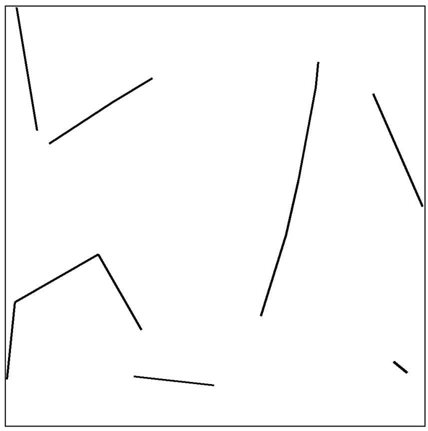

Parameter Name	*Value*
Height	2000
Width	2000
Points	5000
Seeds	50
SearchRadius	150
maxJumps	50

(c) Large Environment

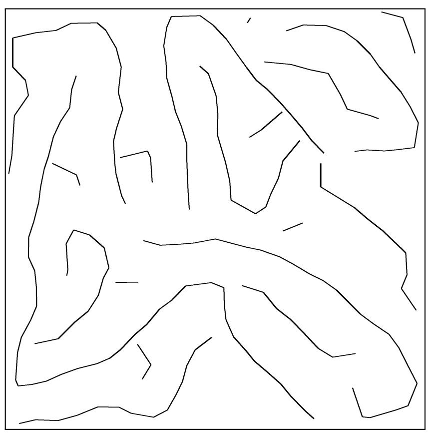

Figure 4.4: Environments used in experiments

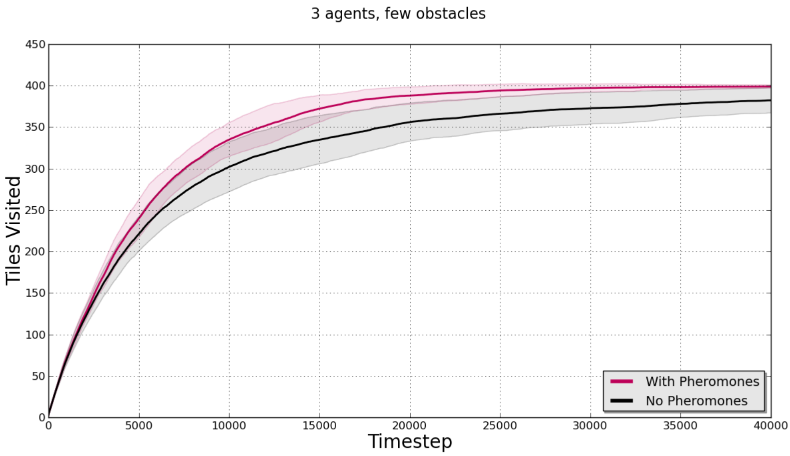

Figure 4.5: Repellent pheromones can improve coverage

to be expected – Their spawn locations are the same and when no pheromones are being detected their behaviours are identical. Since practically any movement at the start of a run leads to the exploration of a new tile, the coverage increases very rapidly. As tiles are being re-visited the increase in coverage slowly dissipates.

The results suggest that the use of pheromones has a positive effect on the agents' ability to cover the environment. Agents acting solely on wall avoidance have their movement entirely dictated by the location of obstacles and the angle at which the agents approach them. This causes the coverage to be highly dependent on their starting position – Some agents may be initialized at locations of close proximity, causing them to inadvertently retrace covered ground. Because of the randomized spawn locations, the deviation from the average coverage is quite large. This deviation is also present when pheromones are being used. However, as the simulations progress and pheromones diffuse throughout the environment, the deviation from the average coverage lessens. The presence of pheromones in the environment provide dynamic influences on the agents, making them more likely to reach the final unexplored areas. After 35,000 timesteps all simulations of agents secreting pheromones have converged to the maximum coverage of 400 tiles.

To illustrate the difference in the behaviour with and without pheromones, consider the events of figure 4.6. This figure depicts the movement history of the agents from one of the 40 simulations presented above.

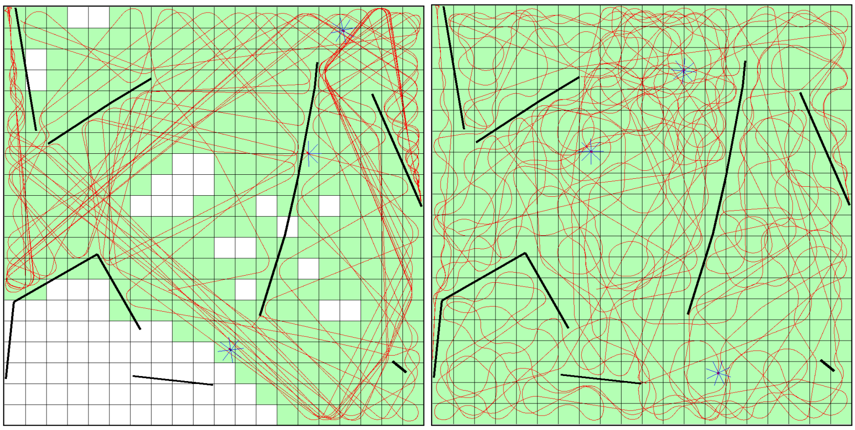

(a) Without pheromones (b) With pheromones

Figure 4.6: Movement trace and covered tiles of 3 agents

Figure 4.6a shows that when pheromones are not being used, the agents may have difficulty exploring certain areas. None of the tiles in the bottom left corner are yet covered. In order to reach this area, the agent are dependent on being guided towards it by nearby obstacles. In certain simulations agents may be starting in this location, causing the tiles within to be covered effortlessly. Yet, in other cases these locations can be difficult to reach for the agents. This causes the coverage of the environment to be highly inconsistent. Figure 4.6b depicts the same simulation, but with the inclusion of pheromones. The agents can be seen to cover the environment much more efficiently. The pheromones enable the agents to reach all areas of the environment irrespective of their starting position. While many tiles are still covered multiple times, the coverage of the environment becomes much more even.

4.4.2 Effect of Increasing Agent Population

When agents are acting solely on local sensory input, each agent constitute a completely independent entity. From a technical standpoint this promotes better scalability of the overall system. Given that the system scales well, it is important to assess the usefulness of such a scaling from a functional perspective.

To address this question the same experiment from the previous section was repeated, however the number of agents was increased to 12. The results of this experiments is presented in figure 4.7

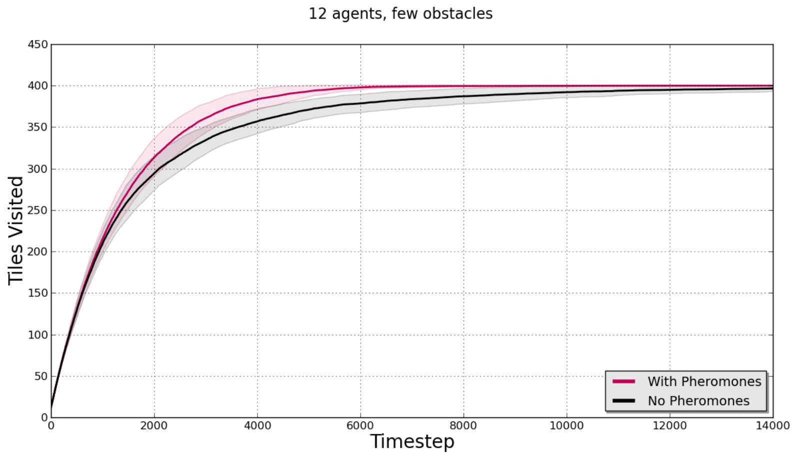

Figure 4.7: Complete coverage is achieved much sooner with more agents

After increasing the population to 12, there is still an advantage when using pheromones in terms of how quickly the environment is covered. After just 1,000 timesteps the coverage exhibited by agents not using pheromones start to subside, while the coverage of agents who are using pheromones continue to rise. This shows that when there is no indication of what areas have been previously visited, agents soon begin to retrace covered ground. Notice that in this graph complete coverage is obtained much sooner than it was with 3 agents. Agents using pheromones achieve complete coverage of the environment after 8,000 timesteps, while agents not using pheromones converge to complete coverage after around 14,000 timesteps.

To see the further effect of increasing the agent population the same experiment was repeated with 24 agents. The results are presented in the graph in figure 4.8.

While there is still a small advantage in coverage when using pheromones, it is apparent that the benefit decreases as the number of agents increase. Complete coverage is achieved after just 4,000 timesteps, irrespective of whether pheromones are in use or not. With the high redundancy in the agents, all areas of the environment are covered so quickly that there is hardly any room for the pheromones to offer any improvements. The effect of pheromones only comes into play when agents are about to retrace previously covered ground. If the environment is completely covered before this occurs, the pheromones become redundant.

As the density of agents is increased the coverage when using pheromones does not

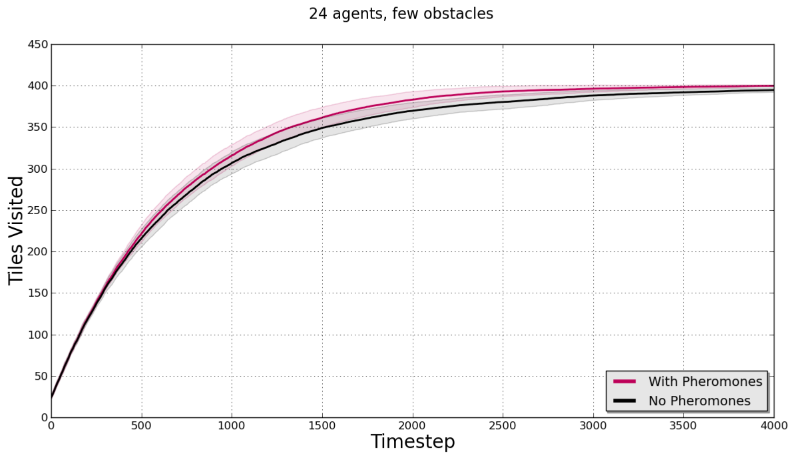

Figure 4.8: The benefit of pheromones decrease in large populations

see significant changes. Naturally, with more agents the environment is covered at a faster rate. However the coverage follows a similar pattern. Agents who are not using pheromones however see a much greater improvement as the size of the population increases. Presumably, since the agents themselves are also considered obstacles, higher density populations can inherently achieve a certain degree of dispersion. As the agents used in the simulations have very small bodies (5 pixels diameter), this effect is somewhat limited. Yet, it can still evoke changes in the movement of the agents, which may subsequently improve their coverage of the environment.

4.4.3 Influence of the Environment

The results of the previous sections suggest that the use of pheromones can indeed improve coverage, however that the effectiveness dissipates when the density of agents is increased. The rapid coverage elicited by the high number of agents is partly a result of the random spawn locations – Much of the environment is likely to be visited very quickly, regardless of future actions. This leaves very little room for the pheromones to evoke improvements in the exploration.

This raises the question of what will happen when the environment enables significant breathing room for the agents to be influenced by the pheromones. Once again using pheromones with a slow evaporation rate, figure 4.9 shows the coverage exhibited by 3 agents in an environment four times the size of the former experiments.

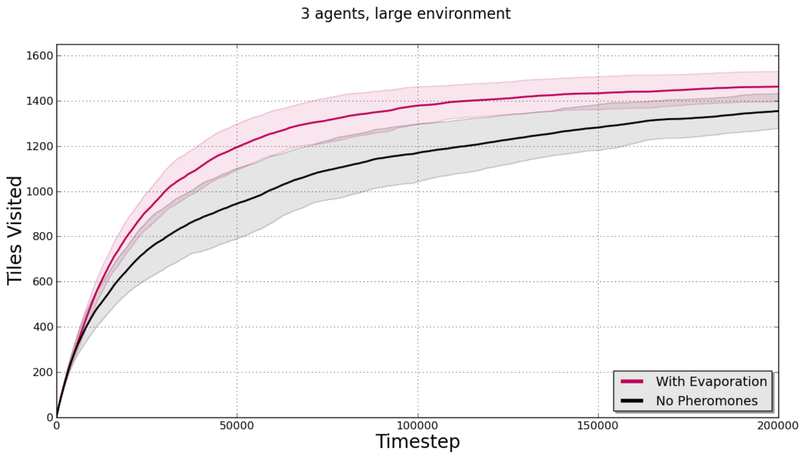

Figure 4.9: Large distances further increases the effect of pheromones

Because of the vast size of the environment, neither types of agent achieve the maximum coverage of 1600 within the allotted simulation time of 200,000. Still there is a clear distinction in how quickly the area is covered when pheromones are being used. After 50,000 timesteps, the agents secreting pheromones have covered an average of 1194 tiles. This accounts for 74.6% of the environment. At the same point in time the agents who are not using pheromones have only covered an average of 944 tiles, or 59.0% of the environment. After 200,000 timesteps the simulations are terminated. At this point the agents have achieved an average coverage of 91.4% with pheromones and 84.2% without.

Since the distances in the environment are so large, a turn towards a previously explored location will negatively affect the overall coverage to a much higher degree than before. When there are no cues to indicate that a location has already been visited, agents will unwittingly retrace covered ground for a much extended period of time. Therefore the reward for avoiding such situations becomes much greater.

To see whether an increase in agent count has a similar impact in large environments, another experiment was performed. While using the same environment, the number of agents was increased from 3 to 24. The results of this experiment is presented in figure 4.10.

The difference in coverage is indeed diminished, however not to the same extent as in the smaller environment. After 10,000 timesteps there is still a noticeable

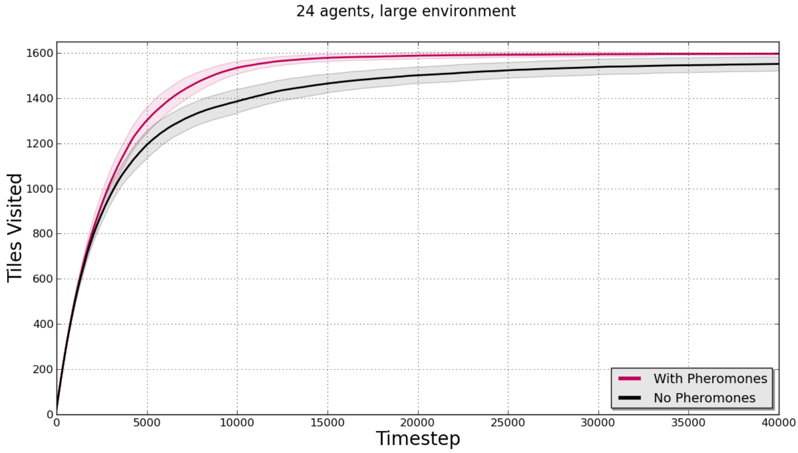

Figure 4.10: Complete coverage may never be achieved when all influences on movement are static.

difference in coverage. At this stage agents using pheromones have covered an average of 1533 tiles, or 95.6% of the environment. At the same time, agents not using pheromones, have covered an average of 1388 tiles. This accounts for 86.8% of the environment. Recall that when 24 agents were deployed in the smaller environment (figure 4.8), complete coverage was achieved by both types of agents after only 4,000 timesteps. In this larger environment, the coverage exhibited by the two types of agents does not start to coincide until after 40,000 timesteps.

Admittedly, as this environment is four times larger, the density of agents is naturally less than in the experiments of section 4.4.2. However the results reveal another peculiar property. While the agents using pheromones converge to the maximum coverage, the agents acting solely based on obstacles do not.

Whenever the coverage of an environment subsides it means that previously covered tiles are being revisited. However, since the measure of coverage make no distinction between the various tiles of the environment, it renders it difficult to determine why maximum coverage is not achieved without pheromones. Movement traces like the one in figure 4.6 is only able to capture the behaviour of the agents in a single simulation. In order to get a non-binary representation of how tiles are covered across a collection of simulations, heat maps were created. These highlight which tiles of the environment the agents are more frequently revisiting.

The heat maps are created by keeping individual counters for each tile. Whenever an agent is within the bounds of a tile, its counter increments. These values are normalized over the global minimum and maximum, before each tile is coloured from green to red according to the increasing time spent in each tile. The heat map of the simulation with 24 agents in the large environment is depicted in figure 4.11.

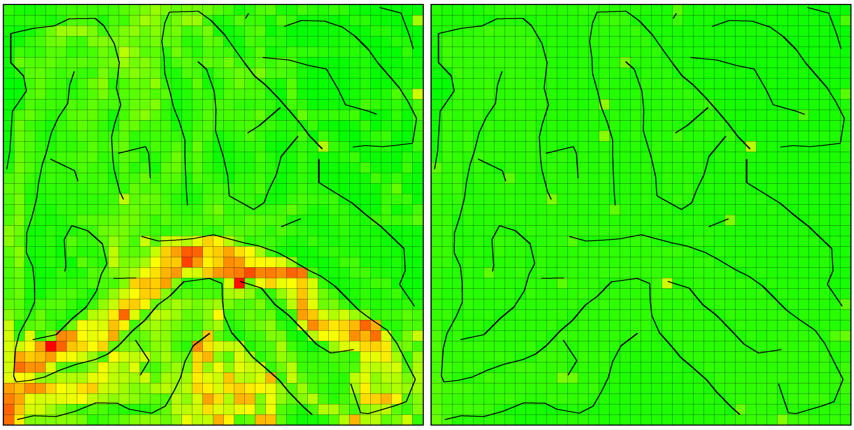

(a) Without pheromones (b) With pheromones

Figure 4.11: Red and yellow tiles are visited more often

It is apparent that when agents are only acting based on static obstacles they can lock themselves in repeating movement patterns. Without any additional influences on their movement, agents may continuously retrace their own steps. When using pheromones the redundant coverage of tiles reduces significantly. Because of the high number of times certain tiles are visited in figure 4.11a the global normalization process causes nearly all of the tiles of figure 4.11b to be coloured green – Only a few of the tiles are revisited a significant number of times.

In order to determine whether this behaviour is limited to one particular environment, a similar heat map was generated for the smaller environment containing many obstacles. This is depicted in figure 4.12. The simulations used to generate this heatmap was also performed with 24 agents.

Figure 4.12a shows that the agents tend to gravitate towards certain areas also in this environment. The presence of obstacles can cause the agents to repeatedly cover certain tiles. This type of behaviour may negatively impact the overall coverage of the environment – Any time spent visiting previously covered tiles,

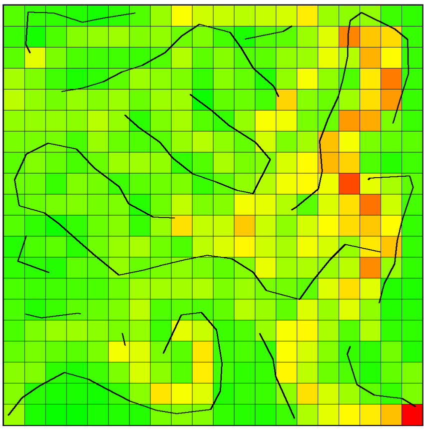

(a) Without pheromones

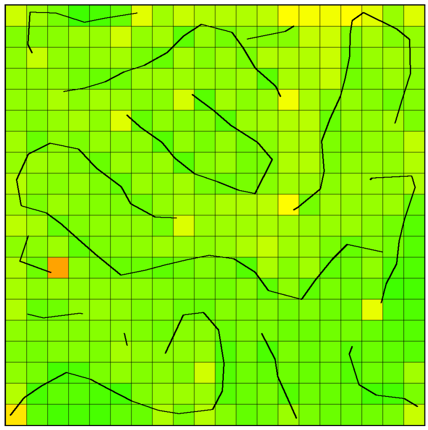

(b) With pheromones

Figure 4.12: Environments are covered more evenly with pheromones

is time that could have been spent exploring new areas. Figure 4.12b shows that when using pheromones the various tiles of the environment are covered much more evenly. The pheromones introduces a dynamic influence on the agents, making them less restricted by the appearance of the environment. This does not necessarily entail that the environment is covered at a faster rate. However, if agents are dispersed more evenly, it can increase the likelihood of new areas being covered.

4.4.4 Effect of Evaporation

To assess how the longevity of pheromone trails affect the agents in the various environments, experiments were performed with four different rates of evaporation. Figure 4.13 shows the difference in coverage between a slow and fast evaporation rate in the environment with many obstacles.

In this scenario the difference in coverage is not huge. As mentioned earlier, in order for pheromones to efficiently function as a collective memory of visited areas, the evaporation must be slow. This is however not achieved without side effects.

Consider the scenario of figure 4.14, where a single agent is moving towards a dead end. Since the passage is narrow the agent zigzags between the opposing walls as it enters. This causes the pheromone trail to block the agent's exit, forcing it back

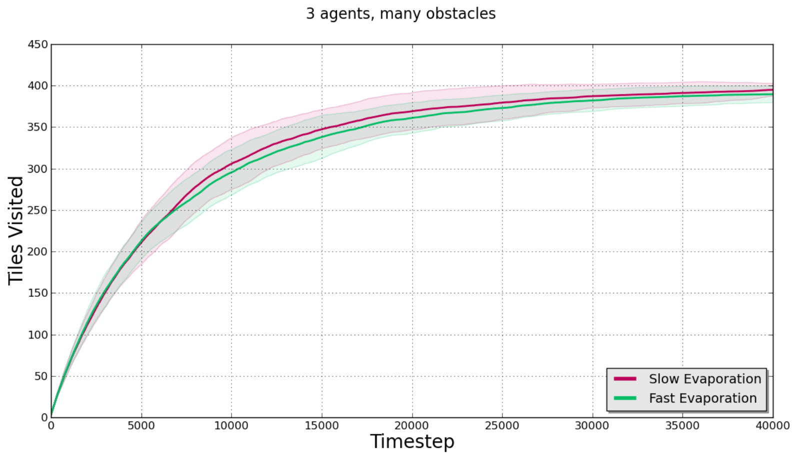

Figure 4.13: Coverage increases with slow evaporation rates

in towards the dead end. While this affects the overall coverage in a negative way, it may have some positive side effects. Recall from the end of section 3.2.1 that the pheromone model utilizes an average filter to create variations in intensity values. Because of the surrounding walls there are less zero valued cells that this average filter can propagate intensity values from. Therefore the pheromones will evaporate slower in confined spaces. The average filter eventually creates a gradient of intensity values towards the exit, allowing the agent to leave the area. When more agents are introduced they are discouraged from exploring this location for an extended period of time.

Figure 4.15 depicts the same scenario with a fast evaporation rate. Since the severity of reactions to pheromones is determined by its intensity, the agent is not forced to make any abrupt turns. This allows the agent to exit the area effortlessly. While the fast evaporation rate makes the agent less susceptible to its own pheromone trail, it sacrifices the prolonged effect of collective memory – The pheromones become limited to only influence agents that are in the vicinity. Despite its shortcomings, the effect of a slow evaporation rate proves more valuable in the long run.

Open environments

When there are few obstacles in the environment it was hypothesized that slow evaporation rates could lead to excessive amounts of pheromone, which would

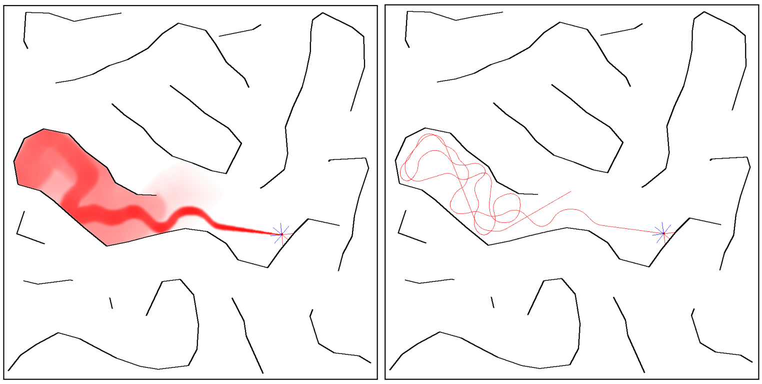

(a) Slow evaporation rate (b) Movement trace

Figure 4.14: Pheromones can slow down an agent's progression

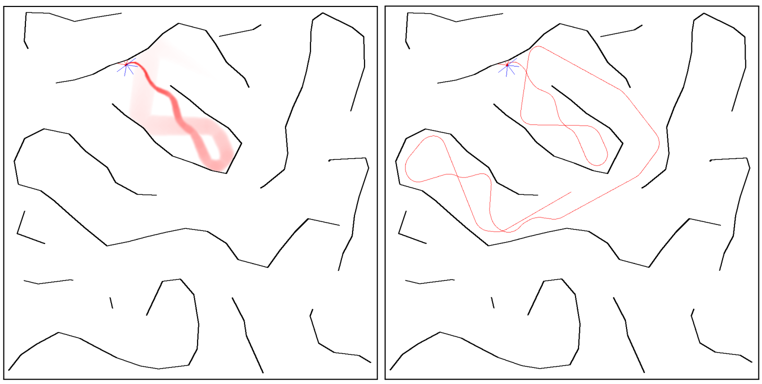

(a) Fast evaporation rate (b) Movement trace

Figure 4.15: Fast evaporation rates only causes small variations in movement

ultimately slow down the progression of the agents. The difference in coverage elicited with a slow and fast evaporation rate in an open environment is depicted in figure 4.16.

Contrary to the hypothesis, the results suggest that even when there are few obstacles in the environment, superior coverage is achieved with a slow evaporation rate. There are a few key factors that could explain this. The uncontained pheromone can diffuse to faraway areas, promoting a coarser dispersion of the agents. Even if the pheromones ultimately cover the entire environment, the

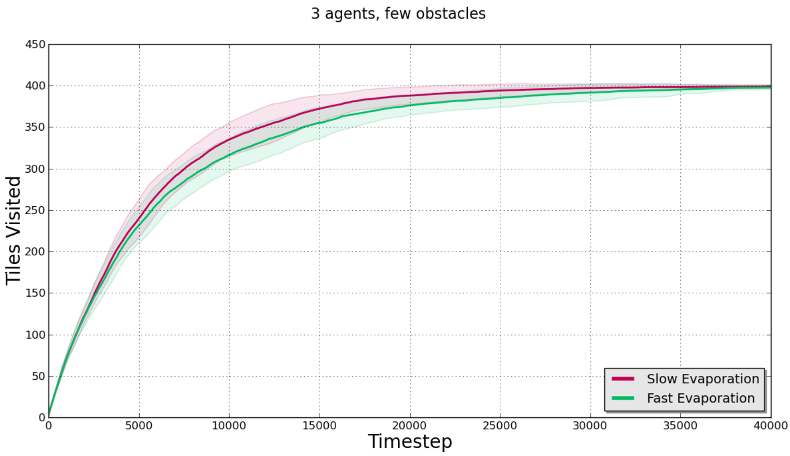

Figure 4.16: Coverage improves with slow evaporation even without obstacles to contain the pheromones

agents are not necessarily affected in a negative way. Similar intensity values will eventually equalize, at which point the agents will simply ignore the pheromones. The worst case scenario is that the behaviour of the agents is reduced to that of agents acting solely based on obstacle avoidance.

This worst case scenario is however unlikely. The only source of maximum intensity pheromones is directly beneath the agents. Furthermore, the intensity decays logarithmically. Thus, pheromones of high concentration rapidly lose their potency, ensuring certain degrees of variation in intensities. Therefore when there are excessive amounts of pheromone in the environment it merely reduces the total range of intensity values. However since the agents react according to maximum values, normal behaviour is retained as long the the evaporation is strictly decreasing.

To illustrate the effect of exaggerated amounts of pheromones, figure 4.17 depicts the coverage of 24 agents secreting pheromones with slow and fast evaporation rates.

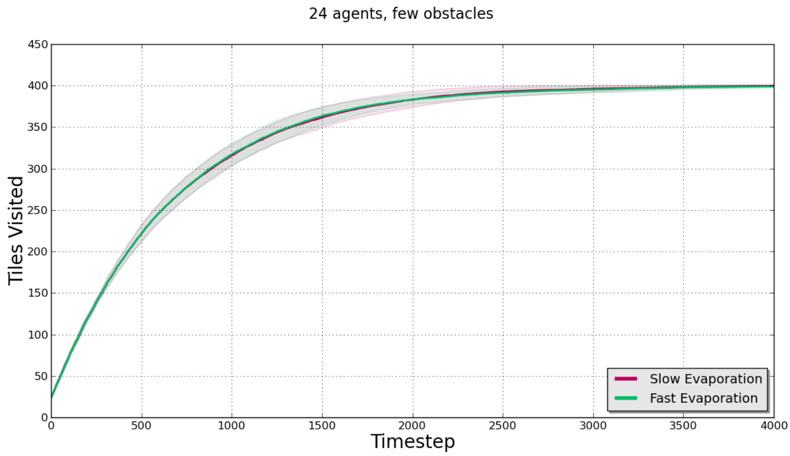

Figure 4.17: Excessive quantities of pheromone does not negatively impact the agents.

The lack of difference in the performance suggests that even when the amount of pheromones in the environment is excessive, it will not negatively affect the agents.

Chapter 5

Conclusion and Future Work

This chapter gives a general assessment of the findings of this project. Section 5.1 presents the conclusion of the experiments with respect to the research question. Following the main conclusion, several aspects of the presented work are elaborated and discussed. The contributions of this report is presented in section 5.2, before finally the future work required on this subject is described in section 5.3.

5.1 Conclusion

The following research questions was posed for this study:

> **RQ:** *Can the use of repellent virtual pheromones limit redundant exploration when robots are deployed in unknown environments?*

The experiments performed shows that the inclusion of repellent pheromones can indeed reduce the redundancy in search. Pheromones remain in the environment as a local collective memory of areas that have been recently visited. As pheromones diffuse they create a gradient of intensity values, guiding agents away from areas recently by themselves or other agents.Despite the experiments showing improvements in coverage when including pheromones, it also showed a similar decrease in the advantage of pheromones as the density of agents increased. When the agent count increases, areas are covered very quickly. High agent density may also evoke natural dispersion to some extent, as the agents themselves act as obstacles.

When agents are basing their movement almost entirely on static obstacles, the location of these obstacles dictate how the agents will move. In some cases the agents can be force to continuously retrace their own steps. Pheromones adds a dynamic influence on the agents that make them less dependent on the appearance of environment. Ultimately, the inclusion of pheromones may not only increase the speed at which the environment is covered, but also enable exploration of areas that would otherwise be unreachable.

Admittedly, while the pheromones shows signs of improvement, the rate at which the environment is covered is still far from optimal. Situations can arise where the pheromones actually counteracts the advancement of the swarm. Due to diffusion and evaporation such delays in the system are only temporary. Regardless, time is still wasted. In spite of such delays, the overall performance is superior when pheromones are present. Presumably, if such delays can be mitigated it will result in significant improvements. This is discussed further in section 5.1.3.

The experiments also showed a consistent improvement in coverage as the rate of evaporation was reduced. Since pheromones are only perceived locally they will necessarily have no effect until agents uncover them. If pheromones evaporate too quickly they will only be able to affect other agents in the vicinity. Additionally their ability to discourage agents from revisiting the same areas will be limited.

Contrary to the hypothesis, the positive effect of slow evaporation was most prominent in environments consisting of few obstacles and open areas. The absence of walls to contain the pheromone was believed to add restrictions to the agents' movement. However the results suggest that even when the pheromones extend to cover the entire environments, agents are not affected in a negative way. The unhindered diffusion actually promotes a wider dispersion of the agents. The absence of walls also results in more distinct gradient towards open areas, resulting in more efficient exploration.

In light of the experimental results, the following sections address various aspects of the presented system. This includes several shortcomings of the system that can be improved, as well as concerns with respect to future work.

5.1.1 Limitations of the Pheromone Model

An important property of pheromones is their ability to blend with each other. When trails intersect the pheromones should amplify each other's intensity, such that multiple weak trails constitute one potent. This property is absent in the presented pheromone model. Multiple attempts were made to incorporate this effect with cellular automata, but they were ultimately unsuccessful.

While the solution might appear as trivial as simply adding intersecting values together, the omnidirectional nature of traditional cellular automata makes this a difficult problem. Without any sense of direction, pheromones diffusing outwards will necessarily also propagate inwards on itself. If the values of neighbouring cells are not decreasing from one iteration to the next, the resulting values quickly spiral out of control.

To illustrate how quickly values propagate, consider the event of figure 5.1. Two distinct trails intersect, which are each secreted with an initial intensity value of 150. Every iteration the cellular automata adds the average intensity of neighbouring cells to the existing cell value.

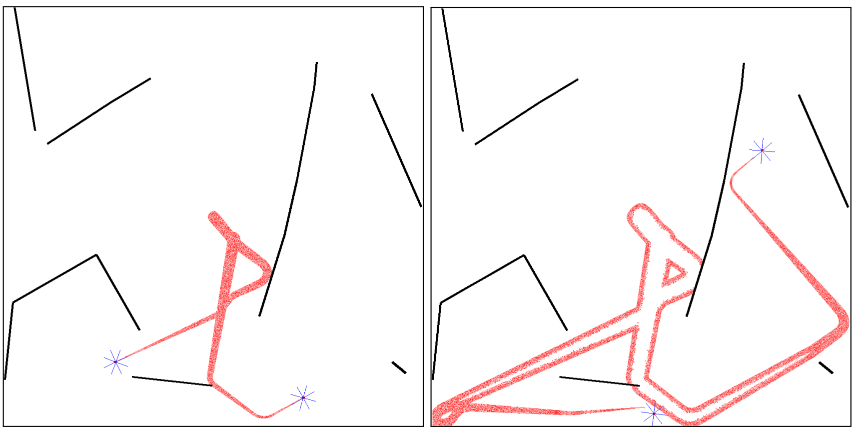

Figure 5.1: Intensity values propagate very rapidly when using cellular automata

As intended the pheromones continually diffuse outwards from its source. However since the cellular automata makes no distinction of direction, the cell values also propagate inwards. As a result the center cells of the trail rapidly increase in value. After just a few iterations the values of these cells have increased from the initial value of 150 to integer overflow (i.e. $> 2^{32} = 4.294.967.296$).

As a compromise, the pheromone model presented in this report propagates the maximum value of neighbouring cells. Rather than accumulating intersecting pheromone trails, the most potent assumes precedence. This enables the pheromone trails to continuously diffuse outward, without the added side affects of the cellular automata. An evaporation factor applied to each cell ensures a decrease in intensity values which confines the effect of this continuous diffusion. Finally an average filter evens out the distribution of the intensity values. While this model maintains a steady diffusion and evaporation of pheromones, it lacks the ability to blend two distinct trails. It is apparent in figure 5.2 that when pheromone trails intersect the trail of highest intensity is preserved at the expense of the existing one.

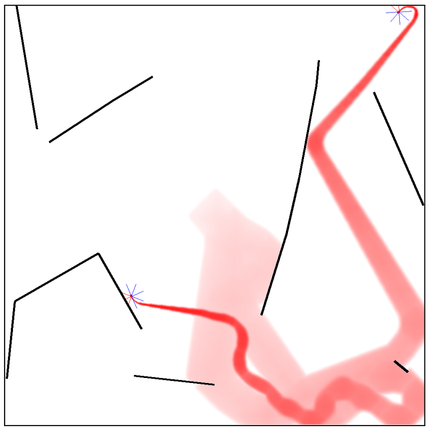

Figure 5.2: Pheromone trails of higher intensity overwrite those with lower.

Even though the average filter partially merges adjacent intensity values, multiple weak trails will never produce a strong one.

If a pheromone model that incorporates proper blending can be achieved efficiently it may be able to evoke more intelligent behaviour in the agents. Agents could release pheromone with weaker initial intensity, which makes them less susceptible to their own trail. Additionally, pheromones of high potency would only be present in areas traversed multiple times, rather than being reserved for areas directly behind agents.

5.1.2 Why Slow Evaporation Always Works

The initial hypothesis was that slow evaporation rates could cause the pheromones to diffuse beyond its intended area, thus limiting the agents' ability to explore the environment fully. The experiments suggest that this is not the case. Quite the contrary, pheromones with slow evaporation rates appear to consistently outperform the use of pheromones that evaporate faster.

When pheromones evaporate slowly they can indeed cause the ground to be completely secluded in pheromones, especially when there are few obstacles to contain the pheromones. However this does not seem to affect the agents very negatively. The key element to this is the robot controller and how the agents react to the pheromones. In order to prevent agents from getting completely entangled in their own pheromone trail, the reaction to pheromones was heavily relaxed. This relaxation causes the agents to ignore the pheromones if they acquire indistinguishable sensor input. Additionally, when agents react to pheromones, they only assess the maximum perceived intensity, which they attempt to avoid. Previous attempts at creating the control mechanism utilized a direct sensor-actuator coupling. This type of control mechanism works fine for avoiding obstacles since sensory input is seldom received across all sensors at once. Pheromones however, offer no physical restrictions on the agents. This makes the agents perfectly able to walk through the pheromones. This can cause the sensor input to be uneven and sporadic across the various sensors, resulting in erratic behaviour. Only assessing maximum intensity values resolved this issue.

Because of the pheromone model's inability to blend with other trails the only source of maximum intensity pheromones are directly behind agents. Additionally, because of the evaporation factor, the pheromone intensities throughout the environment are strictly decreasing. The pheromones also decay logarithmically, such that high intensity pheromones rapidly lose its potency. Therefore even when the environment becomes completely covered in pheromones, there are still variations in pheromone intensities close to other agents. The high amounts of pheromones simply narrows the range of intensity variations, or in some cases cause the agents to completely ignore the pheromones.

5.1.3 Pheromone Stagnation

While the purpose of the pheromones is to help agents cover an environment more efficiently, at times it can have the opposite effect. Agents continuously secrete pheromones as they move, oblivious to their surroundings. Unknowingly agents may in fact temporarily prevent other agents from exploring vital parts of the environment. In certain situations agents may also trap themselves between obstacles and their own pheromone trail. Due to evaporation and diffusion such occurrences were never observed to be permanent, yet time is wasted that could have been spent exploring additional areas.

A pheromone stagnation mechanism was implemented as part of the robot controller, intended to prevent monotonous reactions to pheromone. The details of this mechanism was presented in section 3.3.2. While inclusion of this mechanism showed signs of improvement during preliminary experiments, it is far from robust. The most challenging problem is to efficiently and reliably detect stagnation. When relying solely on the local perception of pheromones it is difficult for the agents to assess whether their actions will have positive or negative effects in the future.

Efficiently resolving stagnation once detected can also be challenging. Since the pheromones are not physical constraints, the most prominent solution is to simply allow the agents to move through the pheromones. This can be easily achieved by merely disabling sensory input for a limited time. The duration for which sensor input is suppressed may however vary depending on the situation the agents are in.

If stagnation by pheromones can be detected and resolved, both consistently and efficiently the system will likely see a significant improvement in performance.

5.1.4 Viability of repellent pheromones

One of the benefits of having agents communicate indirectly through pheromones is to promote better scalability of the swarm. If agents are communicating directly, the number of inter-agent connections increases exponentially with the size of the swarm. As autonomous mobile robots often have very limited computational capabilities, such communication may be infeasible.

The experiments showed that usefulness of pheromones was greatest in swarms with a low population. While larger swarms were not negatively influenced by the pheromone, the overall benefit of the pheromone dissipated when there was a high density of agents. The inclusion of pheromones did however enable the agents to achieve a high degree of coverage, which was otherwise only attained by increasing the agent population. Repellent pheromones may therefore be a viable approach in situations where only a small number of robots is available, or the environment to be covered is large.

The simple nature of the presented pheromone technique is aimed at providing good scalability of the swarm. However, since the benefit was shown to be greatest in small swarms, approaches with more complex behaviours and more computationally demanding communication, may be feasible.

5.1.5 Applicability for physical robots

An obvious limitation of pheromones is the difficulty in representing them for physical robots. The robot controller of this study relies heavily on differences in intensity values. These variations are created by the average filter applied to the pheromones. Since this operation is delimited by walls, it effectively creates gradients toward open areas and locations not recently visited by other agents. Achieving this effect in a real world setting could be a challenge. Furthermore, the variations created by the average filter are often very small. Outside of controlled environments where sensors are susceptible to noise, reliably detecting similarly small variations can be a difficult task.

Multiple different approaches for physically representing pheromones are reported in the literature. In [13] virtual pheromones was simulated by using an overhead projector as proposed in [28]. Agents were top-mounted with luminosity sensors that detected differences in intensities in the projected image. While this technique was proven successful, it is difficult to recreate outside of controlled environments.

Representation of pheromones using chemicals was reported in [20]. The reported chemical diffusion bears resemblance to the pheromones in this project. However even when using airflow to boost the chemical diffusion, both the range and speed of this diffusion is significantly lower than that of the pheromones in the simulator.

As the usefulness of repellent pheromones is greatest with a low number of agents, inter-robot distribution of pheromone data could be a viable approach.

5.2 Contributions

This report has presented a technique aimed at reducing the redundancy in exploration when multiple robots are deployed in unknown environments. By releasing repellent pheromones during exploration, robots can indicate which areas have been visited, and promote future exploration to favour different locations.

Multiple simulations have been performed in randomly generated environments. These have accentuated the limitations of agents who are only influenced by static obstructions, and how the use of pheromones can offer improvements. The results have also shown that while the presented technique enables good scalability of the swarm, the use of pheromones offer the greatest improvement when the population density is low.

Several different pheromone evaporation rates have been assessed. The results have shown that slow evaporation rates present the greatest benefit. Pheromones that linger in the environment also pose a greater threat of slowing down the agents' progression. However, with appropriately relaxed reactions to pheromones the potential of such hazards are mitigated.

5.3 Future Work

Despite its merits, repellent pheromones still offer several pitfalls for the agents. In order to enable the full potential of repellent pheromones, future work will involve finding ways of benefiting from its strengths while mitigating the effect of its hazards.

The continuous secretion of pheromone was often observed to cause long delays in the system. This was especially true when the pheromones evaporated at a slow rate. Yet, the greatest benefit of pheromones was also observed when the pheromones evaporated slowly. If the reliability of the pheromone stagnation mechanism can be increased, the presented technique will likely see significant improvements.

A further assessment of the effect of evaporation would also be necessary. The slow evaporation rate, which consistently produced the best results, was considered an extreme case. It is yet unknown what the actual effect will be when pushing the evaporation rates to further extremes.

Future work will also involve addressing the rate of pheromone diffusion. In order to focus on the effect of evaporation, the diffusion rate was kept constant throughout the experiments. If variations in intensities can be accentuated through other means than the current average filter, higher rates of diffusion could be applied. This could potentially hasten the emergence of dispersion, as well as create more distinct intensity gradients for the agents to follow.

The current representation of the pheromones is incapable of properly blending distinct trails. If a model can be developed that efficiently incorporates this property, different behavioural traits could potentially emerge in the agents.

In its current state, the technique of use of repellent pheromones in exploration is limited to simulated environments. Multiple approaches to represent pheromones for physical robots are reported in the literature. However none of those found are easily applicable in a real life scenario. The ultimate usefulness of pheromones in exploration is dependent on the discovery of an efficient way to represent environmental cues outside of the laboratory.

Bibliography

[1] Marc Barthélemy and Alessandro Flammini. Modeling urban street patterns. *Physical review letters*, 100(13):138702, 2008.

[2] MaximA. Batalin and GauravS. Sukhatme. Coverage, exploration, and deployment by a mobile robot and communication network. In Feng Zhao and Leonidas Guibas, editors, *Information Processing in Sensor Networks*, volume 2634 of *Lecture Notes in Computer Science*, pages 376–391. Springer Berlin Heidelberg, 2003. ISBN 978-3-540-02111-7. doi: 10.1007/3-540-36978-3_25. URL `http://dx.doi.org/10.1007/3-540-36978-3_25`.

[3] Gerardo Beni. From swarm intelligence to swarm robotics. In Erol Şahin and William M. Spears, editors, *Swarm Robotics*, number 3342 in Lecture Notes in Computer Science, pages 1–9. Springer Berlin Heidelberg, January 2005. ISBN 978-3-540-24296-3, 978-3-540-30552-1. URL `http://link.springer.com/chapter/10.1007/978-3-540-30552-1_1`. 00170.

[4] Jannik Berg and Camilla Hauknes Karud. Swarm intelligence in bio-inspired robotics. Master's thesis, Norwegian University of Science and Technology, Norway, 2011.

[5] John G. Blitch. Artificial intelligence technologies for robot assisted urban search and rescue. *Expert Systems with Applications*, 11(2):109 – 124, 1996. ISSN 0957-4174. doi: http://dx.doi.org/10.1016/0957-4174(96)00038-3. URL `http://www.sciencedirect.com/science/article/pii/0957417496000383`. Army Applications of Artificial Intelligence.

[6] V. Braitenberg. *Vehicles: Experiments in Synthetic Psychology*. Bradford books : Psychology. MIT Press, 1986. ISBN 9780262521123. URL `http://books.google.no/books?id=7KkUAT_q_sQC`.

[7] Nicolas Bredeche, Jean-Marc Montanier, Berend Weel, and Evert Haasdijk. Roborobo! a fast robot simulator for swarm and collective robotics. *CoRR*, abs/1304.2888, 2013.

[8] Rodney Allen Brooks. A robust layered control system for a mobile robot. *Robotics and Automation, IEEE Journal of*, 2(1):14–23, 1986.

[9] Scott Camazine. *Self-organization in biological systems*. Princeton University Press, 2003.

[10] J.-L. Deneubourg, S. Aron, S. Goss, and J. M. Pasteels. The self-organizing exploratory pattern of the argentine ant. *Journal of Insect Behavior*, 3(2):159–168, March 1990. ISSN 0892-7553, 1572-8889. doi: 10.1007/BF01417909. URL http://link.springer.com/article/10.1007/BF01417909.

[11] A. Filipescu, I. Susnea, S. Filipescu, and G. Stamatescu. Wheeled mobile robot control using virtual pheromones and neural networks. In *IEEE International Conference on Control and Automation, 2009. ICCA 2009*, pages 157–162, 2009. doi: 10.1109/ICCA.2009.5410442. 00000.

[12] D. Floreano and C. Mattiussi. *Bio-Inspired Artificial Intelligence: Theories, Methods, and Technologies*. Intelligent robotics and autonomous agents. MIT Press, 2008. ISBN 9780262062718. URL http://books.google.no/books?id=hY76AQAAQBAJ.

[13] S. Garnier, F. Tache, M. Combe, A. Grimal, and G. Theraulaz. Alice in pheromone land: An experimental setup for the study of ant-like robots. In *IEEE Swarm Intelligence Symposium, 2007. SIS 2007*, pages 37–44, 2007. doi: 10.1109/SIS.2007.368024. 00038.

[14] Simon Garnier, Jacques Gautrais, and Guy Theraulaz. The biological principles of. *Swarm Intelligence*, 1(1):3–31, June 2007. ISSN 1935-3812, 1935-3820. doi: 10.1007/s11721-007-0004-y. URL http://link.springer.com/article/10.1007/s11721-007-0004-y. 00156.

[15] M. Giurfa. The repellent scent-mark of the honeybeeApis mellifera tigustica and its role as communication cue during foraging. *Insectes Sociaux*, 40(1):59–67, March 1993. ISSN 0020-1812, 1420-9098. doi: 10.1007/BF01338832. URL http://link.springer.com/article/10.1007/BF01338832. 00063.

[16] J. Hardy, O. de Pazzis, and Y. Pomeau. Molecular dynamics of a classical lattice gas: Transport properties and time correlation functions. *Phys. Rev. A*, 13:1949–1961, May 1976. doi: 10.1103/PhysRevA.13.1949. URL http://link.aps.org/doi/10.1103/PhysRevA.13.1949.

[17] Achim Lilienthal and Tom Duckett. Experimental analysis of smelling braitenberg vehicles. *environment*, 5:10, 2003.

[18] David Payton, Regina Estkowski, and Mike Howard. Pheromone robotics and the logic of virtual pheromones. In *Swarm Robotics*, page 45–57. Springer, 2005. URL `http://link.springer.com/chapter/10.1007/978-3-540-30552-1_5`. 00046.

[19] Janice L. Pearce, Bob Powers, Chistopher Hess, Paul E. Rybski, Sascha A. Stoeter, and Nikolaos Papanikolopoulos. Using virtual pheromones and cameras for dispersing a team of multiple miniature robots. *Journal of Intelligent and Robotic Systems*, 45(4):307–321, April 2006. ISSN 0921-0296, 1573-0409. doi: 10.1007/s10846-006-9038-4. URL `http://link.springer.com/article/10.1007/s10846-006-9038-4`.

[20] Anies Hannawati Purnamadjaja and R. Andrew Russell. Guiding robots' behaviors using pheromone communication. *Autonomous Robots*, 23(2): 113–130, August 2007. ISSN 0929-5593, 1573-7527. doi: 10.1007/s10514-007-9035-x. URL `http://link.springer.com/article/10.1007/s10514-007-9035-x`. 00027.

[21] E. J. H. Robinson, K. E. Green, E. A. Jenner, M. Holcombe, and F. L. W. Ratnieks. Decay rates of attractive and repellent pheromones in an ant foraging trail network. *Insectes Sociaux*, 55(3):246–251, September 2008. ISSN 0020-1812, 1420-9098. doi: 10.1007/s00040-008-0994-5. URL `http://link.springer.com/article/10.1007/s00040-008-0994-5`. 00014.

[22] John A. Sauter, Robert Matthews, H. Van Dyke Parunak, and Sven A. Brueckner. Performance of digital pheromones for swarming vehicle control. In *Proceedings of the Fourth International Joint Conference on Autonomous Agents and Multiagent Systems*, AAMAS '05, pages 903–910, New York, NY, USA, 2005. ACM. ISBN 1-59593-093-0. doi: 10.1145/1082473.1082610. URL `http://doi.acm.org/10.1145/1082473.1082610`.

[23] Thomas Schmickl and Karl Crailsheim. Trophallaxis among swarm-robots: A biologically inspired strategy for swarm robotics. In *Biomedical Robotics and Biomechatronics, 2006. BioRob 2006. The First IEEE/RAS-EMBS International Conference on*, page 377–382, 2006. URL `http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=1639116`. 00033.

[24] G. Silva, J. Costa, T. Magalhaes, and L.P. Reis. CyberRescue: a pheromone approach to multi-agent rescue simulations. In *2010 5th Iberian Conference on Information Systems and Technologies (CISTI)*, pages 1–6, 2010. 00000.

[25] O. Simonin, T. Huraux, and F. Charpillet. Interactive surface for bio-inspired robotics, re-examining foraging models. In *2011 23rd IEEE International*

Conference on Tools with Artificial Intelligence (ICTAI), pages 361–368, 2011. doi: 10.1109/ICTAI.2011.60. 00001.

[26] Christian Skjetne, Pauline C Haddow, Anders Rye, Håvard Schei, and Jean-Marc Montanier. The chirp robot: A versatile swarm robot platform. In *Robot Intelligence Technology and Applications 2*, pages 71–82. Springer International Publishing, 2014. doi: 10.1007/978-3-319-05582-4_6. URL `http://dx.doi.org/10.1007/978-3-319-05582-4_6`.

[27] Keith N. Slessor, Mark L. Winston, and Yves Le Conte. Pheromone communication in the honeybee (apis mellifera l.). *Journal of Chemical Ecology*, 31(11):2731–2745, November 2005. ISSN 0098-0331, 1573-1561. doi: 10.1007/s10886-005-7623-9. URL `http://link.springer.com/article/10.1007/s10886-005-7623-9`.

[28] K. Sugawara, T. Kazama, and T. Watanabe. Foraging behavior of interacting robots with virtual pheromone. In *2004 IEEE/RSJ International Conference on Intelligent Robots and Systems, 2004. (IROS 2004). Proceedings*, volume 3, pages 3074–3079 vol.3, 2004. doi: 10.1109/IROS.2004.1389878.

[29] Brian M Yamauchi. Packbot: a versatile platform for military robotics. In *Defense and Security*, pages 228–237. International Society for Optics and Photonics, 2004.

[30] Erol Şahin. Swarm robotics: From sources of inspiration to domains of application. In Erol Şahin and William M. Spears, editors, *Swarm Robotics*, number 3342 in Lecture Notes in Computer Science, pages 10–20. Springer Berlin Heidelberg, January 2005. ISBN 978-3-540-24296-3, 978-3-540-30552-1. URL `http://link.springer.com/chapter/10.1007/978-3-540-30552-1_2`.

Appendices

A Source Code

A.1 Random Environment Generator

```
from Tkinter import *
import random as r
import Image, ImageDraw

class Point:
  def __init__(self, x, y):
    self.x = x;
    self.y = y;
    self.visible = 1;

  def insideCircle(self, centerPoint,radius):
    squareDist = (centerPoint.x-self.x)**2 + (centerPoint.y-self.y)**2
    return squareDist < radius**2

  def removePoint(self, canvas):
    self.visible = 0;
    canvas.create_oval(self.x, self.y, self.x, self.y, outline='white')

class Vector:
  def __init__(self, x, y):
    self.x = x;
    self.y = y;

def getVector(p0, p1):
  return Vector(p1.x-p0.x, p1.y-p0.y)

def averageVector(vectors):
```

```
  x = 0
  y = 0
  i = 0
  for v in vectors:
    x += v.x
    y += v.y
    i +=1

  return Vector( round(x/i), round(y/i) )

def drawLine(point, vector, canvas, lineWidth):
  w.create_line(point.x, point.y,
                point.x + vector.x,
                point.y + vector.y,
                width=lineWidth)

  #Draw the same line in memory
  draw.line((point.x, point.y,
             point.x + vector.x, point.y + vector.y),
             fill=(0,0,0), width=lineWidth)



def findNextSeed(seed, jumps, canvas, draw, lineWidth):
  vectors = []
  for p in points:
    if (p.visible == 0):
      continue
    if( p.insideCircle(seed, searchRadius) ):
      p.removePoint(canvas)
      vectors.append( getVector(seed, p) )
  if (len(vectors) != 0 and jumps != 0):
    v = averageVector(vectors)
    seed.removePoint(canvas)
    drawLine(seed, v, canvas, lineWidth)
    findNextSeed(Point(seed.x + v.x, seed.y + v.y),
                 jumps-1, canvas, draw, lineWidth)



#### PARAMETERS ####
canvasWidth = 1000
canvasHeight = 1000

noofPoints = 2000
```

```
noofSeeds = 10

searchRadius = 200

maxSeedJumps = 10

points = []
seeds = []
master = Tk()



###################
###### START ######
###################
w = Canvas(master, width=canvasWidth, height=canvasHeight, bg='white')
w.pack()

#create an empty image in memory
white = (255, 255, 255)
img = Image.new("RGB", (canvasWidth, canvasHeight), white)
draw = ImageDraw.Draw(img)

#generate points
for i in range(noofPoints):
  p = Point(r.randrange(0, canvasWidth), r.randrange(0, canvasHeight))
  w.create_oval(p.x, p.y, p.x, p.y)
  points.append(p)
#generate seeds
for i in range(noofSeeds):
  p = Point(r.randrange(0, canvasWidth), r.randrange(0, canvasHeight))
  w.create_oval(p.x, p.y, p.x, p.y, outline='red')
  seeds.append(p)


# Find points within maxRadius of each seedpoint
# Average the vector between seed and points to create a new seed.
# Connect the seeds with a line

for s in seeds:
  jumps = r.randrange(1, maxSeedJumps+1)
  findNextSeed(s, jumps, w, draw, 5)

#delete all remaining points
for p in points:
```

```
  if (p.visible == 1):
    p.removePoint(w)
for s in seeds:
  if (s.visible == 1):
    s.removePoint(w)

#save image to working directory
filename = "env.png"
img.save(filename)

mainloop()
```