# RAGE AMONGST THE MACHINES: A BIOLOGICAL APPROACH TO REDUCING SPATIAL INTERFERENCE IN MULTI-ROBOT SYSTEMS

by

Mauricio Zuluaga
Systems Engineer. Universidad de los Andes, 1995

A THESIS SUBMITTED IN PARTIAL FULFILLMENT
OF THE REQUIREMENTS FOR THE DEGREE OF
MASTER OF SCIENCE
in the School
of
Computing Science


SIMON FRASER UNIVERSITY
Summer 2005

# APPROVAL

**Name:** Mauricio Zuluaga

**Degree:** Master of Science

**Title of thesis:** RAGE AMONGST THE MACHINES: A biological approach to reducing spatial interference in multi-robot systems

**Examining Committee:** Dr. Binay Bhattacharya, Professor
School of Computing Science
Chair

Dr. Richard Vaughan, Professor
School of Computing Science
Senior Supervisor

Dr. Greg Mori, Professor
School of Computing Science
Supervisor

Dr. Lawrence M. Dill, Professor
Biological Sciences
Examiner

**Date Approved:** June 21, 2005

# Abstract

Interference is a common problem for animals, and can be characterized as a competition for different types of resources such as food, mates or territory. In the case of multi-robot systems, similar problems arise. Many species have evolved aggressive displays as a more efficient alternative to physical combat to solve conflicts over resources.

This thesis considers a transportation task in which a team of robots with no centralized control frequently interfere with each other. This thesis describes two new, principled approaches to selecting an aggression level, based on a robot's investment in a task. The methods are economically rational.

Simulation experiments in an office-type environment and a smaller-scale real world implementation show that under some special circumstances, the methods are able to significantly improve system performance compared to a similar competition with a random outcome.

# Reader's Summary

Interference is a common problem for animals; it can be characterized as a competition for different types of resources such as food, mates or territory. In the case of multi-robot systems, similar problems arise; robots may compete for a charging station or use of a shared tool or sensor.

A solution often seen in nature, when conflicts over resources arise, is fighting. But physical combat is potentially costly for the individual agents involved and to the overall system as robots may become damaged. Similar costs apply in nature, and many species have evolved aggressive displays as a more efficient alternative of physical combat to solve interference problems.

Spatial interference can reduce the effectiveness of teams of mobile robots. A team of robots with no centralized control performing a transportation task, in which robots frequently interfere with each other, is examined. The robots must work in the same space, so territorial methods are not appropriate. In [57] it was shown that a stylized competition inspired by aggressive displays in various animal species can reduce interference and improve overall system performance. However, none of the methods previously devised for selecting a robot's 'aggression level' performed better than selecting aggression at random.

In our methods, when robots come into competition for floor space, each selects an *aggression level* and the competition is resolved in favour of the more aggressive. This thesis describes two new, principled approaches to selecting an aggression level, based on a robot's investment in a task.

Simulation experiments with teams of six robots in an office-type environment show that under some special circumstances, a *global investment* method that takes into account the effort the agent has put into finishing a task, is able to significantly improve system performance compared to a random competition. A second method based on the concept of *local*

*investment* improves upon the first, by making the aggression level of a robot proportional to the time it has spent inside an area where spatial interference is likely.

The local investment method is evaluated and shown to be effective in a real-world robot implementation.

Finally, benefits and limitations of the methods proposed as well as future directions of research are discussed.

*Para mi gordis,*

*"... y eran una sombra larga ... "*

*"I've seen things you people wouldn't believe.*
*Attack ships on fire off the shoulder of Orion.*
*I watched C-beams glitter in the dark near Tannhauser gate.*
*All those moments will be lost in time, like tears in rain.*
*Time to die..."*

— *Blade Runner*, RIDLEY SCOTT, *1982*

# Acknowledgements

First, I would like to thank Richard Vaughan, for the invaluable contributions and time spent in the development of this work. Also, all the members of the thesis committee for their observations and comments. To all my colleagues of the Autonomy Lab for their suggestions and critiques, especially Sarah and Carl, for sharing the pain and success of the early work. I wish them all luck in their research and personal endeavours.

Thanks to Julia Vaughan for her help in making this thesis sound and read better.

Finally I would like to express my gratitude to Binay Bhattacharya who trusted me and gave me the opportunity of being part of this program.

A Sonia y Dario, mis muy queridos padres, por la intervención telepática de sus consejos invisibles. Y a quien pronto llegará, por la felicidad que anticipa.

Vancouver June 15th 2005

# Contents

# List of Tables

# List of Figures

# Chapter 1

# Introduction

## 1.1 The Autonomy Lab

The Autonomy Lab is part of the Computing Science department at Simon Fraser University. Its goal is the increase of the autonomy (i.e. self-direction and self-maintenance) of robots and other machines.

An overview of the principles and mission of the lab can be found at `http://autonomy.cs.sfu.ca/` and is included in this thesis as a frame of reference:

> "There are two main reasons to study autonomous machines:
>
> 1. Scientific: an autonomous machine is a concrete test of hypotheses about mechanisms of intelligent behaviour. We come at this from two directions: a bottom-up approach examining the necessary and sufficient conditions for rational behaviour; and a top-down approach realizing and testing models of human and animal intelligence.
> 2. Economic: by definition, autonomous machines can do more work than those which require human supervision. Autonomous robots could make society more efficient by enabling new kinds of industry, science and exploration.
>
> The acquisition and management of resources such as energy and space is a fundamental, unavoidable task for all living things. From an ecological perspective,

intelligent behaviour can be seen as rational manipulation of resources. This observation underlies our approach to building autonomous systems" [33].

## 1.2 The RAGE Project

RAGE (Robot AGgression Experiments) is a project of the Autonomy Lab (`http://autonomy.cs.sfu.ca/projects/rage/`). Its goal is to develop scalable techniques to reduce destructive interference in groups of robots.

RAGE makes use of aspects of animal behaviour that can be incorporated into multiple-robot experiments. Aggression displays are used by animals to signal and help in their conflict resolution over different resources like food, mates, or territory. In general aggression displays are evolved behaviours that enable animals to resolve conflicts while avoiding the costs of physical combat [40, 32].

Previously [57] showed that a simple stylized competition can reduce interference in a multiple-robot transportation task. The resource for which the robots compete is work space, i.e. the right of way in a narrow corridor where only one robot can pass. The RAGE project extends this work to investigate economically rational strategies for deciding an individual's level of aggression such that the system-wide performance is increased [47].

## 1.3 Goal

Animals, even the simplest ones, are examples of successful autonomous agents. Nowadays it is a principle accepted by the robotics community that the replication of animal behaviour can be a powerful tool for the development of adaptive, robust and autonomous robots.

[57] showed that, through the use of a stylized competition inspired in animal fighting displays, a team of autonomous robots could reduce robot-to-robot interference and thus achieve more work in a resource transportation task. It was found however, that none of the techniques used to select the level of aggression of a robot performed better than using randomly assigned aggression.

Consider a person controlling a team of robots with a global picture of the environment and exact location of the robots. She would be able to decide which robot should win the fight when two of them encounter each other going in opposite directions inside an area in which conflict is likely. The person would choose to let one of the robots win the fight either

(i) because there was an advantage for the system reflected in time, work or energy or (ii) because the outcome of the fight looked irrelevant to the overall system performance. We can say that the decisions made by the person in control are *rational* because by measuring aspects of the environment-robot system, conflicts for space are handled so the performance is increased. Additionally, it is likely that in similar situations, the same decisions are taken. Humans, however, have better or at least different sensing capabilities than any robot.

A non-rational method, like a random approach, could successfully handle conflicts in which there is no advantage for the system in one robot winning over the other. However, it would be incapable of dealing with the opposite situation and even more, of distinguishing between the two. We would like to have a team of robots that in a rational way can set their aggression using their sensing, without the need for communication or human assistance and that performs better than non-rational approaches.

In summary the goals of this thesis are:

- Find a rational aggression mechanism that is decentralized, independent of a navigation strategy, makes use only of existing sensors (sonar and laser), works in heterogeneous robots systems and is simple to compute.
- Show that a rational aggression mechanism outperforms a random scheme in a team of robots.
- Show that the robots using the aggression mechanism are autonomous and do not require human intervention.
- Demonstrate that a stereotypical competition, inspired by animal behaviour, can be implemented in the real world.
- Show that the results obtained in simulation are carried to the real world.
- Show that a rational aggression mechanism is robust and performs well in different world configurations.

## 1.4 Organization

The use of aggression as a mechanism to solve spatial interference problems was first proposed and demonstrated by Vaughan et al. in [57]. It has been extended by Vaughan and students at the School of Computing Science at Simon Fraser University. The Project comprised three MSc students: Mauricio Zuluaga, Sara Brown, and Carl Zhang.

The team worked together in [14] which is used in Section 3 of this thesis. Subsequently each student approached different extensions of the project that complemented each other and did not overlap. The central goal of the project was the demonstration of rational aggression mechanisms to increase the amount of work done by a team of robots working in the same space. This thesis describes the development of those ideas.

## 1.5 Thesis outline

### Chapter 2: Biologically Inspired Robotics

An historical introduction into the area of autonomous robots and the influence that biology has had in the design and construction of such robots is presented. Different robotics architectures are explained and related to the ideas of Maturana and Varela in the theory of autopoiesis.

Finally the utility of aggression is discussed and evidence is given to show its benefits in the development of agents and autonomous robots.

### Chapter 3: Rational AGgression Experiments RAGE

In [57] it was shown that a competition inspired by aggressive displays in various animal species can reduce interference and improve the performance of a system. When robots come into competition for space, each selects an 'aggression level' and the fight is solved in favour of the more aggressive robot.

This chapter presents a novel concept of 'global investment' as a way of rationally selecting the level of aggression of a robot in a resource transportation task.

Experiments done in simulation with teams of 6 robots in two different environments show the utility of the method. Robots using the global investment approach were able to complete more work than robots using random aggression.

Advantages and problems of the proposed method are discussed.

## Chapter 4: Local RAGE

This chapter extends and improves upon the work presented in Chapter 3 on the use of stereotypical aggressive display behaviour to reduce interference in robot teams, thus improving their overall efficiency.

A new method, 'local investment', for selecting the robot's aggression level is presented. The method performs better than any previous method and relies only on local sensor data.

Experiments in simulation and in the real world confirm the advantages of the local investment method over other methods. The experiments also show the first real world implementation of a stereotypical aggressive display behaviour to reduce spatial interference.

Adavantages and problems of the local investment method are discussed at the end of the chapter.

## Chapter 5: Discussion and Future Work

Several observations concerning the work in this thesis as well as directions of future research are presented. The following is a list of the most important ideas discused:

- The environment.
- The use of communication.
- Fights between different types of robots.
- A theory of mind (anticipation).
- Other forms of fight.
- A global fight.
- Territorial subdivision plus rational aggression.
- Evolution and the Theory of Games.

## Chapter 6: Conclusion

The conclusion compares the goals to the results. It is suggested that the hypotheses stated in Chapters 3 and 4 were supported by the results. Further, it is argued that aggression methods as a way to deal with interference problems can be used in the real world.

This thesis provides:

1. An introduction into the area of biologically inspired robotics.
2. Two different rational approaches to selecting the robot's aggression that make a team of robots complete more work in an environment where they have to share the working space.
3. The best method found so far for setting the aggression of a robot.
4. The first real world implementation of an aggressive display behaviour used in a resource transportation task.

Suggestions of possible extensions to the work in this thesis are also presented.

# Chapter 2

# Biologically Inspired Robotics

The first approach to artificial intelligence (hereafter AI), in what is now called classical AI, is in general based on representations of the world through symbols that are manipulated in some variant of logic or the predicate calculus. The use of classical AI in robotics could be summarized into building a model of the world and then using the model to prove or disprove some theories that later on are used to execute action commands. Sensory input is quickly forgotten inside the AI layers. For example, an image of an object in a camera would be replaced by a symbol representing it.

Shakey the robot [44] (Figure 2.1) is perhaps the most famous robot built with a classical AI architecture. Shakey has been highly influential in the robotics community and many of the tools and theories created in the project are still used in different areas of computing science, for example, the popular STRIPS planner [48].

Figure 2.1: Shakey the robot [26]. Printed by permission of SRI International ©.

The world where Shakey lived was, however, extremely simple. It was made of simple polyhedra; the colours of the walls and objects as well as the illumination of the rooms were carefully chosen and controlled to enable cameras to process and detect them. Even with all these simplifications Shakey was not able to work in real time. The time elapsed between receiving an order, planning and execution was too long, in some cases 15 to 30 minutes.

Shakey was therefore not robust for real world situations. It also suffered from the 'Frame Problem' [18]: all the objects in Shakey's world had names, and planning was done in a model of the world which used those names. This way of proceeding proved to be not scalable at all. There are just too many objects and situations in the real world to be named; in the case of Shakey, if acting in the real world, it would not know whether the climate in Bogota should be taken into consideration when deciding to move through an office door in a building in Vancouver, unless specifically told.

The failure of classical AI to provide robust and autonomous systems called for new research directions. These new ways would later be known as *the new AI*. The following sections present an introduction and description of these new ideas.

## 2.1 Behaviour based Robotics

The importance of the environment in the creation of autonomous robots is translated to a bottom-up approach in which, from the very beginning, robots are built and tested in the real world. In [17] Dennett asks "why not the whole iguana?", later on Brooks [13] starts talking about 'creatures'.

Creatures will provide ways of dealing with the implications of the real world and, ultimately, provide a method of knowing more about ourselves. This is the starting point of 'New AI' [11, 12, 13]. Creatures have the following properties: (i) They exhibit different behaviours from simple to complex. (ii) They have a body and live in the real world. (iii) They have to cope in an appropriate and timely manner within the environment. (iv) They have multiple goals (search for food, mark territory). (v) They have a purpose (reproduction etc). (vi) They are also robust to the environment (small changes in the world should not lead to the collapse of the creature) [13].

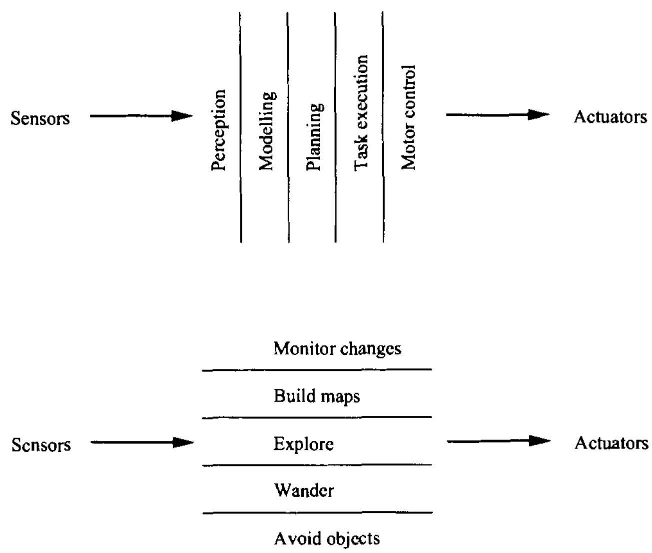


Figure 2.2: Classical Model (top) and The Subsumption Architecture (bottom) [11]. Printed by permission of Rodney Brooks ©.

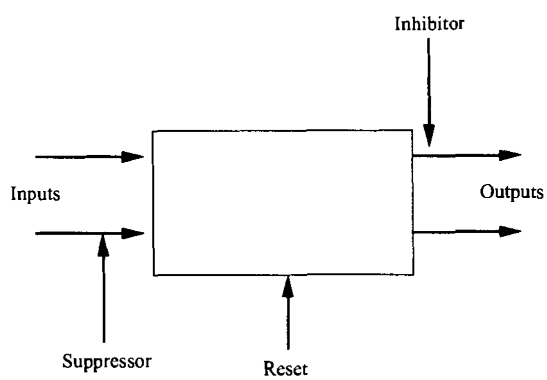


Figure 2.3: Subsumption Module [11]. Printed by permission of Rodney Brooks ©.

## 2.2 Subsumption Architecture (SA)

Brooks introduced the Subsumption Architecture in 1986 [11], as an alternative to the classical AI approach [44]. In classical AI a robot controller consisted of vertical layers of functional units. The subsumption architecture is a bottom-up approach in which the controller is made of layers of task-achieving behaviours (Figure 2.2). One crucial difference of the architecture compared to classical AI is that sensing and action were always available to any of the layers. Brooks proposed that one should start with the development and extensive testing of simple behaviours in the real world, and then start adding more complex ones. For example, a robot could have the following layers and behaviours: the first layer is obstacle avoidance, the second is wandering, the third is explore, the fourth is left wall following etc.

Each of the task achieving behaviours is made out of subsumption modules (Figure 2.3). These modules are connected and can inhibit outputs or suppress inputs. For example, a wandering module would inhibit the outputs of an obstacle avoidance module in order to make the robot wander and not only stay still.

Situatedness and embodiment are two of the characteristics stressed by the subsumption architecture. An agent is situated in the world and has a body. Intelligence is something that emerges from the robot's interactions with the environment.

In particular, in 'Elephants Don't Play Chess' [12], Brooks asserts that intelligence is based on physical grounding, and mainly it is an emergent phenomena caused by the interaction and coexistence of more simple behaviours. Classical AI has been searching for the *language of thought.* In general, this language of thought has been represented by some variant of the predicate calculus but one can only wonder what would be the language of thought in simple creatures like dogs and elephants or even bacteria. After all, many of their behaviours can be perceived as intelligent.

Brooks robots, built with the subsumption architecture, were able to execute simple tasks in real time in the real world, without human assistance; something that the 'classical AI' school had not been able to achieve up to that point. Robots like 'Herbert' [12] showed that through the interaction of simple behaviours more complex ones could be obtained. Herbert stole empty soda cans from an office environment. Still, after almost two decades, SA has not been as extensible as promised and the emergence of intelligence is still something not seen. In particular, behaviours that require anticipation have not been dealt with by

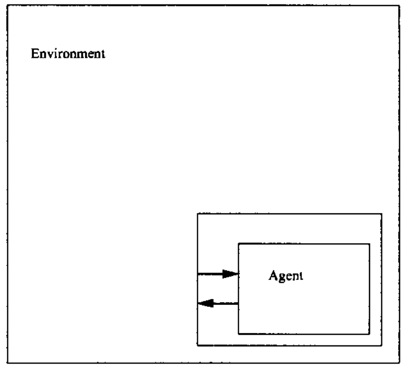


Figure 2.4: Coupling between Agent and Environment in DST [4]. Printed by permission of Randall Beer ©.

SA. Many researchers think that anticipatory or more complex behaviours require the use of representation and, therefore, the mix of classical AI and new AI approaches. They also say that SA is mainly good in the domain of reactive behaviours [30, 8, 22, 15].

## 2.3 Dynamical Systems Theory (DST)

Dynamical systems theory has its background in physics. Van Gelder presented a non-representational and non-computational account of cognition through the use of DST. He proposed that relations between agent and environment could be described as two coupled dynamical systems, described by a system of differential equations. In general, a dynamical system is described by a set of initial state variables and a dynamical law that says how those variables change through time [4].

Different robust and autonomous robots have been built using DST [7]. In DST, organisms and environment are the two parts of a dynamically coupled system (Figure 2.4). For example, when an insect lifts one of its legs, the other legs touching the ground feel more weight. Similarly to SA, embodiment and situatedness are of vital importance.

Modeling of complex dynamical systems is extremely difficult with DST in the same way that it is in SA. Behaviour is supposed to arise out of the coupling of the dynamical systems, but complex behaviours have not been successfully demonstrated; in general, there is a tendency to include representational elements together with DST for the generation of more complex behaviours [30, 8, 22, 15].

In the next section there is a presentation of a biological theory where the ideas exposed by Brooks in the subsumption architecture and by Van Gelder with the dynamic coupling of systems can live; these ideas are valid even though biology was not the main concern in the development of both architectures.

## 2.4 A Biological Theory of Living Organizations

The term autopoiesis (self-production) was invented by Maturana and Varela [38, 39, 54]. It came from an analysis of what is common to all living systems. Autopoiesis takes place when there is a circular organized network of interactions which maintains itself over extended periods, and which has a well defined border or membrane limiting the organization in space, and this border is also maintained by the organization. An autopoietic system is a network of component-producing processes in which the interaction of the processes generate the network itself, and also constitute the system as a differentiable entity in the space in which it exists.

In [34] Maturana et al. showed how the retina of the frog's eye did not send a complete copy of the image to the brain. Instead, what was sent was only some specific features like moving dots (the frog's prey) and large shaded edges (predators). Behaviour and cognition are consequently aspects of the living organization.

A fundamental property of an autopoietic system is *operational closure*, that defines the system as an entity that has no inputs or outputs. This property does not make the system independent of its environment, however, external events act only as perturbations that may trigger internal processes. The size of the internal changes caused by the perturbations depends only on the internal dynamics of the system. In [6], Beer presents a typical example of an autopoietic system, *the cell*:

> The paradigmatic example of autopoiesis is a cell, in which the components are molecules, the interactions are chemical reactions, and the cell membrane serves as a physical boundary that spatially localizes these reactions into an entity (or unity) distinguishable from its environment.

Cognition in the autopoiesis theory is a biological phenomenon. It is an ability that results from a *structural coupling* (the history of interactions) of the organism and its environment. Cognition is also embodied in the organism and therefore cannot be transferred from one to another. *Cognitive domain* is the actions that an agent can perform without ceasing to exist. For Maturana and Varela *nervous systems* are not a requirement for cognition. Any system capable of selectively interacting with its environment possesses a rudimentary cognitive domain, i.e. plants or cells. However, as discussed in [6], nervous systems enrich the cognitive domains of the animals that possess them by increasing the internal state that can be maintained, and thus the structural changes that can be handled. As a result, nervous systems expand the cognitive domain of an organism [30].

The autopoiesis theory does not forget about the environment in which an agent lives. The point is that the nervous system of the organism operates independently of the environment even when it is involved in a constant interaction with it. When the organism is perturbed by an environmental feature to which it is structurally congruent, the organism changes its dynamics in such a way that the perturbation is handled. To illustrate the difference between perturbations and inputs, the process of putting pressure in the throttle of a car does not set the car in motion. The foot's pressure perturbs the car in a way which causes it to move [30].

In summary, in autopoiesis theory, behaviour results from the interaction between an embodied system and the specific structure of the environment. The agent together with its environment become a single behavioural system. Behaviour, as such, will only be perceived by an external observer, and is invisible to the behaving system.

## 2.5 Autonomous Agents Research: animats

In the previous sections two different robot architectures, SA and DST, have been presented. The autopoiesis theory has been shown to be a possible biological ground where both architectures could live. Next is a small introduction into *Autonomous Agents Research*, which can be considered as an extension of the previous architectures. The goal of autonomous agents research is the modeling and creation of *animats*. An animat is a complete animal-like system, modeled and created by artificial means.

In [46] a summary of design principles for the development of animats is presented. The hypothetical creatures to be built are 'fungus eaters' whose mission is on a distant planet where no human intervention can be provided (Figure 2.5). They have to perform their tasks autonomously and maintain their energy supply. The animat approach is synthetic; "understanding by building".

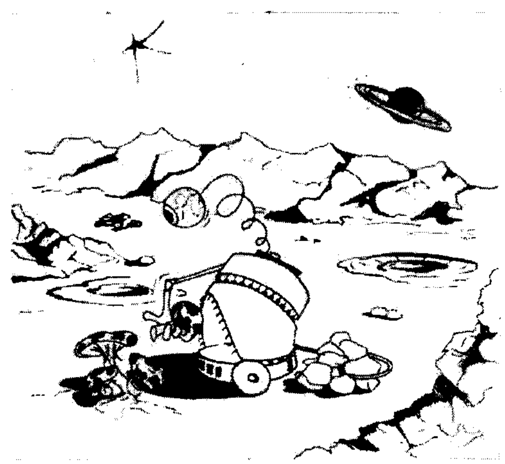

Figure 2.5: A "Fungus Eater" ingesting fungus on a distant planet. It has to perform its task autonomously while maintaining its energy supply[46]. Printed by permission of Rolf Pfeifer and the MIT Press (Book Title: "From Animals to Animats 4") ©.

Principles for the design of animats, adapted from [46]:

1. The "Complete System" principle: The agents are autonomous, self-sufficient, embodied and situated.

2. The "Ecological Niche" principle: The designer should know in advance the task that the creature has to accomplish and the place where those tasks have to be done. An underwater garbage picking robot will be different from a terrestrial one.

3. The "Parallel, loosely coupled Processes" principle: The agent's intelligence emerges from a large number of loosely coupled processes (for example as seen in the Subsumption Architecture). There is no central unit deciding on what to do next.

4. The "Value" principle: An autonomous and situated agent should have a way of evaluating what is good and bad for itself (the interpretation of value is only in the eyes of the designer, for the agent it is only about reflexes).

5. The "Sensorimotor coordination" principle: Interaction with the environment is through sensorimotor coordination. Memory, classification and perception should be seen as coordination rather than individual modules. There is evidence that recognizing an object implies re-enacting a sensorimotor coordination [43].

6. The "Ecological Balance" principle: There should be a match between sensors, actuators and the environment. Perception and Action mechanisms are tied and related to each other.

7. Principle of "Cheap Design": Cheap means having to capitalize on the system-environment interaction. (Compare a robot sheepdog that makes ducks sleep, then picks them up and moves them to a goal location vs. a sheepdog that manipulates ducks' behaviour to accomplish its goals [55]).

These design principles are only propositions and not statements written in stone. Still they capture very important ideas about the way artificial creatures should be created.

There are several objections to the principles that I would like to make: (i) Other architectures different than the subsumption architecture have proved to be equally powerful in the development of autonomous robots, for example, coupled dynamical systems [7]. (ii) Self-supervision and self-organization should not be strictly required in an agent; one could think of robots in which the goals are specific and everything is prewired so that learning is not necessary [57]. (iii) The sensorimotor coordination principle may not hold against social behaviours and emotions.

## 2.6 The Study and Use of Animal Behaviour Pays Off

The design of autonomous agents has proved to be quite hard to achieve. However, nature has a great variety ranging from bacteria to insects, to birds and humans.

Biologists have studied animal behaviour for thousands of years. There is reason to believe that though nature's solution to most problems that animals have to face during their life time may not be the best way, it is nevertheless robust and adaptive [7]. Nature itself has not created wheels or propellers, which seem to be very useful devices, yet it has created many other powerful mechanisms that we are still trying to understand.

Animals can interact with an incredible amount of different situations without collapsing or dying. In autopoietic terms, animals are able to handle a vast range of perturbations in an appropriate way, which means their cognitive domain is extensive.

Simple beings like cockroaches outperform state-of-the-art robots at navigation tasks, but recent robots built using animals as models have been successful at dealing in robust

Figure 2.6: Lobster robot at the Marine Science Center Northeastern University [37]. Printed by permission of the Office of Naval Research ONR ©.

ways with real problems. Some use 'bee' vision to navigate down corridors avoiding obstacles. Other cricket-like robots can track down male crickets by their songs [58, 35]. A robot lobster can follow an underwater chemical gradient to its source [37] (Figure 2.6). 'Rover' the robot sheepdog is able to make a flock of ducks move to a desired location [55].

Studies of animals have also helped in the creation of computational models that may inherit the robustness and adaptability of their natural cousins. In [2], a computational model is created that mimics recent findings about head-direction neurons in the brains of rodents. These so-called 'place-neurons' are only active when the subject was located at certain specific places in the world.

Vaughan [55] proved that it is possible to use computer simulations of animal behaviour (a flock of ducks) and build robust robot controllers that perform properly in the real world (sheepdog). More importantly he proved that animal behaviour can be manipulated by a robot to achieve some goals (Figure 2.7). In simulation, only one behaviour of the ducks was present, the flocking behaviour. Vaughan called the simulated creatures ducklets to stress the difference between them and the real ducks. Ducklets did not need to exhibit all the complex behaviours that ducks have. Vaughan's approach had advantages in time, cost and animal welfare. The robot sheepdog project was also the confirmation of a model of flocking behaviour in animals and therefore of potential interest to biologists. In [28, 29] Jakobi showed that it is possible to use simulation in the development of robust robot controllers. Though not in the same way as in [55], it still proves that simulation should not be discarded as a tool for the development of robust autonomous robots. This goes against some of the radical ideas stressed by architectures like SA or DST related to the situatedness and embodiment principles.

Holland and Melhuish [25] showed one of the first real-world examples of emergent

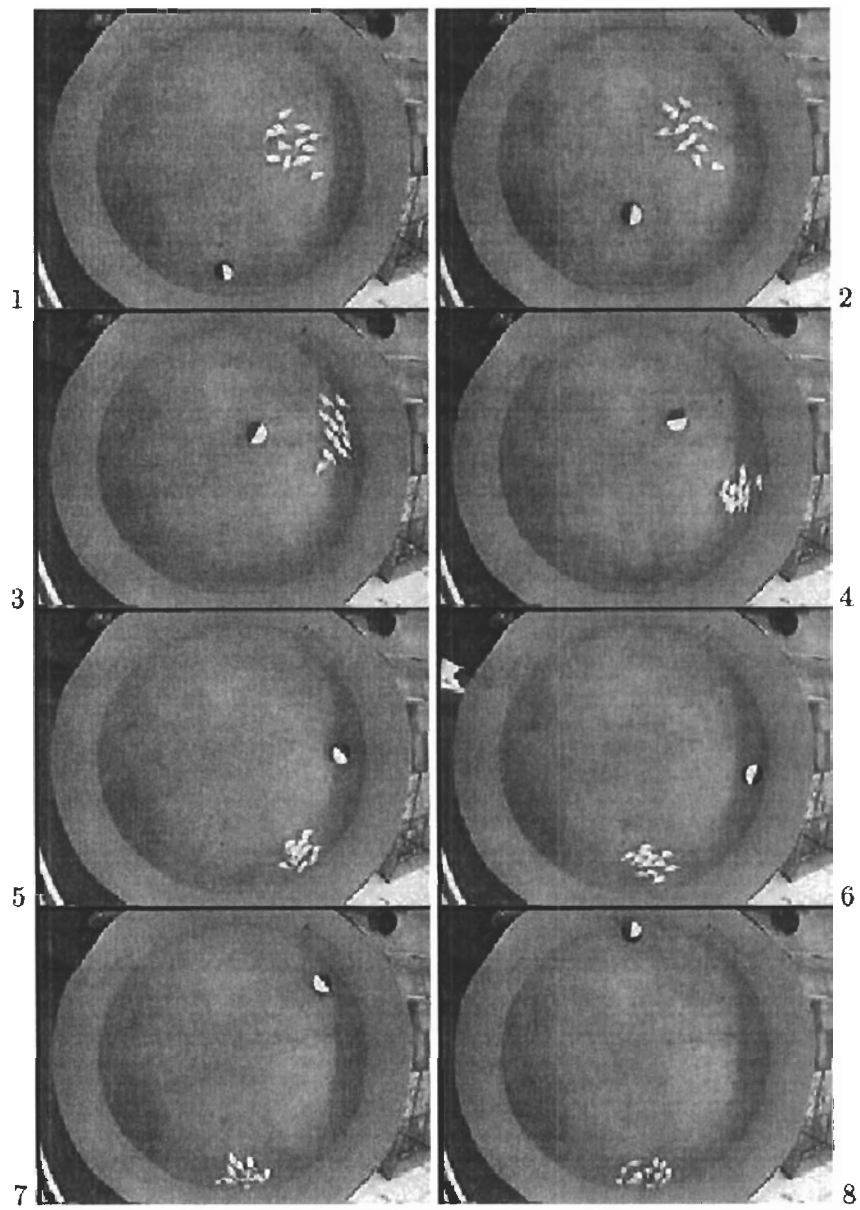


Figure 2.7: Sheepdog robot controlling and moving a flock of ducks to a goal location [55]. Printed by permission of Richard Vaughan ©.

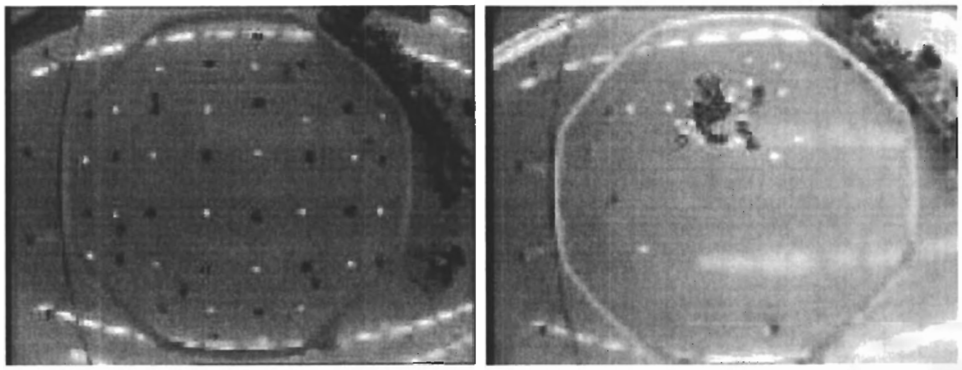

Figure 2.8: Emergent sorting of disks by termite-like robots [25]. Printed by permission of Owen Holland ©.

behaviour in a team of robots. They followed a biology inspired approach related to observations on termite building behaviour. The experiments showed how sorting appeared as an emergent behaviour of a group of termite-like robots (Figure 2.8).

In [7], Beer showed a biologically inspired approach to robot legged motion: a neural network modeling distributed gait control. Insects like ants can adopt new walking patterns when they loose one of their limbs. Beer's legged robots were able to deal with this type of problems plus others found in rough terrain. Note that Beer's legged robots (Figure 2.9) are quite similar to the MIT legged robots (Figure 2.10) but they were built following different architectures. While the MIT robots used the subsumption architecture, Beer's robots followed a dynamical systems approach. The behaviours, robustness and adaptability of the robots was very similar.

One recent concern in the robotics comunity is *social awareness*. We do not only want servants that do some of our boring daily tasks, we also want company and the possibility of social engagement. Our face is the most important organ of social communication. Some muscle-skeletal models have been incorporated from human anatomy and artificial agents that can mimic the gestures we make have already been built. However, that is only half of the problem, we also need robots that can react autonomously to our emotions and expressions. In [10, 9] a robot that conveys intentionally, 'Kismet', was developed using the subsumption architecture. Kismet exhibits infant-like responses through a motivation system consisting of drives and emotions (Figure 2.11).

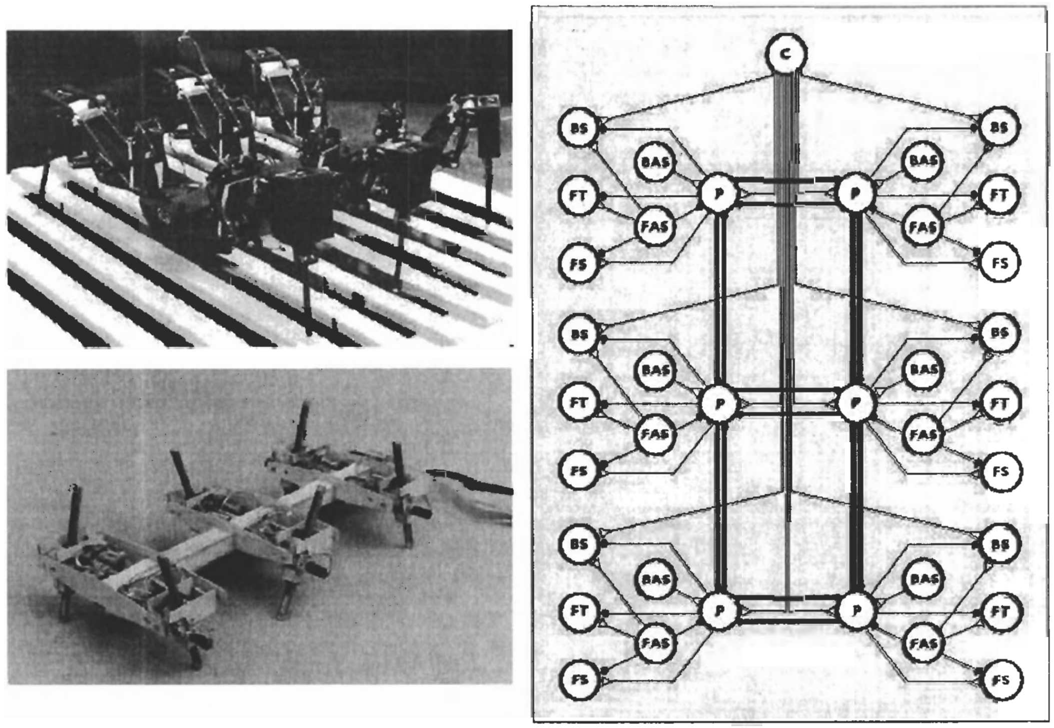


Figure 2.9: Beer's insect like robots, and Neural Network that controlls them [7]. Printed by permission of Randall Beer ©.

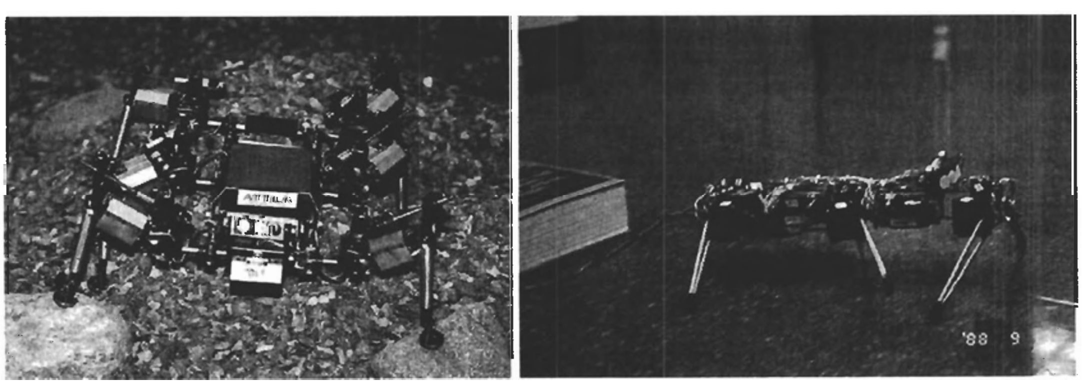


Figure 2.10: Attila and Genghis, two of the insect-like robots at MIT [50, 51]. Printed by permission of Rodney Brooks ©.

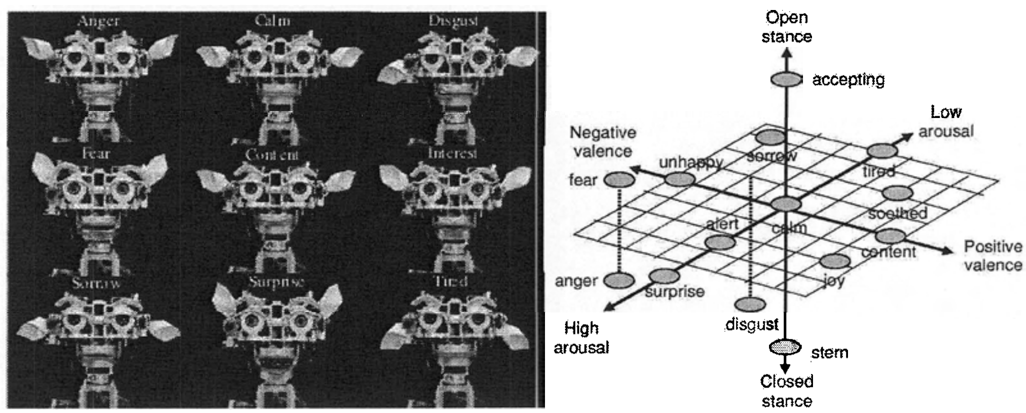


Figure 2.11: Kismet and its emotional space [9]. Printed by permission of Cynthia L. Breazeal and the MIT Press (Book Title: "Designing Sociable Robots") ©.

All the examples presented here demonstrated biology as a source of knowledge for the design and development of adaptable and robust robots. It also works the other way around for biologists as a method for understanding how animals behave and interact. After all, robots provide a way to test ideas and to check if they result in the generation of behaviour [58, 55, 30].

But why is it that the dream of affordable robots helping us in daily tasks like mail delivery or the delicate art of chopping vegtables are not here yet? Autonomous agents have proved to be hard to build. There are many reasons why this is the case. Below, I will enumerate the most important:

- The world has many constraints.
- Real life requires real time. AI programs are often computationally intensive and therefore run slow on the current available hardware.
- Compared to even simple creatures like insects, current sensing devices are extremely poor. The vision system of a bee, or the sense of smell of a dog are capabilities as yet not replicated. We may have sophisticated machines that outperform a dog in finding some type of odour, but making them small, and able to recognize the incredible variety of odours that dogs can, has yet to be achieved.
- It is hard to create complex behaviours, especially emergent and anticipatory behaviours. In the real world there are not many examples of robots that exhibit them.

The promises of complex and sophisticated behaviours arising from the interaction of very simple ones have not been fulfilled.

- The cost of sophisticated devices is prohibitive for mass production.
- Humans see the world through human sensors and not robot sensors. Because of this we are likely to incur design mistakes or just be incapable of understanding what is really going on.

Nevertheless we are in fact starting to see robots helping in some simple tasks (Figure 2.12). The best is yet to come.

## 2.7 The Utility of Aggression

The information presented in the previous sections shows that the use of animal behaviour as a model for the development of robust and autonomous systems is a good path to follow. As Beer says "it almost always pays off. While nature's way may not be the only way, or even necessarily the best way, time and again we have found unexpected benefits to paying close attention to the design of biological systems" [7].

In nature, many species have evolved aggressive behaviours useful for the resolution of conflicts over different resources (food, territory, mates, etc). Though physical fighting is one of the aggressive behaviours present in animals, its extremely high cost makes it a strategy of last resort. It is not worth fighting to the death for a sip of water, especially when it is almost guaranteed that water can be found somewhere else or at the same location but at a different time. Fighting not only wastes a lot of energy but may be the last thing that an animal does. Because of this, many animals have evolved aggressive displays that allow them to solve conflicts in a safer way (Figure 2.13). A display allows an animal to signal information about its own state (health, strength, etc.) and to perceive the same information from an opponent. In this way, two animals in a conflict can measure some properties of the opponent and decide if it is worth fighting for the resource. The threat of escalation to a genuine fight is usually implied, if the display is not enough to resolve the conflict [40, 32, 3].

The use of aggression in robotics has not been well studied. It has usually been used as a control parameter in evolution experiments. In [52] an evolution experiment with thousands of agents showed that the most successful ones were the most aggressive. In [20] it was found

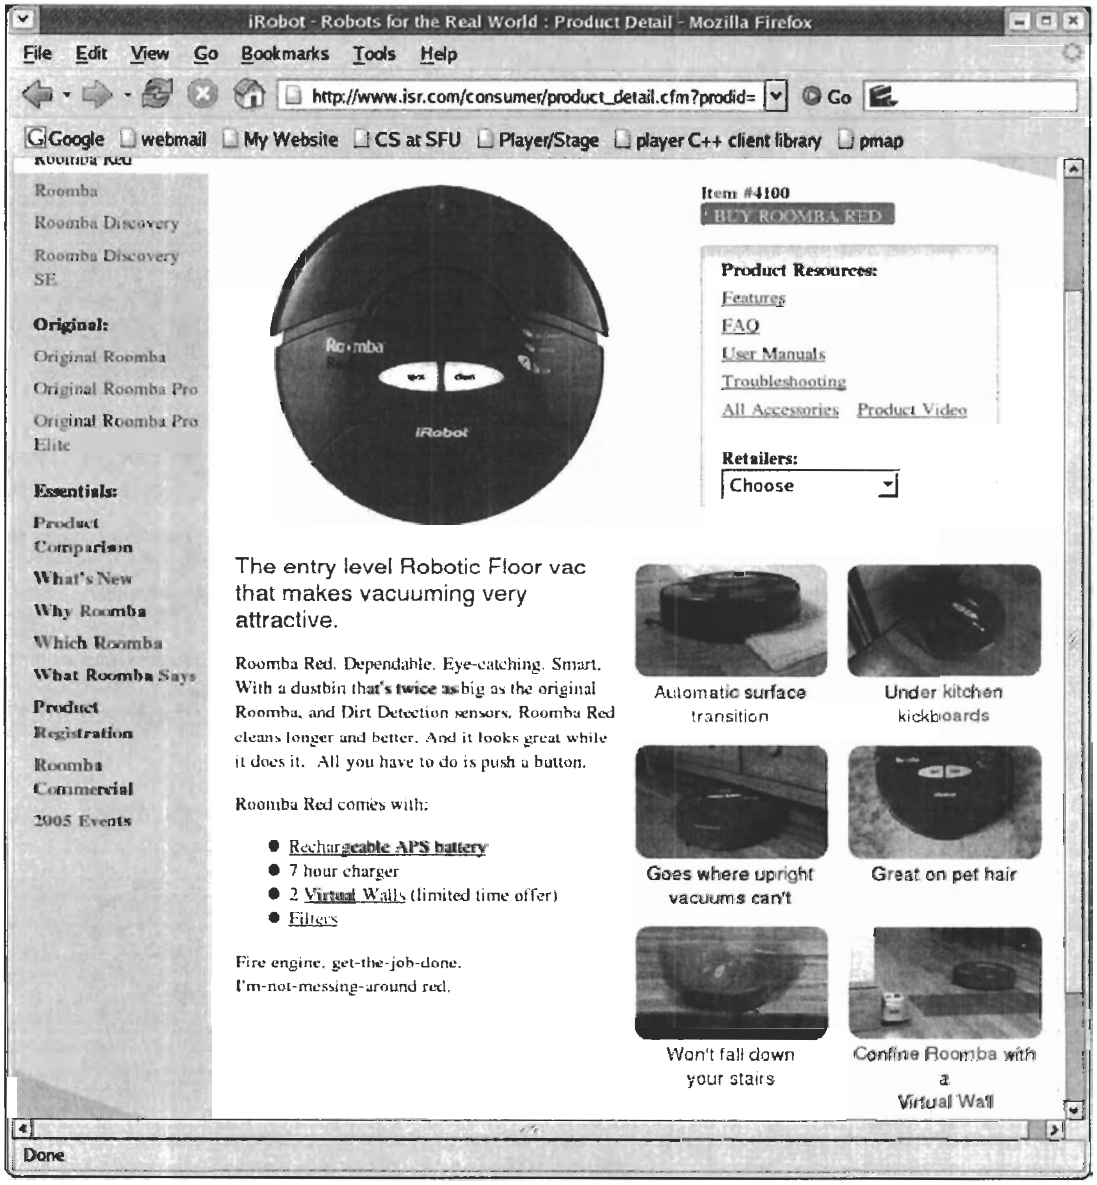


Figure 2.12: A commercial vacuuming robot [27]. Printed by permission of iRobot ©.

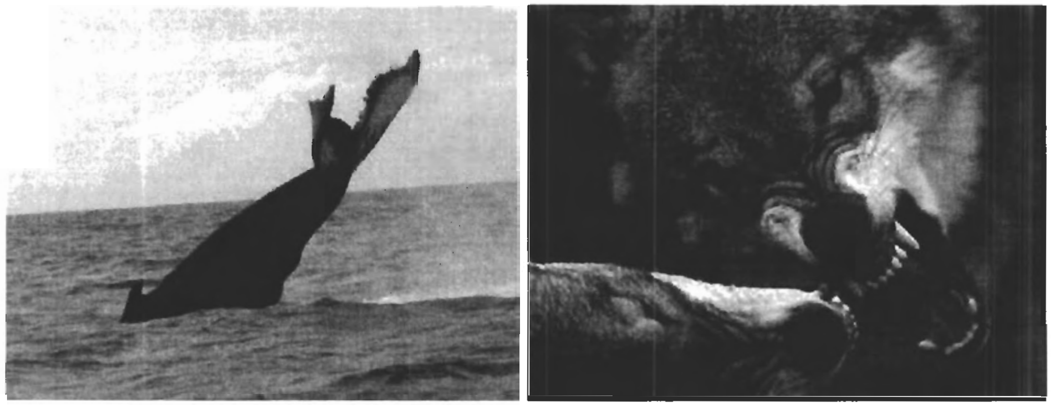

Figure 2.13: Aggressive display showed by different species of animals [60]. Pictures by Ken Bondy and Bo Holmberg ©.

that the aggression of an agent had a direct impact on the way the agent evolved. While prey (non-aggressive) evolved noisy controllers that moved at random, predators (aggressive agents) evolved directional controllers to improve pursuit behaviour.

In [57], animal-inspired aggressive behaviour is proposed as a way of handling the real life problem of 'spatial interference'. In the experiments a team of robots has to share a work area and complete a transportation task. An aggressive display mechanism in which the robots were able to signal their level of aggression to the other robots helped them achieve more work as a team.

The use of aggression presented in [52, 20] is different from the one in [57]. While in the first case the experiments presented seem to be designed to understand the way nature and evolution work, in the second the way nature works is already known and put to use for the benefit of a group of robots doing a resource transportation task.

The use of economical models has also made an impact on the study of animal behaviour for a good number of years. In [59], a cost-benefit study into the economics of fleeing from predators is presented. The results support the theory that prey does not immediately flee from predators when detected, but that there are other variables that affect the fleeing behaviour. Some of these include food availability, distance to a refuge, group size, sex, etc. In a similar study [19], it is shown how a worm can adjust hiding times in response to the availability of food. When the worm is not hiding it can feed but its vulnerability increases.

In the experiments it was shown that the hiding times of the worms changed with food availability and that worms can track short-term changes in it.

Conflict resolution, and consequently the use of aggression and fighting displays in animals, has also been extensively studied. Maynard-Smith was one of the pioneers of the use of game theory in the study of animal behaviour, especially fighting behaviours [40]. One of his proposed models is called *war of attrition*. It is a model in which conflicts are settled without escalation with the winner being the agent willing to pay the most (via display) for the resource. In general there are three ways of settling a contest over a resource without the need of fighting [31]:

1. By cooperation: Split the encounters. This may be used only when the contestants recognize each other and can keep track of their previous contests so that new ones can be split in a fair way. Unfortunately, this is not something that most animals can do, and even if possible, what if finding the same opponent again in the world is something unlikely?

2. By asymmetry: In this case an asymmetry, normally detectable in the fighting display will determine who wins and who losses. If the display is not honest, i.e. either of the agents are not displaying their 'real' state, it is possible that one agent may call the other's bluff. It is suggested that in this type of contest there has to be escalation, which, in the case of animals, means physical fight.

3. Waiting game or *war of attrition*: In this case the conflicts are settled by immediate withdrawal or by non-escalated display contests (Doves strategy in [40]). In a waiting game, the agent willing to wait the longest wins. One example of war of attrition found in nature is the case of territorial contests in damselflies (Calopteryx maculata) [36].

The capabilities of a robot plus the type of task or mission commissioned to it will define the conflict model to use. In conflicts solved by cooperation or waiting games the agents never fight, this is attractive because it guarantees the physical safety of the robots. Conflicts with asymmetries could be used but only if there are other means of escalation safe for the robots. Cooperative conflicts require the robots to be able to recognize other robots and to communicate with them. This is not scalable for groups of many robots. On the other hand, waiting games are very attractive from the scalability, safety and cost points of view, as robots would not require ID's, special sensors, or the need of physical combat. Waiting games may be inefficient if the waiting time is long. However, in nature one important characteristic of the randomness of the times animals wait before withdrawing from a contest is that it follows an exponential distribution. This causes the majority of conflicts to be solved quickly, while a small amount of them take longer, reducing the costs associated to the display behaviour.

In [57], the robots are clones that solve their conflicts by display and never escalate to physical fight. The contest model used by Vaughan et al. is an example of a waiting game. The work presented in this thesis is intended to extend on the use of aggression as presented in [57].

# Chapter 3

# Rational AGgression Experiments RAGE

The work presented in this chapter is based on work done in collaboration with Sarah Brown, Carl Zhang and Richard Vaughan [14], accepted to the International Conference on Advanced Robotics ICAR2005, Seattle, Washington, July 18th-20th, 2005; and [60], Last-Minute-Results Poster at the International Conference on Simulation of Adaptive Behaviour (SAB'04), Santa Monica, California, 23-17 July 2004.

## 3.1 Introduction

Spatial interference is a frequently encountered problem in multi-robot systems, especially those without centralized control, that can seriously degrade their performance.

Interference can be characterized as competition for resources. In animals these may be food, mates or territory; in robots these may be access to a charging station or use of a shared tool or sensor. Most commonly, robots simply block each other's way during normal navigation in the environment.

An example of the problem is getting two Pioneer robots (0.5m in diameter) through a doorway (0.8m wide) from opposite directions; a symmetry-breaking mechanism is required to decide who goes first. This is a real-world problem for robot applications such as mail delivery, factory and warehouse AGVs (autonomous guided vehicles), and assisted-operator wheelchairs. An example scenario is shown in Figure 3.1, where two robots are driving in

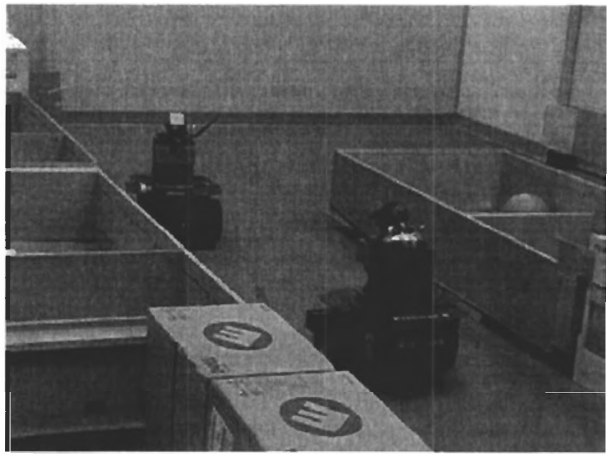

Figure 3.1: Spatial interference: two robots block each others' way in a narrow corridor.

opposite directions down a narrow corridor, blocking each other's progress.

In [57], the authors proposed the use of animal behaviour as a way to handle a real life problem 'spatial interference'. Because the robots had to work in the same space, territorial methods [21, 1] were not appropriate. An important advantage of the system was that communication of the 'aggression level' was performed using only existing sensors and actuators, there were no special-purpose sensors, no wireless communication and no need for unique identifiers for each robot. Thus, the method was found to be perfectly scalable and useful in heterogeneous systems, and even in human-robot interaction; humans can easily understand and manipulate the behaviour of the aggressive robots.

The symmetry-breaking provided by the aggressive competition was shown to produce better overall system performance, in terms of the number of transportation trips completed. compared to an otherwise identical system that lacked the aggression mechanism. Changing the behaviour of the robots in this way did not eliminate interference. In a typical one-on-one competition, the 'winning' robot certainly interferes with the immediate progress of the losing robot. Yet, as the overall system performance is increased, it can be said that the overall negative interference is reduced.

Several strategies for determining a robot's aggression level during a 'fight' were evaluated: random aggression; a linear dominance hierarchy; and a 'personal-space' method, where aggression was determined by the amount of free space visible to the robot. All methods were shown to have statistically similar performance. Neither the dominance hierarchy

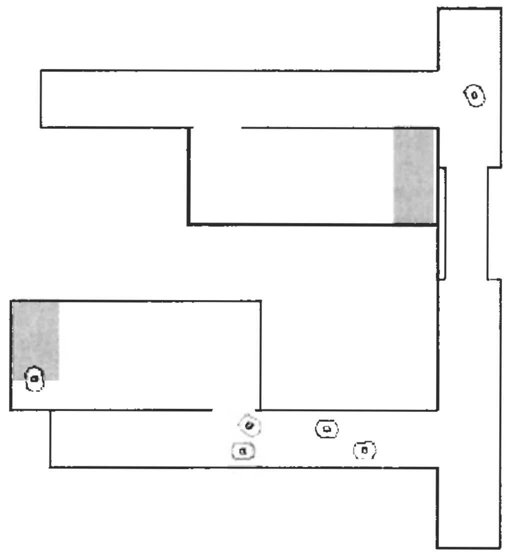

Figure 3.2: The simulated world containing six robots used in Experiment 1 and in[14]. It closely resembles the environment used in [57].

nor the personal space method offered any improvement over a random outcome.

This chapter describes a new, principled approach to selecting an aggression level, based on robot's investment in a task. The term 'investment' towards achieving a goal is fundamental in models of autonomy in animals [41, 59, 19]. Simulation experiments with teams of six robots in an office-type environment show that, under certain conditions, this method can significantly improve system performance compared to a random competition and a non-competitive control experiment.

## 3.2 Rational Aggression

To improve performance compared to a random outcome, the outcome of an aggressive interaction must reflect some relevant state of the world. As argued in [57], a hierarchy of robots with fixed aggression levels can not encode any information relevant to the outcome of a particular competition; when two robots meet at a doorway, their status in the hierarchy does not matter, so long as one of them gives way. Adding memory of past robot/robot interactions does not help; similar arguments apply to dynamic hierarchies.

In general, to control some parameter of a system, it must be measured or estimated

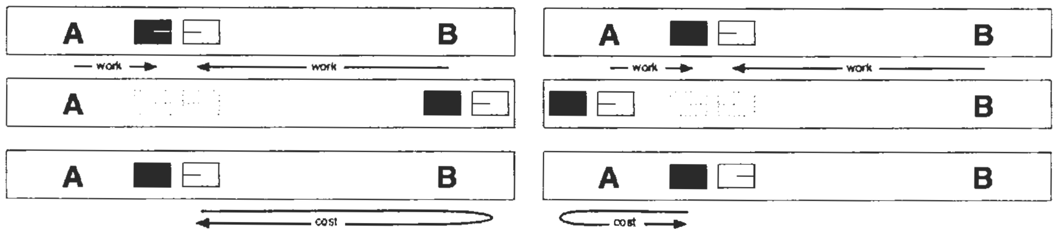


Figure 3.3: Motivation for the *global effort* strategy. The robot that has invested the most work in a task should win.

from its correlates and the environment. To maximize the amount of work done by the robot system, an estimate of how much work a robot is doing is needed as an input to the control system. This principle leads to the following economic approach to this problem.

Consider a system of two robots, Black and White, working in a narrow corridor as shown in Figure 3.3. They have the same task; transporting widgets from A to B at either end of the corridor. Assume that it is not practical for robots to transfer widgets between themselves. Black starts at A, White at B. At some moment, shown in the top row of the figure, Black and White block each other's progress. Assume the robots have an internal aggression level, and can perform a stereotypical behaviour sequence called a *fight*, in which each robot displays its aggression to the other. If a robot perceives that its rival has a higher level of aggression, it goes into retreat mode and allows the other robot to move forward. By performing a *fight*, Black and White can resolve their conflict; the more aggressive robot will push the other backwards and out of the way. From the point of view of system efficiency, which robot should be more aggressive?

### 3.2.1 Investment Aggression

In Figure 3.3, the arrows beneath the robots indicate how far the robots have travelled towards their goals. This travel inevitably has real cost in terms of time, energy and computation. These are *sunk costs*; they can not be recovered. In the left column, Black wins the fight and pushes White along the corridor until Black reaches its goal (middle row). Then Black switches to goal A, and proceeds down the corridor followed by White. At some point (bottom row), White is now back where it started to fight, after travelling the distance indicated by the arrow. The cost of the fight is the sum of White's sunk cost plus the cost to get back to its start position. The right column shows the outcome if White wins the fight. The steps are the same, but the total cost of Black losing the fight (total length of arrows) is much smaller. In this thought experiment, the robot with the higher sunk costs should be more aggressive as it has more to lose. With this scheme, the system will achieve more trips from A to B in unit time than with randomly chosen aggression. The method is economically rational; it makes decisions based on the expectation of a favourable outcome. This method will be called the 'investment' method from now on.

The investment method can be implemented very simply by adding a minimal memory to the robot: a counter. The counter is incremented each control loop cycle. On reaching a goal the counter is reset to zero. The value in the counter reflects the amount of time the robot has spent on reaching its current goal. The aggression level is set proportional to this value.

## 3.3 Hypotheses

The experiments presented in this chapter were designed to test the following hypotheses:

1. The results obtained in [57] can be replicated.

2. Investment-based aggression will allow a team of robots achieve more work than when a random aggression function is used.

3. The implementation of the investment based aggression mechanism is decentralized, independent of a navigation strategy, makes use only of existing sensors (sonar and laser), works in heterogeneous robots systems and is trivial to compute.

## 3.4 Experimental Design

This section describes the experiments carried out to evaluate the advantage of investment based aggression over other non rational aggression methods.

### 3.4.1 Task

Robots have the task of transporting resources back and forth between two goal locations (shaded areas in the rooms, Figure 3.2). When a trial is executed all robots go to the same starting position, then they proceed to the first goal area (bottom room). Once the goal is

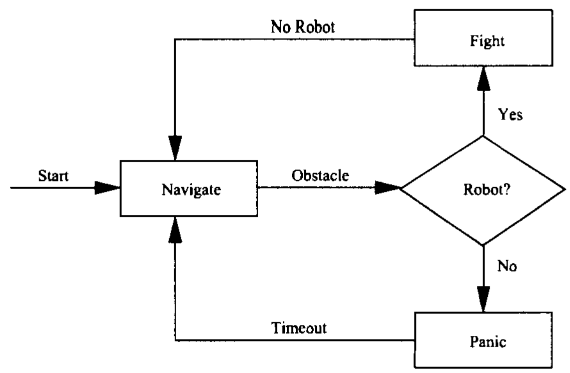


Figure 3.4: Control architecture [57]. Printed by permission of Richard Vaughan and the MIT Press (Book Title: "From Animals to Animats 6") ©.

reached they proceed to the second goal area (top room). Each time a robot reaches a goal area it drops one unit of resource and increases the amount of work done by one unit as well.

The world in Figure 3.2 closely resembles the world used in [57] therefore the results presented in this document can be compared with those described in that paper.

### 3.4.2 Control Architecture

Each robot runs the same control program, shown schematically in Figure 3.4. Each mode is described below.

#### Navigate

*Navigate* is the default behaviour for a robot and the only way that robots achieve work. A robot will navigate until it finds an obstacle in its path. The navigate mode used here is an adaptation of the one presented in [57], however, instead of using a crumb trail, the robot uses a two dimensional map of the environment which provides a heading direction for any possible location of the robot, as shown in Figure 3.5. The direction vectors change depending which goal area the robot is seeking.

Robots perform left wall-following and obstacle-avoidance using a modified version of the robust 'sliding box' algorithm [57]. The 'sliding box' algorithm works in the following way: A virtual box slightly larger than the robot is moved from the robot's left to the robot's right until no obstacles from the laser scan are detected within it as shown in Figure

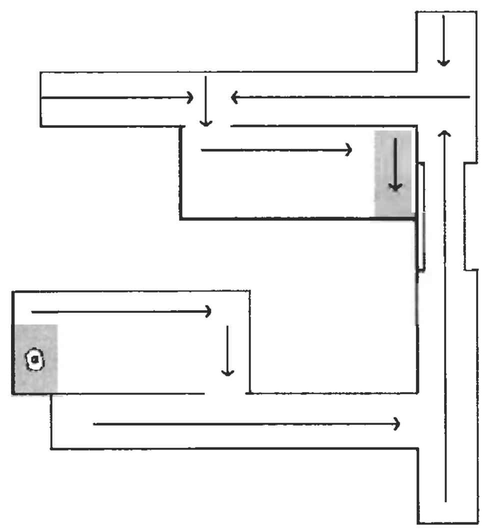

Figure 3.5: A trail map with general directions to go from the bottom room to the top room.

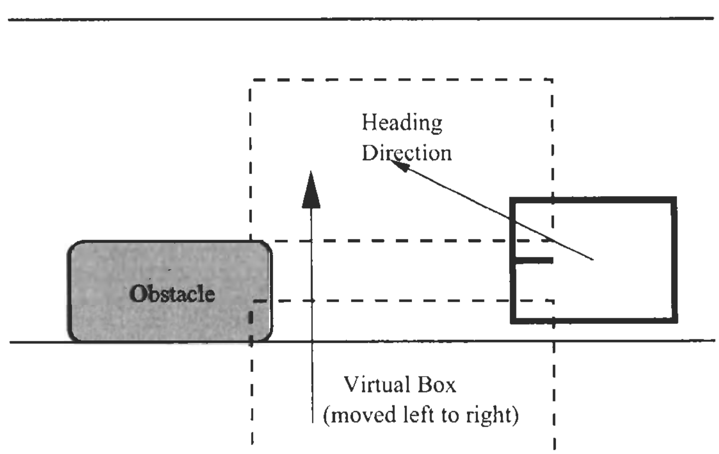


Figure 3.6: Robot negotiating an obstacle inside a corridor and using the sliding box algorithm to find a desired heading.

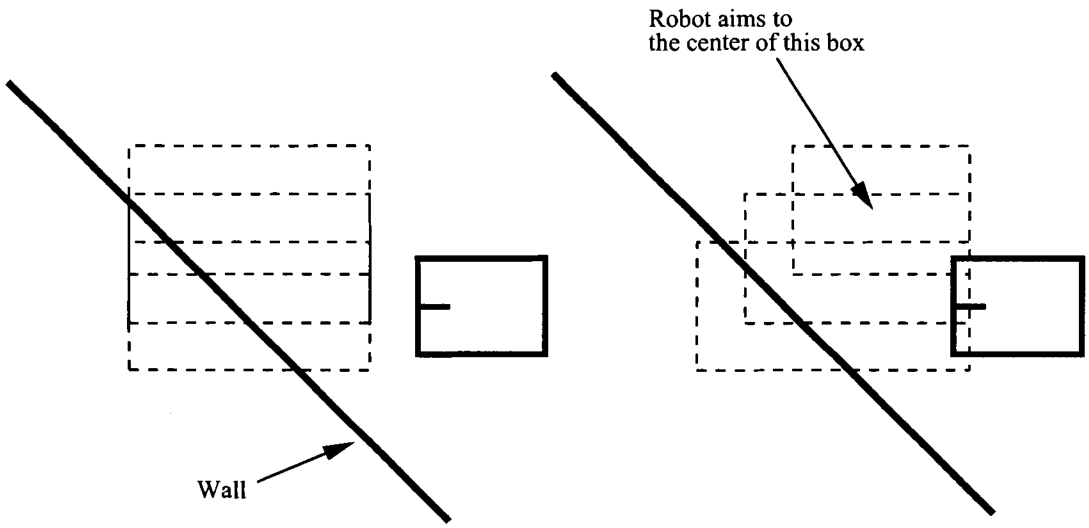


Figure 3.7: In the left figure the robot does not find an empty virtual box. In the right figure the robot finds an empty box and aims to the center of it.

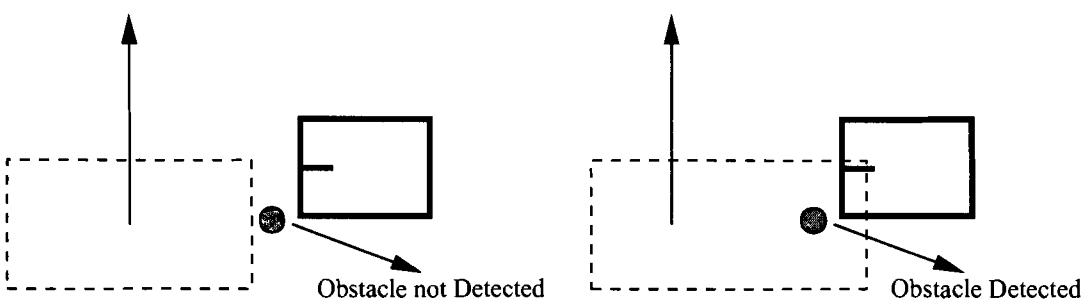


Figure 3.8: In the left figure the robot is blind to a near obstacle in its path. In the right figure, the robot always sees the obstacle.

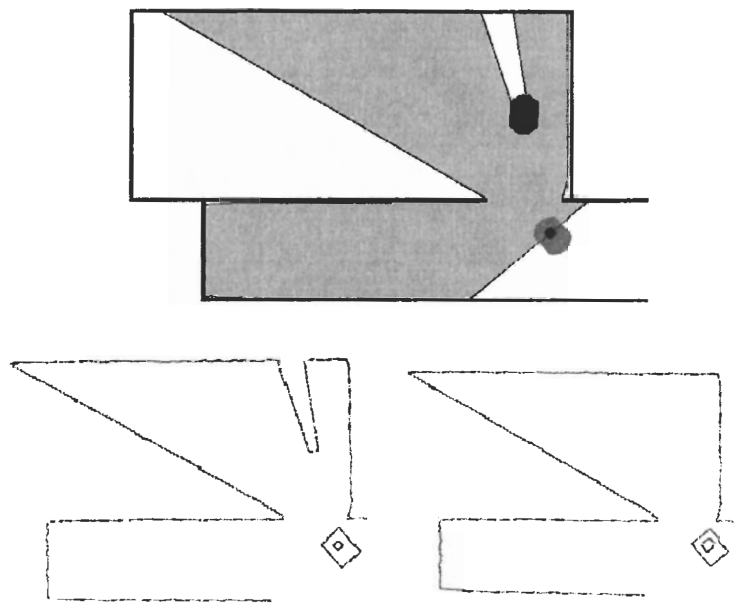

Figure 3.9: The top figure shows two robots in simulation. The bottom figures show the laser scan of the robot trying to get into the room before and after the other robot is removed from it.

3.6. Once an empty box is found, the robot aims to the center of the box. This has the overall effect of producing smooth obstacle avoidance.

An extension of this technique is presented in this document. It is called 'variable length sliding box' and provides the following additional advantages over the 'sliding box' algorithm: (i) Instead of a fixed length virtual box, a 'variable length' box is used. This slight modification allows the robots to handle situations when they are not parallel to the wall (See Figure 3.7). (ii) The virtual box always intersects the robot and in that way very close obstacles are always detected (Figure 3.8).

Given that the world is populated by several robots, other robots need to not be treated as obstacles. To achieve this, robots are erased from the laser scan and the gaps filled by interpolation from the rest of the scan (Figure 3.9). The filtered laser scan is only used by the navigation mode, all other modes (i.e. emergency stop, panic) use the original laser scan.

In the simulated world most of the corridors are big enough for two pioneer robots to pass comfortably. When turning corners, however, robots that swing too wide, or not wide enough, can interfere with each other. To solve this problem, robots which are turning left make very sharp left turns and robots turning right use the far right corner of the wall, as detected in the laser scan, to make the widest turn possible. If a robot turns too sharply (or not sharply enough) to make a turn correctly, this may trigger an emergency stop which would normally lead to the robot entering the *panic* mode. This is inefficient so a small corrective turn was added to the navigate behaviour for this situation. Robots in the navigate behaviour will first attempt to make a very small rotation towards the direction in which the most space is perceived from the laser scan then resume navigation. If this fails, navigate will exit abnormally and the *panic* mode will become active.

### Panic

When a robot finds itself in a situation in which it has an obstacle it cannot negotiate or it has been fighting with a robot for a period of time without any resolution in the fight, the robot switches its behaviour to panic. Panic will, in general, take a robot and make it move to a random location for certain amount of time. Once the way is clear, the robot goes back to the navigation mode.

The panic behaviour can be summarized in the following steps:

1. Sleep for a random period of time and check if way is free, if not go to next step.
2. Check if surrounded by obstacles, if not, go to next step.
3. Turn for a random amount of time, if not able to rotate, move slightly to the place with more empty space and try to rotate again.
4. Navigate for a random period of time, if successful, get out of panic and switch to navigation. If not successful go back to step 2.

The panic behaviour was made very robust, by use of randomness. This guarantees that deadlock situations are never encountered. Nobody wants two robots on a distant planet in a deadlock loop. However, this design creates some disadvantages for the robot, because the behaviour is not goal oriented, all the energy spent during panic is normally wasted energy; a robot while panicking does not know if it is getting closer or farther from the goal. Also as shown in the trials run in the experiment a robot panicking takes a considerable amount of time before switching back to navigation. This is a characteristic that can seriously degrade the performance of a team of robots. To reduce the amount of panic in the system and to obtain better statistics in shorter periods of time, the navigation and fighting behaviours were design to be adaptable to many different situations so that the panic behaviour was used as little as possible.

#### Fight

The *fight* behaviour is triggered when a robot detects another one in front at a certain distance, usually because the other robot is blocking the way. The fight model used in the experiments in this chapter is the same as in [57], and is called *the inverse chicken game.* When entering the *fight* mode, the robot calculates its *fear threshold*, which is the minimum distance at which it will tolerate another robot. Robots start fighting by moving backwards; they continuously back away as long as the other robot is within its fear threshold. If the rival is outside its fear threshold, the robot switches back to the *navigate* behaviour. If the robot is too close to an obstacle while it is moving backwards, the emergency stop mechanism is invoked. In this case the robot stops and switches to *panic.* During a fight the robot with the smaller fear threshold will be the first to start driving forward again - this robot is the *winner.* Consequently, it will push its rival (the *chicken/loser*) backwards until there is enough room to pass. Once the winner moves outside the loser's fear threshold, the loser starts moving forward again and the fight is over.

### 3.4.3 Aggression Function

The fear threshold is determined by a robot's *aggression* $\alpha$, where $0 < \alpha < \alpha_{max}$ and $\alpha \in \mathbf{R}$, selected at the start of the *fight* procedure. *Fight* is designed so that the more aggressive robot is likely to be the winner.

The fear threshold is the minimum distance one robot can tolerate another robot and is inversely proportional to a robot's aggression, plus some offset distance.

$$\phi = K_1 + \frac{K_2}{\alpha} \tag{3.1}$$

Constants $K_1$ and $K_2$ were chosen to give a fear threshold between 450mm and 2450mm in the experiments. A tiebreaker mechanism is employed in order to reliably resolve the fight between two robots with very similar fear thresholds. It adds a small random distance to a robot's fear threshold, ensuring that two robots with the same aggression have different fear threshold values. It breaks the symmetry between two robots by preventing them from switching from *fight* to *navigate* at the same time.

Next is an explanation of each of the aggression functions used in the experiments.

#### None

In this case there is no fighting behaviour. When a robot detects another one blocking its path, it panics. This non-aggression technique is used as a control test.

#### Random

Aggression $\alpha$ is chosen at random in the range $0 < \alpha < \alpha_{max}$, $\alpha \in \mathbf{R}$.

#### Investment

Aggression value is proportional to the time a robot has spent approaching the goal in each trip. A robot's aggression increases with the time it has spent on the *navigate* behaviour within the current trip. Specifically, the aggression $\alpha$ is calculated using the formula:

$$\alpha = min(K_3 \times \frac{T}{T_{normal}}, \alpha_{max}) \tag{3.2}$$

where $T$ is the time spent approaching the current goal, $T_{normal}$ is a normalization constant reflecting the expected time to reach the goal, and $K_3$ scales the aggression to the desired range. $\alpha_{max}$ sets the upper bound of $\alpha$, so that $0 < \alpha < \alpha_{max}$.

### 3.4.4 Procedure

There are $n$ robots living in a simulated world $W$. The world has rooms and corridors where the robots can move. The doors and some sections of the corridor are narrow and only allow one robot to pass by. All other sections of the corridors are wide enough to allow two robots to pass through when going in opposite directions.

For each aggression function of 'none', 'random' and 'investment', where 'none' means that fighting is disabled, a total of $n_{trials}$ trials that last for a number of seconds $trialLength$ are run. Every time a new trial is started, the location of the robots is reset to the same initial position. All trials have the same starting conditions, but the experimental system has non-deterministic features, so that each trial is different.

These parameters allow for control of the degree of spatial interference between the robots in the world, in a way that the performance of each of the different aggression functions can be meassured. In the following sections, two experiments in two different worlds $W1$ and $W2$ are presented. Setting the number of robots $n = 6$ produced a good degree of robot interference without saturating the world, which means that the robots can still achieve their goals.

### 3.4.5 Performance Metric

As soon as a trial starts, all robots begin to log information regarding the type of aggression function used, the trial number, the robot number, the number of trips completed, and the total time spent in *navigate*, *fight* and *panic* behaviours.

To measure the success of a trial the sum of all the trips $T_{team}$ performed by the team of robots is calculated (Equation 3.3).

$$T_{team} = \sum_{i=1}^{n} T_i \tag{3.3}$$

where $n$ is the number of robots in the team and $T_i$ is the number of trips performed by robot $i$. This value is easy to obtain and represents an objective measurement of the performance of the system as a whole. In the resource transportation task, each trip completed by a robot is equivalent to one unit of resource transported. It is not the goal to improve the number of trips that a single robot does but rather the number of trips that the entire team of robots completes.

Using the values of $T_{team}$ in a different number of trials, the performance of each aggression function can be compared with one another.

Another useful comparison for $T_{team}$ is against a hypothetical maximum number of trips, $T_{max}$, that is the total number of trips that a team of robots could execute given that none of them interfere with each other. To compute this value the following is done:

1. Find $T_{one}$, the average of the total number of trips that a single robot does during a $trialLength$ for $n_{trials}$ trials. In this experiment, the robot uses the same controller

of the multirobot experiments. Also, given that it is the only robot in the world, it never fights or panics.

2. Multiply $T_{one}$ times the number of robots ( $n$ ) and obtain the upper-bound on the number of trips a team of robots could do $T_{max}$ (See Equation 3.4). In reality, this would only be possible if none of the robots would interfere with each other during the trial. Note that ($T_{max}$) is also the upper-bound on the number of resources which could be transported in the system.

$$T_{max} = T_{one} \times n \tag{3.4}$$

$$0 \leq T_{team} \leq T_{max} \tag{3.5}$$

For a given trial duration, there is a relation between the number of trips performed in a trial and the total time spent in navigation, fighting and panic. The greater the number of trips, the greater the navigation time and the lower the fighting and panic times. Because the goal is to reduce interference between robots as much as possible, most of the time spent by a robot during a trial should be navigation time. Fight and panic times are wasted time; they do not directly contribute towards achieving a goal. Therefore, a good interference reduction method should increase navigation time and reduce one or both of the panic and fight times.

### 3.4.6 Statistical Tests

A value of 1.5 times the interquartile range (IQR) has been used to set up fences to remove trials that are outliers. In most cases, the outlier trials are caused by the robot controller not being robust enough, or because of the amount of panicking in the trial.

On the filtered data, two-tailed t-tests have been run to show the difference between the aggression functions. A significance level, $\alpha < 0.05$, has been used in order to mark two distributions to be statistically different. This method is a standard technique used by biologists when the number of samples is small and the distributions are known to be Gaussians.

A table of the distribution of values of t is found in Appendix C.

## 3.5 Experiment 1

In this experiment the following parameters are used:

- $n = 6$, the total number of robots.
- $W$ like Figure 3.2, the world where the robots live.
- $trialLength = 1800$, the length of each of the trials in seconds.
- $n_{trials} = 24$, the total number of trials executed for each aggression function.
- $Trips_{one} = 21$, the average number of trips performed by a single robot in the world $W$.
- $Trips_{max} = n \times Trips_{one} = 6 \times 21 = 126$, the upper bound on the number of trips for a team of robots.
- $\alpha_{max} = 10$, the maximum aggression level of a robot.

The aggression function is set such that $0 < \alpha < \alpha_{max}$, with the aggression reaching 80% of maximum in the normal time taken to reach a goal. This time was determined empirically by measuring the mean time taken by a single robot to drive between goals, in the absence of interference, and was found to be 84 seconds. Substituting these parameters, the aggression function of $T$ becomes:

$$\alpha = min(K_3 \times \frac{T}{T_{normal}}, \alpha_{max}) = min(8 \times \frac{T}{84}, 10) \tag{3.6}$$

$$\alpha = F_{expl}(T) \tag{3.7}$$

### 3.5.1 Results

From Table 3.1 and the histograms presented in Figure 3.10 it can be seen that a team of robots with an aggressive display behaviour, investment or random, performs better than a team of robots with no fighting behaviour. The performance of random and investment though similar is actually a small amount better for random. The t-tests showed that the random and investment trials belong to two different distributions. Also in Figure 3.11 it can be seen that the distribution of fighting, panicking and navigation time is quite similar

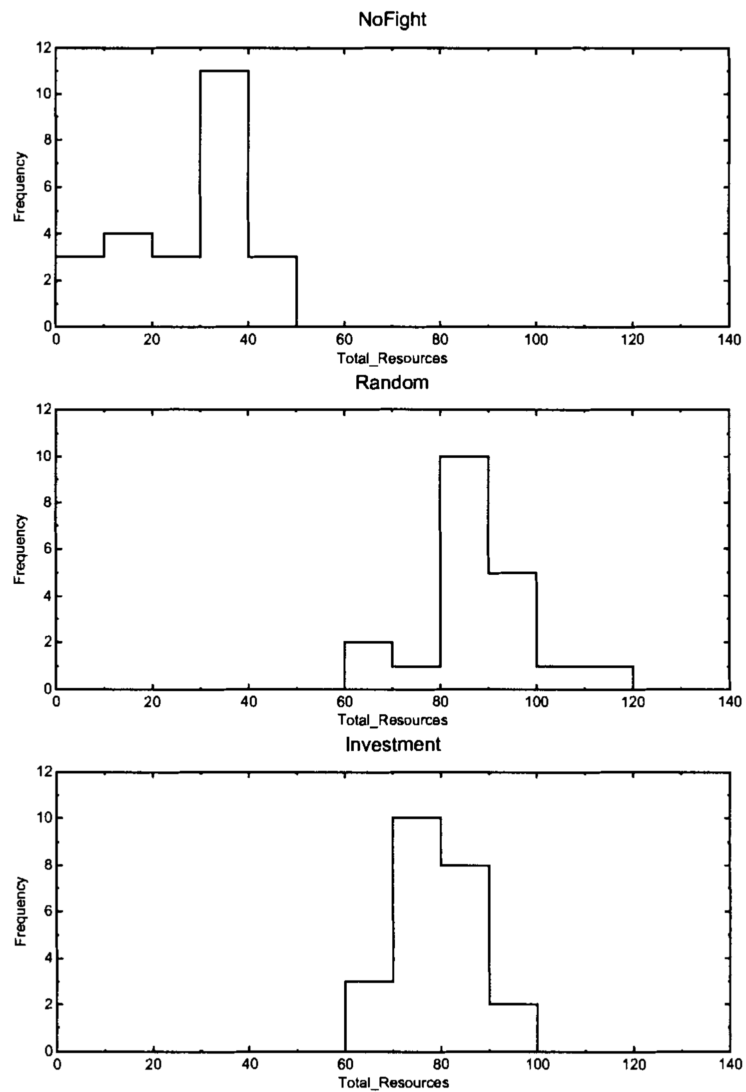


Figure 3.10: Expl results: Histograms showing distribution of performance scores for three different controllers: no fight (top), random aggression (middle), investment aggression (bottom). The graph shows the number of trials (Frequency) in which an amount of resources (Total Resources) was transported. For example, in the Random graph it is shown that in 10 of the 24 trials run in the experiment, the team of robots was able to transport between 80 and 89 units of resource.

| TrialType | Mean Resources | $\sigma$ | $N$ | Outliers | $t - test_{random}$ |
|---|---|---|---|---|---|
| NoFight | 28.0 | 11.4 | 24 | 0 | -17.25 |
| Random | 87.3 | 11.2 | 20 | 4 | 0.0 |
| Investment | 79.5 | 8.6 | 23 | 1 | -2.53 |

Table 3.1: Exp1 results: Summary of performance scores from three different controllers.

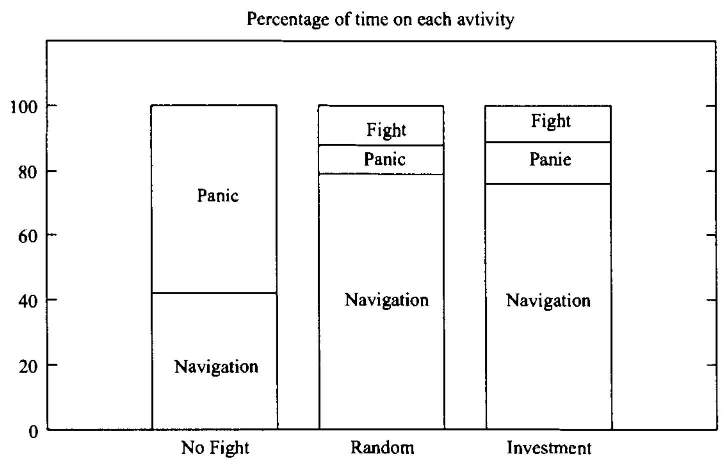


Figure 3.11: Exp1 results: Proportion of time spent in each activity.

for the random and investment methods, while the navigation time in the 'no fight' trials is small as expected.

Though it was expected that the investment method was going to show a performance increase in the team of robots when compared to the random method, there are several observations that give clues about why this was not the case in the experiment:

(i) When robots interfere in the doorways, if they have enough space to safely manoeuvre when getting in or out of a room, there is no real advantage for one robot winning the fight over the other (Figure 3.12). In a different situation, one of the robots has other robots behind, going in the same direction and with the same goal. This type of spatial configuration is called a *worm of robots.* Because the robots do not have any communication mechanism other than their perception (laser and sonar) a robot does not have a way of detecting that they are members of a worm. As a result robots cannot use that property to modify their aggression (Figure 3.13). The long term effect of this lack of communication is that sometimes a worm will win a fight, and sometimes it will lose it. In the long term both effects on the system performance cancel out and are equivalent to a random outcome.

(ii) The narrow corridor in the world (Figure 3.2) is about halfway between the goals and as a result robots with similar aggressions compete for the right of way. The robot controller is implemented in a way that guarantees to resolve fights even between robots in this situation, however the fights take longer to resolve than when two robots with very different aggressions compete. A random approach usually generates different aggression values for the robots and consequently has an advantage over the investment method when robots meet in the narrow corridor. I believe this to be the main reason why the performance of the investment method is slightly worse compared to a random outcome.

(iii) The narrow corridor is very short in length, and as a result, in order to obtain a significant statistical difference the trials would require to be run for a very long time (this would require a more robust controller) or the number of robots would need to be increased so as to have more fights (but experiments run with more than 6 robots showed that the world would become saturated).

All the points presented above have one thing in common, they show how important the role of the environment is when trying to evaluate one aggression function over another. This insight into the role of the environment may help to explain the results in [57], where a 'personal space' strategy was also shown to be equivalent to random.

At this moment in time, no previous work, [57] and this thesis, has found if there even

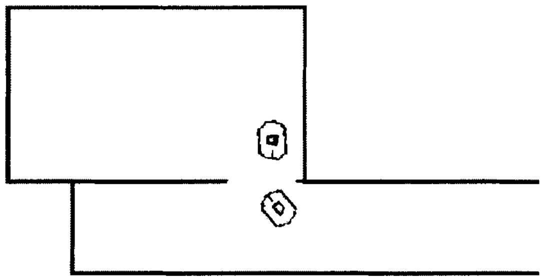

Figure 3.12: Robots fighting for the right to pass across a door. There is no advantage in one robot winning the fight over the other.

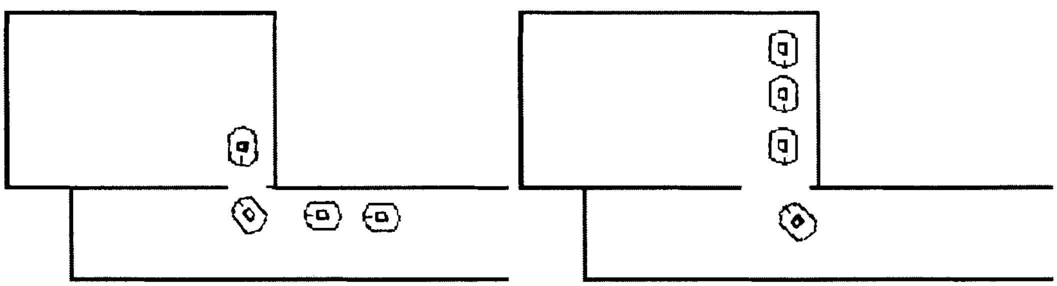

Figure 3.13: Robots fighting (one against a worm of other robots) for the right to pass across a door. If robots were able to communicate it would be possible to make the worm win and improve the performance of the system. However, without communication any door fight can be either of the situations presented in the figures, and the controller would not be able to know how to properly modify the aggression of a robot.

exists an environment in which a rational method can outperform a random scheme. I believe that it is certainly possible to create a simple environment with some special characteristics that would cause one aggression scheme to succeed or fail. However, we do not want 'any' environment, we would like environments that are still *realistic*, that is: (i) environments that can be found in the real world, and at the same time, (ii) environments in which the task given to the team of robots makes 'sense'. For example, putting 20 robots into a small room where a task has to be executed, does not make sense, as the environment is clearly over-crowded.

Next, a second experiment is presented that takes into account the arguments presented above.

## 3.6 Experiment 2

It is one of the goals presented in this thesis to find whether a rational method of setting the aggression of a robot can outperform a random approach. At the end of experiment 1 there are several hypotheses that explain why a rational approach based on investment did not perform better than the random scheme. It was hypothesized that the environment plays a very important role in the performance of any given aggression scheme. Following this idea a second environment (Figure 3.14) was designed to show the features, if any, of the investment method. This new simulated world has longer narrow corridors that increase the penalty for the 'wrong' robot losing a fight. Also the narrow parts of the corridors are not located at the halfway point between the goals, this causes the fights to be solved more quickly as the difference in aggression between the robots is increased. The new environment still shares many similarities with the previous one.

The new world was designed to be approximately twice the size of the old world. Also the average time for a robot to go from source to goal is about twice the time. For this reason, the length of the trials was doubled.

The following is a list of the parameters used in this experiment:

- $n = 6$, the total number of robots.
- $W$ like Figure 3.14, the world where the robots live.
- $trialLength = 3600$, the length of each of the trials in seconds.

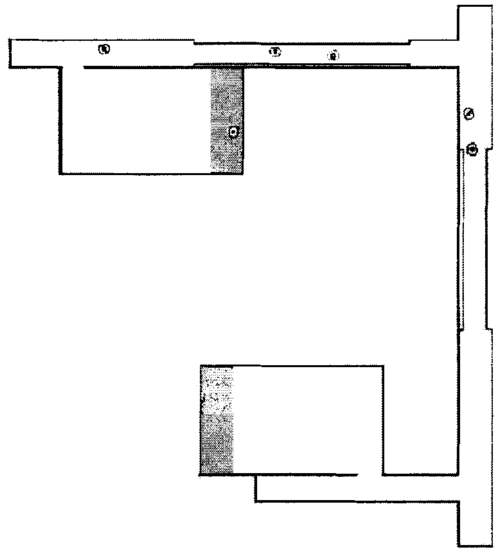

Figure 3.14: The world used in the second set of experiments.

- $n_{trials} = 24$, the total number of trials executed for each aggression function.
- $Trips_{one} = 22$, the average number of trips performed by a single robot in the world $W$.
- $Trips_{max} = n \times Trips_{one} = 6 \times 22 = 132$, the upper bound on the number of trips for a team of robots.
- $\alpha_{max} = 10$, the maximum aggression level of a robot.

The Investment aggression function was set similarly to experiment 1, except that $T_{normal}$ was larger due to the larger environment.

$$\alpha = min(K_3 \times \frac{T}{T_{normal}}, \alpha_{max}) = min(8 \times \frac{T}{152}, 10) \tag{3.8}$$

$$\alpha = F_{exp2}(T) \tag{3.9}$$

The configuration of this experiment allows some comparison between both experiments as the total number of trips possible for a robot is also approximately the same.

### 3.6.1 Results

The t-tests run on the experiment show that the investment based trials belong to a different distribution than the random based trials. The investment aggression outperforms the random and no fight schemes. This means that the team of robots in the investment aggression trials achieve more work on average. Table 3.2 and the histogram presented in Figure 3.15 give a summary of the performance of the three aggression schemes. The proportion of time spent in each behaviour is shown in Figure 3.16. As stated before, the investment scheme is the one which enables the robots to spend the longest time in navigation.

## 3.7 Conclusion

Some of the results obtained in [57] have been validated. It was confirmed that in the slightly different environment used in this chapter (Figure 3.2) none of the different aggression methods performed better than random aggression. In the first experiment it was found that an aggression function based on the concept of investment in a task did not perform better than a random aggression approach. Hypotheses about features of the environment and length of the trials were presented, and a second experiment was designed as a result.

In the second experiment a new environment was built (Figure 3.14). This environment is similar to the old one but has longer narrow corridors where the robots interfere. The idea was to increase the penalty for the wrong robot winning a fight. The new world was still a realistic office space. Tests run in this environment showed that the rational investment aggression scheme performed statistically better than random aggression. The robots using investment aggression were able to complete more work (second of the hypotheses). It also suggests that other rational techniques previously used in [57], like 'personal space', could have performed better than random if used in a different environment.

The experiments in this chapter also showed that the investment method could be easily implemented by the use of a simple counter that was increased while going from a source position to a goal position. The robot controller and the investment method required the use of only existing sensors (sonar and laser), any extra computations were trivial to do. Though no experiments were done in any other robot platform there is no reason to believe that the investment approach could not be used in other robots with similar sensing capabilities.

In summary, a novel interference reduction technique using aggression display has been

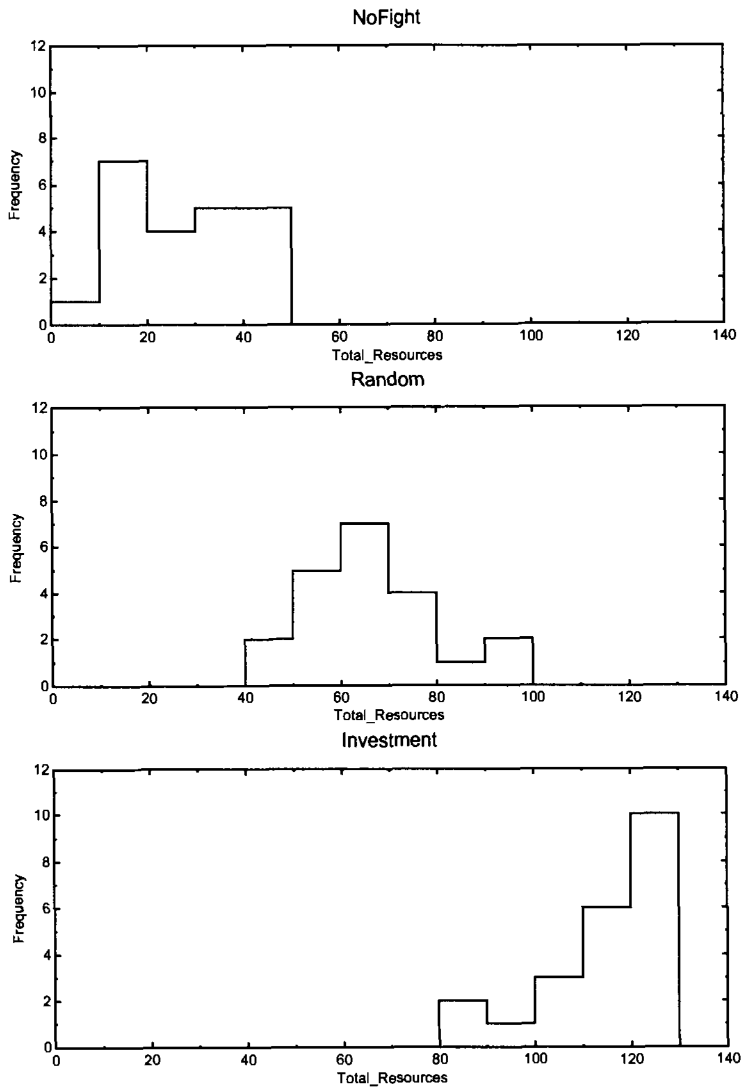


Figure 3.15: Exp2 results: Histograms showing distribution of performance scores for three different controllers: no fight (top), random aggression (middle), investment aggression (bottom). The graph shows the number of trials (Frequency) in which an amount of resources (Total Resources) was transported. For example, in the Random graph it is shown that in 7 of the 21 trials run in the experiment, the team of robots was able to transport between 60 and 69 units of resource.

| TrialType | Mean Resources | $\sigma$ | $N$ | Outliers | $t - test_{random}$ |
|---|---|---|---|---|---|
| NoFight | 25.1 | 13.7 | 23 | 1 | -9.88 |
| Random | 65.5 | 13.4 | 21 | 3 | 0.0 |
| Investment | 112.9 | 13.4 | 22 | 2 | 11.57 |

Table 3.2: Exp2 results: Summary of performance scores from three different controllers.

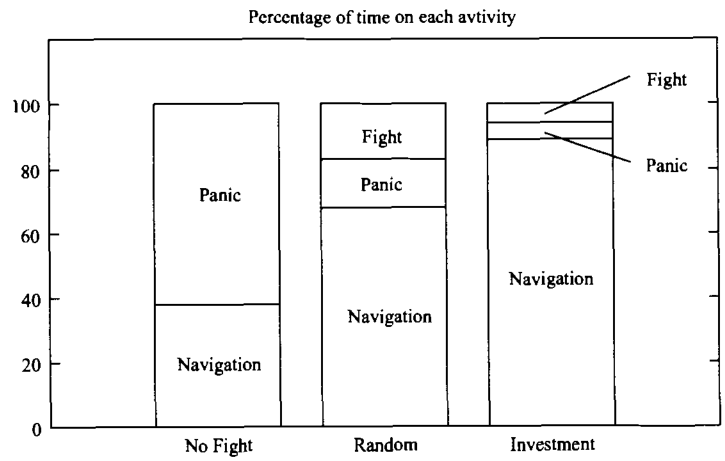


Figure 3.16: Exp2 results: Proportion of time spent in each activity.

presented. The aggression function is based on the idea of investment in a task. A simple implementation was described, and the method was shown to increase the performance of a simulated robot team at a classical transportation task. Unfortunately the method requires some special characteristics about the world that may result in the technique not being sufficiently robust for real world situations. To verify or falsify the previous statement, many more experiments in different worlds would need to be performed. Still, the investment method should be applicable as an interference reduction technique in mobile robot teams.

# Chapter 4

# Local RAGE

The work presented in this chapter is based on work done in collaboration with Richard Vaughan [61], accepted to the IEEE/RSJ International Conference on Intelligent Robots and Systems IROS2005, Edmonton, Alberta, August 2-6, 2005.

## 4.1 Introduction

In section 3.1, an approach to selecting an aggression level based on a robot's investment in a task was introduced. The concept of 'investment' of work done towards achieving a goal is fundamental in models of autonomy in animals [41]. Simulation experiments with teams of six robots in an office-type environment showed that, under certain conditions, the method was able to significantly improve system performance compared to a random competition and a non-competitive control experiment.

In this chapter, a new approach to selecting the aggression level based on the concept of 'local investment' is presented. In this method, the aggression level is proportional to the effort that a robot has put recently into crossing areas where there is a high probability of interference. This method was designed to overcome limitations of the previous investment approach (hereafter 'global investment' approach).

For the first time a 'real world' implementation of a stylized fight is presented in a multi-robot resource transportation task.

Experiments presented in this chapter confirm that the local investment method is the best performer in different simulated and real environments. They also show that the method can be easily implemented in real world situations.

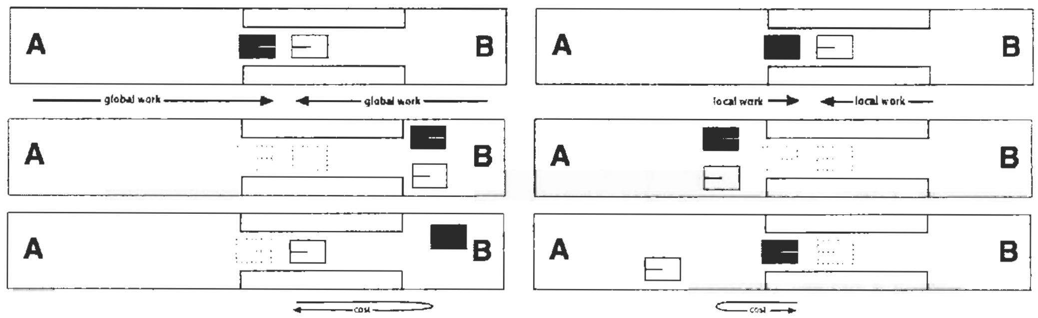


Figure 4.1: Motivation for the *local investment* strategy. The robot that has invested the most work passing the narrow part of the corridor should win.

## 4.2 Local Investment Aggression

In a similar way to the global investment approach, the local investment method is economically rational. To understand the concept of local investment, we need to first define an *area of interference* as an area of the environment where spatial interference between robots is likely, for example a narrow corridor or a door. The local investment method accounts for the effort put into passing an area of interference while the global investment method accounts for the effort put into completing a task.

Now consider the system of robots presented in Figure 4.1. The task is the same as in the global approach explained previously, the Black robot transports resources from A to B and the White robot from B to A. About two thirds of the path are wide enough for two robots to pass without interfering with each other. There is however a narrow part of the corridor (about one third of the total length) where only one robot can pass.

Figure 4.1 presents the two possible outcomes of a conflict happening in the left part of the narrow corridor. Either the White robot retreats and allows the Black robot to pass or the Black robot retreats and the White robot moves forward. Once both robots are in the wide part of the corridor they can resume navigation without any interference.

It is seen in the figure that the sunk costs are higher for the White robot losing the fight, consequently it would be better to solve the fight in favour of the White robot. In the situation presented in this figure, a global investment approach would choose in favour of the Black robot, because it has put more work into getting to its current location compared to the White robot (left column of the figure). On the other hand the local investment

Figure 4.2: Sequence of a Fight in which the cost is minimized, (sequence order is left to right, top to bottom).

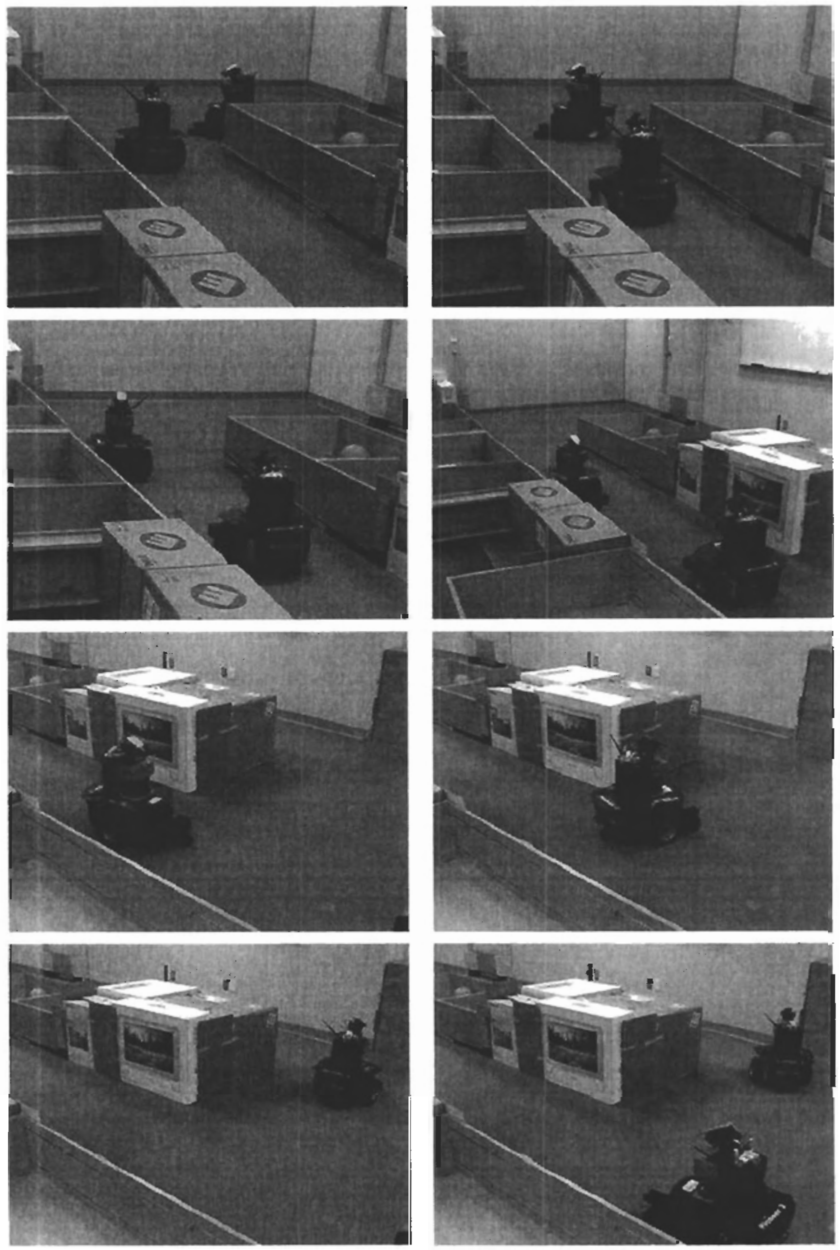

Figure 4.3: Sequence of a Fight in which the cost is very high. one of the robots was close to finish passing through an area of interference and yet had to retreat (sequence order is left to right, top to bottom).

approach would allow the White robot to win, as it has spent more time inside the narrow part of the corridor (right column of the figure). Therefore, in this situation the local investment approach handles the conflict in an optimal way.

The cost of a fight is something that can only be measured after the fight is over. As we cannot measure the future cost directly, the best we can do is to predict the cost of fighting. The better the prediction, the higher the likelihood of making the right decision in a conflict situation. The experiments described in Chapter 3 showed that global investment was not very good at predicting the cost of fights, we believe that the better performance of local investment is due to it being a better predictor of the future cost of losing a fight.

Figures 4.2 and 4.3 show two examples of real-world fights implemented on Pioneer 3-DX robots. The top row shows the sequence (left-to-right) of a fight with a favourable outcome, in which the interference is minimized because the robots back up as little as possible. The bottom row shows an unfavourable fight, in which one of the robots has to back up a long way before resuming its normal path.

## 4.3 Hypotheses

The experiments presented in this chapter were designed to test the following hypotheses:

1. Local investment aggression outperforms the previously studied random and global investment aggression when looking at the amount of work done by a team of robots in a realistic office environment.

    - The team of robots using the aggression mechanism can be autonomous for the length of the trials.
    - The Local Investment aggression mechanism is decentralized, independent of a navigation strategy, makes use only of existing sensors (sonar and laser), works in heterogeneous robots systems and is trivial to compute.

2. Local investment does not require special features in the environment as global investment does.

3. A stylized competition mechanism can be successfully implemented in the real world to reduce interference in a resource transportation task.

4. The results found in simulation are carried to the real world.

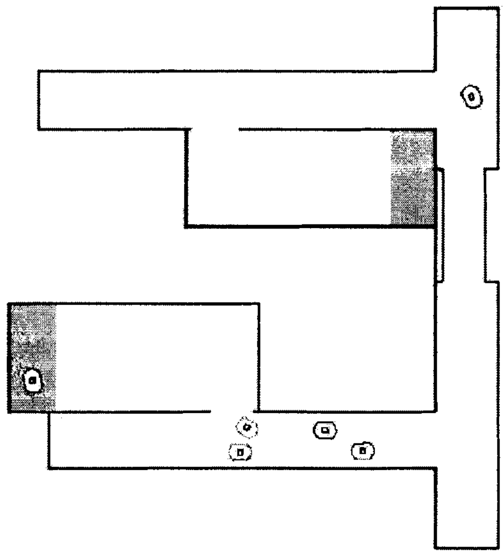

Figure 4.4: The Stage world containing six robots used in Experiment 1 and in [14]. It closely resembles the environment used in [57].

## 4.4 Experiment 1: Local Investment *vs.* Global Investment and Random Aggression in a Simulated Environment

This section describes the experiments carried out to compare the performance of the three methods when used in a team of robots. The experimental design, robot controller and environment used and described in section 3.4 are also used in the experiments presented in this section. The world is shown in Figure 4.4.

### 4.4.1 Aggression Function

The Random and Global Investment aggression functions are described in section 3.4. Next is the description of the new proposed method.

#### Local Investment

The local investment method is based on three different aggression functions that are used depending on the robot perceptions and the environment. Formally, the following equations are used:

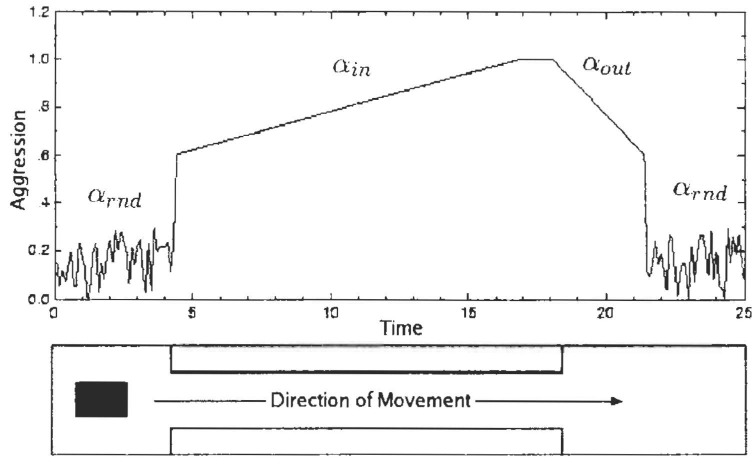


Figure 4.5: Local Investment Aggression

$$\alpha_{in}(t) = (m_{in} \times t_{in}) + C_{in} \tag{4.1}$$

$$\alpha_{out}(t) = C_{out} - (m_{out} \times t_{out}) \tag{4.2}$$

$$\alpha_{rnd}(t) = m_{rand} \times random_{uniform}[0 \dots 1] \tag{4.3}$$

Aggression $\alpha$ is equal to $\alpha_{in}(t)$ if inside an area of interference, equal to $\alpha_{out}(t)$ if just comming out of an area of interference and equal to $\alpha_{rnd}(t)$ if $\alpha_{out}(t)$ is below $C_{in}$. In the experiments $C_{in} = 0.6$ , $C_{out} = 1.0$ , $m_{in} = 0.008$ , $m_{out} = 0.03$ and $m_{rnd} = 0.03$. The final value of aggression was also thresholded between 0 and 1.

Figure 4.5 shows a graph of a robot's aggression while moving left to right in the environment. The aggression is initially random because the robot is not in an interference area. Then, as soon as the robot starts to move inside the narrow corridor, its aggression starts to increase linearly until reaching the maximum ($\alpha = 1.0$) around the end of the narrow part of the corridor. Once the robot is out of the corridor, its aggression decreases quickly and when below some threshold it reverts to random aggression. Note that the random aggression is always smaller than the minimum aggression inside the corridor. Therefore a robot inside the corridor fighting with a robot outside the corridor is more likely to win the fight.

In summary, local investment aggression is a small random number or a fast decreasing linear quantity when outside an area of interference, and proportional to the time spent inside an area of interference if inside an area of interference.

Next is a list of the parameters used in the experiment:

- $n = 6$, the total number of robots.
- $W$ like Figure 4.4, the world where the robots live.
- $trialLength = 1800$, the length of each of the trials in seconds.
- $n_{trials} = 24$, the total number of trials executed for each aggression function.
- $Trips_{one} = 21$, the average number of trips performed by a single robot in the world $W$.
- $Trips_{max} = n \times Trips_{one} = 6 \times 21 = 126$, the upper bound on the number of trips for a team of robots.

### 4.4.2 Results

The results presented in Table 4.1 show that the local investment method is significantly different and better than the global investment and random based aggression schemes. This can also be seen in the shape of the distributions presented in Figure 4.6. While random and global investment have similar distributions, the local investment distribution is shifted to the right which means that the team of robots was normally able to complete more trips during the length of the trials. In Figure 4.7 it can be seen that the local investment strategy reduces the amount of time spent fighting and panicking while increasing the amount spent in navigation.

In the work presented in [57] and [14] the authors spent a long time trying to come up with an aggression function that outperformed the random approach without any success. In fact, many aggression functions not documented in the papers were also tested (inverse personal space, inverse global investment, etc). An important aspect of the world used in these experiments is that it corresponds to a real office building somewhere in California, and therefore its constraints are real. The success of the local investment method is therefore worthy of further study and suggests that the method may be robust enough for many other real world situations.

Note that the standard deviation of the local investment distribution seems large compared to the random and global trials. This is because some trials are marked as outliers for

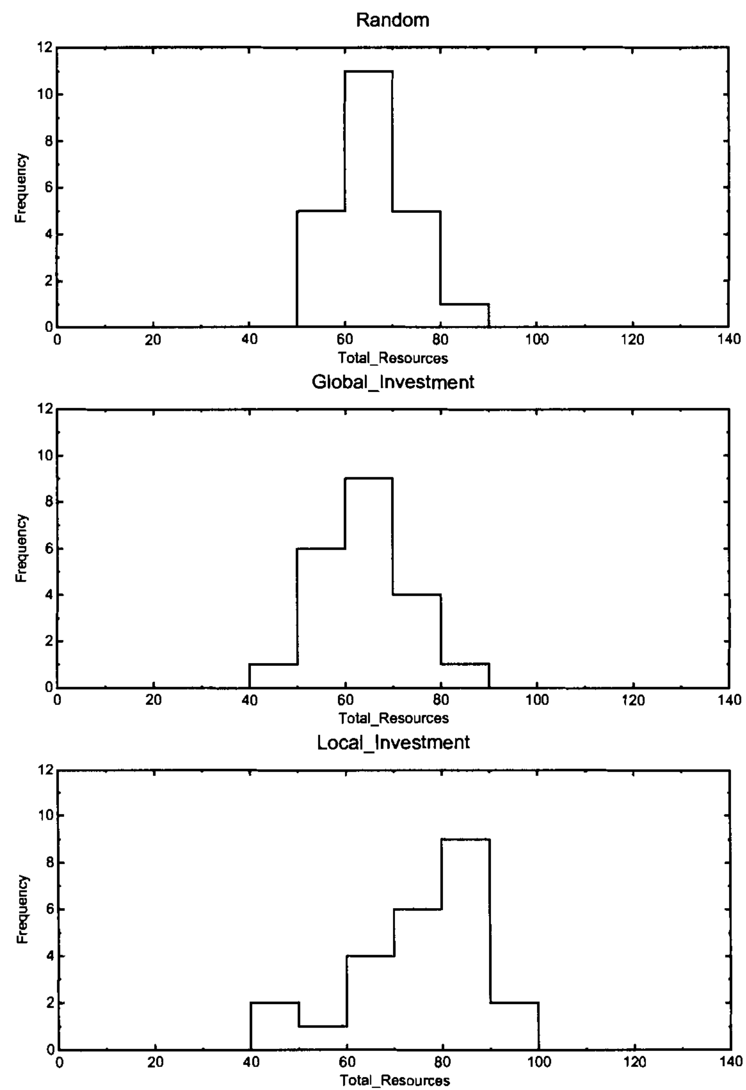


Figure 4.6: Exp. 1 results: Histograms showing distribution of performance scores for three different controllers: random aggression, global investment and local investment aggression.

| TrialType | Mean Resources | $\sigma$ | $N$ | Outliers | $t - test_{random}$ |
|---|---|---|---|---|---|
| Random | 65.54 | 7.93 | 24 | 2 | 0.0 |
| Global Investment | 63.57 | 9.8 | 24 | 3 | -0.72 |
| Local Investment | 75.29 | 14.22 | 24 | 0 | 2.9 |

Table 4.1: Expl results: Summary of performance scores from three different aggression functions.

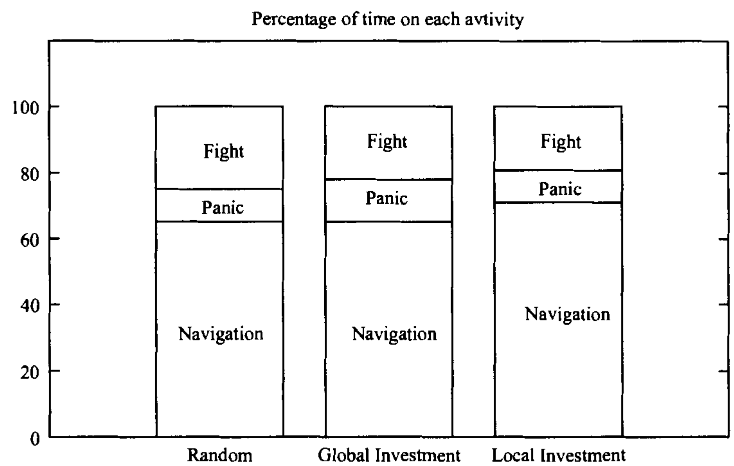


Figure 4.7: Expl results: Proportion of time spent in each activity.

random and global investment, and in this way their standard deviations become smaller. The removal of the outliers did not greatly affect the significance of the t-tests in any case.

The good results obtained by the local investment approach are caused by two properties: (i) As intended by the method, fights occurring in the narrow part of the corridor are won by the robot who has the most to lose. This small advantage accumulated over time and multiple fights makes a significant difference in the performance of the team. (ii) There is an emergent property of the system that is generated by the local investment approach: robots quickly form "chains" or "worms" going in the same direction. Once a worm is created, the frequency of interference is reduced.

A careful reader may notice that the data in the average-number-of-resources column in Table 4.1 is different from the one found in Table 3.1. The reason is that the machine used for the experiments in this Chapter had a different operating system and architecture than the one where the first experiment was run. This is like having a different robot because the real time characteristics of the simulated robot and its controller change. Consequently, it is not right to compare the absolute values but the relations between the results. In both experiments the correlation between the results is the same. That is: (i) random, global investment and local investment perform better than no-fight, (ii) random and global investment perform similarly, and (iii) local investment performs better than the others.

Another experiment, in simulation, was carried out using the local investment method in the second environment (Figure 4.8) described in section 3.4. Note that the world has longer areas of interference and bigger rooms. The results obtained with the local investment approach were statistically similar to those obtained by the global investment approach and better than the ones using the random aggression mechanism. These results can be explained by observing that, in a world with very large areas of interference, the local investment method closely approximates the global investment method. Table 4.2, Figure 4.10 and Figure 4.9 show a summary of the results.

Similarly to experiment 1, the absolute values presented in Table 4.2 are different from the ones in Table 3.2. However, the correlation is maintained. That is: (i) random, global investment and local investment perform better than no-fight, (ii) local investment and global investment perform similarly and better than random.

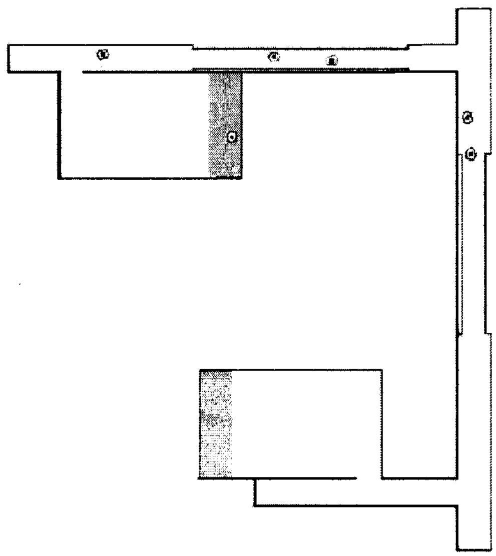

Figure 4.8: A simulated world with longer corridors and more areas of interference.

| TrialType | Mean Resources | $\sigma$ | $N$ | Outliers | $t - test_{random}$ | $t - test_{globalInv}$ |
|---|---|---|---|---|---|---|
| Random | 64.2 | 13.6 | 24 | 0 | 0.0 | -6.21 |
| Global Investment | 90 | 14.3 | 22 | 2 | 6.21 | 0.00 |
| Local Investment | 90.1 | 5.5 | 24 | 4 | 8.49 | 0.04 |

Table 4.2: Expl(big world) results: Summary of performance scores from three different aggression functions.

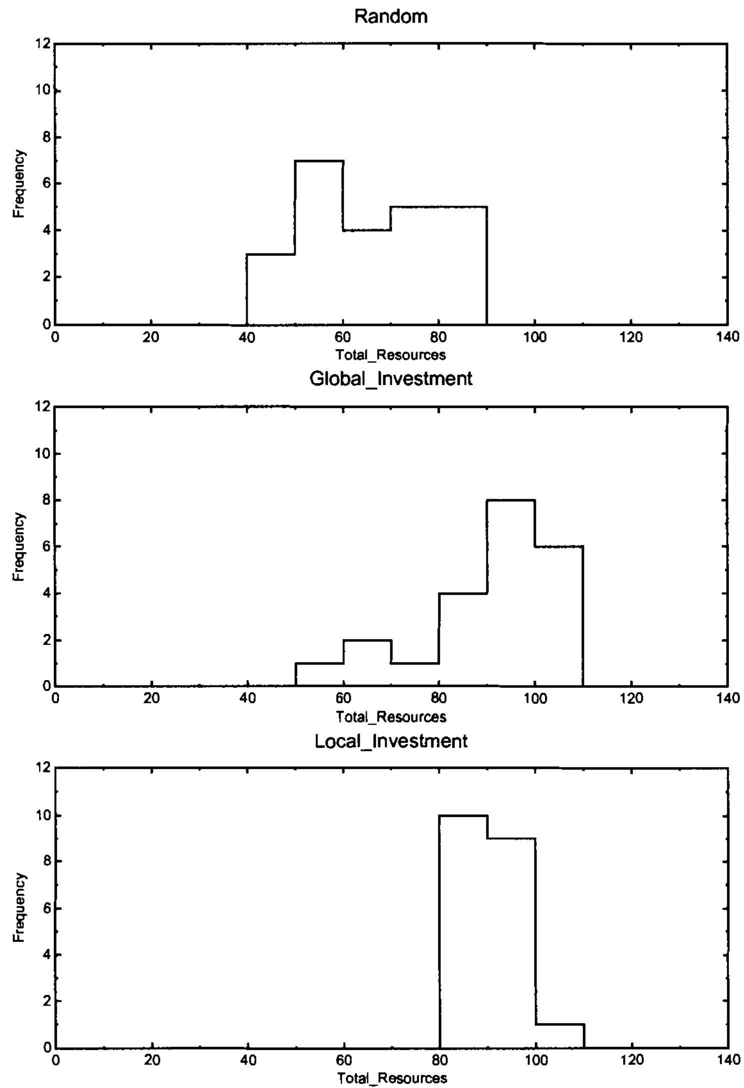


Figure 4.9: Expl(big world) results: Histograms showing distribution of performance scores for three different controllers: random aggression, global investment and local investment aggression.

## 4.5 Experiment 2: Local Investment *vs.* Random Aggression in a Real World Environment

A second experiment was performed in order to (i) verify that a stylized competition can be implemented in the real world with real robots, and (ii) to demonstrate the effectiveness of the local investment strategy in a real world implementation.

This experiment was done in two parts: a simulation to test the experimental procedure and obtain benchmark results, followed by the real world trials. The experiment was smaller-scale than the simulations presented before because only two Pioneer robots were available.

### 4.5.1 Task

The robots must perform laps in an 'O' shaped world, shown in Figure 4.11. The dimensions of the world are 2.85 by 8.5 meters. The narrow corridors are 1 meter wide and the inner block dimensions are 0.85 by 3.5 meters. The dimensions are the same for both the simulated and real world environments. The robots complete loops of the world, returning to their start position on each loop. They advance in opposite directions following their left wall and therefore interfere with each other frequently. Given the shape of the world, the robot going in the clockwise direction makes big loops while the robot going counter clockwise makes small ones. It takes 45 seconds for one robot to complete the small loop, while it takes 66 seconds for the other to complete the big loop.

The 'O world' was designed with the following properties in mind: (i) It is a world that increases the penalty of the wrong robot losing a fight, as the areas of interference constitute the majority of the world (long corridors). (ii) It naturally leads to many fights in a short period of time. (iii) Only two robots were available. In practice, it was found that in 20 minute trials with two robots an average of 30 fights were obtained. All of these characteristics were chosen to demonstrate a difference in performance when using each of the different aggression functions.

The starting condition for each of the trials is always the same, as shown in the simulated world in Figure 4.11 (left).

Figure 4.12 shows a graph of a robot's aggression while completing a loop in the O world environment. The aggression is initially random because the robot is not in an interference area. Then, as soon as the robot starts to move inside the narrow corridor, its aggression starts to increase linearly until reaching the maximum around the end of the narrow part

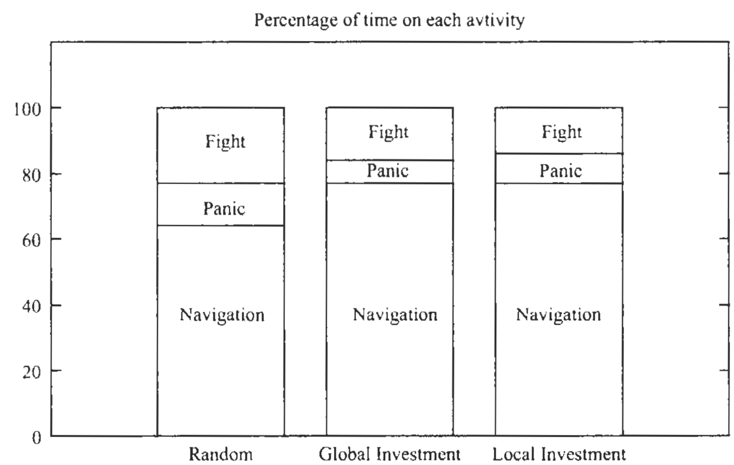


Figure 4.10: Exp1(big world) results: Proportion of time spent in each activity.

Figure 4.11: Simulated (left) and real-world (right) environments. Only one robot at a time can pass across the narrow corridors on the left and the right.

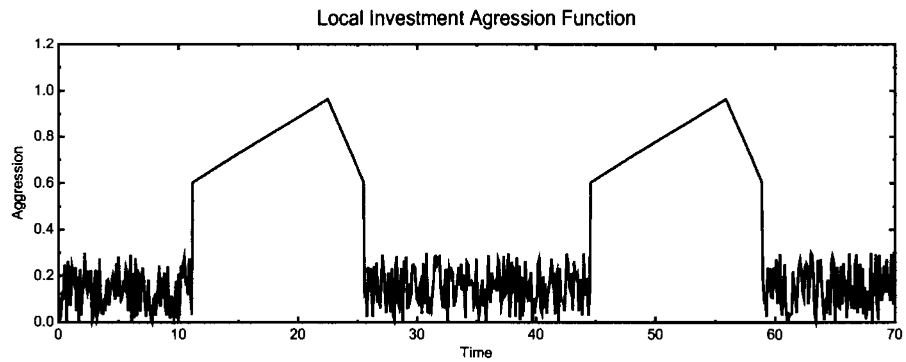


Figure 4.12: Local Investment Aggression through one loop in the O world.

of the corridor. Once the robot is out of the corridor its aggression decreases quickly and when below some threshold it reverts to random aggression. The two non-random sections of the graph correspond to the two narrow corridors present in the world.

### 4.5.2 Control Architecture

The details of the controller used in the 'O world' experiments are different from the one used in the simulation experiments presented in Chapter 3 and Experiment 1 in this chapter. However, the same general control architecture is used (Figure 3.4). The more simple structure of the 'O World' permitted a simplified controller, which proved to be robust; the robots were completely autonomous for the length of the trials. The simulated and real world experiments presented in this experiment made use of the same robot controller.

A simplified robot controller does not mean easier to implement. Certainly the O world is a simplified environment, however, it is a real world with real constraints that cannot be ignored. The main goal of the work presented in this chapter is to show how a local investment aggression function allows a team of robots to achieve more work compared to any previous aggression method. Another goal is to present for the first time an implementation of an aggression mechanism in a real world experiment. In general, it is always easier to develop controllers in simulation because many of the complications of the real world quickly disappear. This is also the reason why normally a controller that works in simulation does not work in the real world. On the other hand, a robust robot controller that works in the real world usually works in a simulated environment.

The idea of the world as its own best model has made quite an impact on the robotics community. Brooks presented this idea as well as the introduction of the subsumption architecture in the 80's [11]. The subsumption architecture emphasizes the situatedness of the robots and therefore has to cope with the real world problems from the very moment a robot and its behaviours are designed.

In the development of the robot controller used in the O world experiments presented in this chapter, several 'real world' problems had to be properly handled. These complications significantly increased the development time of the robot controller. Next is an explanation of the main difficulties encountered and the way they were solved.

### Localization

While in simulation experiments perfect localization is always available, in the real world accurate localization is hard to obtain. Nevertheless, in many cases a robot will need to know where it is in order to decide where to go next.

In real robots, localization based on odometry is quite poor; a method commonly used to get better localization is *Monte Carlo localization*. The problem with this technique is that it requires significant processing power, and does not perform well in symmetrical worlds with long narrow corridors. Another technique commonly used in real world experiments is based on the detection of beacons. This technique is effective but not scalable. It also requires the preinstallation of beacons which may not always be possible, for example, in interplanetary exploration, or missions where human access is difficult.

In the O shaped world used in the experiments in this section of the thesis, I tried with moderate success to use the two techniques previously mentioned. Monte Carlo localization had problems with the long narrow corridors in the O world. Additionally, because of its computation requirements it could not be run in the processors on the robots and had to be executed in an external workstation accessible to the robots via a wireless network. This created a new problem that made the robots too slow in their response time to the dynamics of the environment.

A second technique was through the detection of features in the world. This is similar to using beacons but does not require special sensors. The features in the world that the robots were able to detect were the corners of the inner block, and the walls. The corners allowed the fixing of the position and heading of the robots, while the walls allowed only the fixing of the heading of the robots. Though the method worked well in simulation, in the real world it did not. When panicking, a robot could move and turn unexpectedly and the localization system would lose track of the corner it was detecting (this is because of the symmetry of the world). For example the robot was in the south-east corner, then the robot panicked for some time, went back to navigation and though close to a different corner, the robot would still think that it was at the south-east corner. Similar problems occurred when fixing the heading. This method could have been made more robust through the use of the history of the previous positions, but even that was not fault-proof.

In the end, as a result of all the problems encountered, it was decided to use another interesting feature of the world: given the simplicity of the O world, all that a robot needs to do is to follow the left-hand wall. As a result, localization is not required for navigation, it is only needed for the automatic count of loops completed by each of the robots. This is something that can be done in simulation, where many trials are run, and ignored in the real world, where only a few trials are executed. The counting of trips completed by the robots in the real world experiments was not automated; they were counted by the experimenter.

### Sensors

While in the real world sensors, especially sonars, are quite noisy, in simulation sensors always return perfect values unless models of noise are added [1]. The sonar readings especially depend heavily on the materials of the walls in the world and the angles at which they are hit by the sound beam. A common technique to obtain more credible readings is to use the values of previous readings and in some way detect the ones outside the normal distribution. This has the inconvenient effect of slowing the response time of the robot as decisions based on the readings could be taken only after having received a number of them (effectively a low-pass filter that limits frequency response). The laser on the robots is a less noisy sensor, and, because of this, is the one I tried to use the most. Still, sometimes when two robots face each other, their lasers interfere with one another, and in addition the geometry and location of sensors in the robot did not allow the use of only laser information. The reason is that the laser is located in front of the robot and gives distance readings only in a 180 degree range. The sonar gives information in a 360 degree range but it also has empty areas between the sound beams[2]. Though moving forward is the preferred way of navigation for

[1] Stage, the simulation environment used in the experiments in this thesis, does not have models of noise yet.

[2] A close obstacle may not be detected by the sonar because it is located in an angle not covered by the beams of two sonar rangers.

the robots in the experiment, the fighting behaviour requires robots to navigate backwards as well. Given that the world is designed to cause many fights, a robot spends a good part of the length of each trial doing backward navigation, and as a result, the unreliable sonar sensors have to be used.

### Robot Detection

Robot detection is a central part of the aggression techniques. When two robots face each other they fight, but using only lasers and sonars how does a robot detect another robot in front?. In simulation experiments it is quite simple: the values returned by a laser scan are not only the distance to the obstacles but the intensity of the reflected beams. In simulation it is possible to define a robot to be made of reflective material. Detecting a robot in the laser visual field becomes only a matter of checking a value in an intensity array. In the real world, the reflective material has to be in a specific angle range to the laser beam in order to be detected (almost perpendicular), but lasers in the robots are located at the same height so reflective tape cannot be put in the laser as it would block the laser beams (Figure 4.13). Laser reflective tape is also quite expensive and hard to get.

Using the geometry of the lasers it is possible to detect other robots. In this thesis a Player driver was developed[3]. By using an original laser scan, it provides a new scan in which objects of certain width (i.e. 30 cm for a SICK laser) and at certain distance of other objects (i.e. walls) are removed. With both the original and the new filtered scan, it is possible to detect the angle and location of the objects removed (Figure 4.14). This enables a robot not only to detect and fight other robots but also to interact with humans (it would treat the 2 legs of a person as robots in its path).

The geometrical robot detection technique described here was also used in the simulation experiments. In practice, the robot detection technique was very reliable.

In the next sections there is a description of the different modules of the robot controller.

### Navigate

Given the shape of the O world, a left wall follower is all that is needed to navigate through the environment. Robots do not use a map or any localization information if in the real

[3]The driver can be used by the Player/Stage server [23].

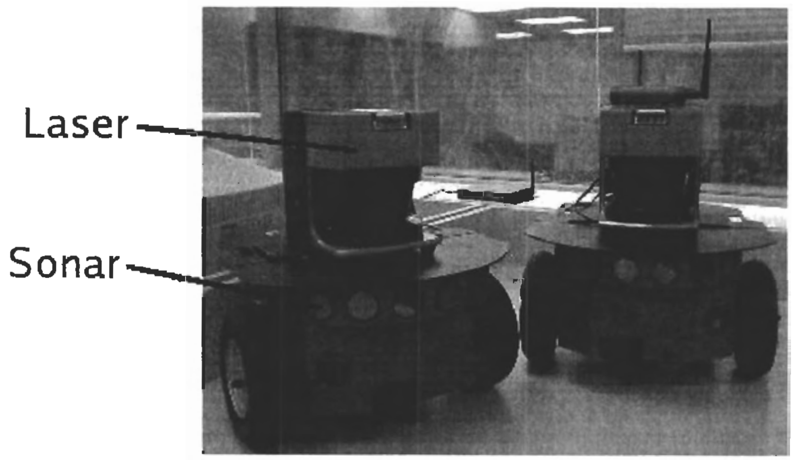


Figure 4.13: Location of the sonar and laser devices in Roy and Priss (the SFU autonomy lab robots). Printed by permission of Richard Vaughan ©.

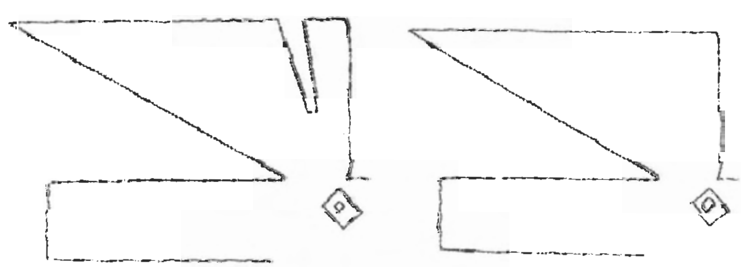

Figure 4.14: Laser scans of a robot with another robot in its laser field. The left scan is the original scan, the right scan is the one in which the robot is removed. By comparing both of them it is possible to know the location of another robot in the robot laser field of view.

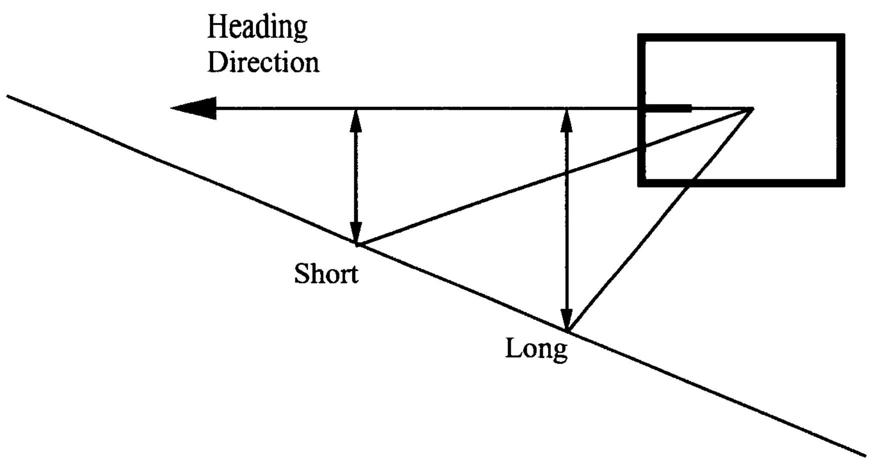


Figure 4.15: Forward Heading Correction. Robot corrects its heading to the right because the front laser distance to the heading axis is shorter than the other one (black double-headed arrows).

world. In the case of the simulations, navigation uses localization information only to correct situations in which the robot would go against its intended rotational direction. Note that if a robot is 180 degrees misaligned, what was before the left wall becomes the right wall. These situations are rare but if present they are generally caused by the panicking behaviour.

The left wall follower is a robust left wall follower; i.e. it knows when to turn left or right at the end of a wall, and also knows how to do it moving forward or backward.

Given that the sensor data, particularly sonar, is noisy, and that the robots would spend a lot of time forward and backward navigating, it was decided to make the forward navigation and the backward navigation quite similar. In Chapter 3 the robust 'variable length sliding box algorithm' was presented. It works extremely well when one has a laser device to guide the location of empty boxes. However, the same cannot be done with a sonar sensor. The algorithm use in this experiment is new and simple. It works in a similar way as a person moving inside a dark room or a maze. The person keeps their hands in constant contact with the wall; this way it knows when to turn or keep going in the same direction. In the case of robots, instead of hands they use the laser data and compare the distances at different angles. Then, the robot can decide if it needs to correct its heading to the left or the right, or if it is time to turn. Figure 4.15 presents a case in which the robot has to correct its heading to the right.

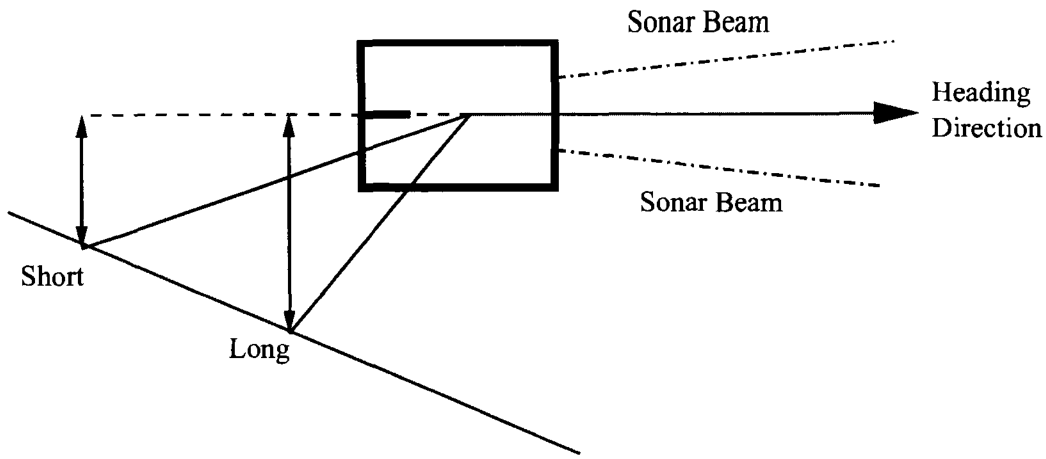


Figure 4.16: Backward Heading Correction. Robot corrects its heading to the right because the front laser distance to the heading axis is shorter than the other one (black double-headed arrows).

There are other parameters that affect the left wall follower: the minimum and the maximum distance to the left wall. If the robot distance to the wall is between the minimum and the maximum distance then the heading corrector presented in the previous paragraph is used. If the robot's distance to the wall is smaller than the minimum distance, the heading is corrected to get away from the wall; the opposite is done if the robot's distance to the wall is greater than the maximum distance. All of this has the overall effect of maintaining the robot at a safe distance from the wall while providing smooth navigation.

When the robot has to navigate backwards, the sonar is used to detect the back proximity to obstacles, and the laser is used to correct the back heading in the same way as the forward case (Figure 4.16). Backward navigation is more difficult for the robot because of the use of the sonar device, especially when deciding to turn left or right.

### Panic

The panic behaviour is also simplified to two sub-behaviours: a 'corrective panic' which corrects the heading of a robot when, for example, the robot does a turn and gets too close to the wall. The response is then just to move in the opposite direction and turn a little. If this is not enough to solve the problem, a second sub-behaviour 'force panic' is used. In this case, the sonar readings are used as an array of forces acting on the robot and proportional to the distance. The robot moves in the direction of the resultant force.

#### Fight

This behaviour is responsible for deciding how to solve interference problems between robots.

The fighting behaviour used in the simulation experiments presented in Section 3.1 did not work well in the real world because it required too much free space behind the robots in order to decide which robot won a conflict. This is because the aggression level of a robot was used to generate a tolerance distance. But two different aggression levels needed to be perceived as two different tolerance distances. For example, if 10 aggression levels are used and 0.4 meters are necessary to differentiate between each one, then at least $0.4 \times 10 = 4$ meters plus some safety distance would be required in order to allow the robots to detect who won a fight safely (these are the actual values used in the simulation experiments in Chapter 3). Reducing this distance causes the robots to get into oscillations. That is, two robots with similar aggressions start to move back and forth because they both think they won or lost a fight. Another problem of the method is that it wastes energy because the robots have to back up a long distance. The 'O World' was too small for this technique to be used.

To solve the problem an alternative scheme has been developed. In the 'staring contest' scheme the aggression level is converted to a waiting time. The method is simple: when two robots going in opposite directions find the other robot too close they both stop. They then use their aggression level to calculate how much time they are willing to wait before retreating (losing the fight). While waiting, they are constantly checking the distance to the other robot. If they perceive that the robot is backing up then they know they won the fight and move forward. This method has several benefits over the previous back-up fight: (i) it is more energy efficient; (ii) it requires less space to manoeuvre; and (iii) robots do not have to back up too much and as such the navigation is less prone to get in situations that require the use of the panic behaviour.

### 4.5.3 Performance Metric

In the previous experiment the work done by the team of robots was directly proportional to the number of trips they completed. In this case however, there is a difference in time or energy spent between big and small loops. Small loops were chosen to be equivalent to the

| Trial Type | Work | | | Trips | | $R_{bigloop}$ | | $R_{smallloop}$ | |
|---|---|---|---|---|---|---|---|---|---|
| | Total | $\sigma$ | $t - test_{random}$ | Total | $\sigma$ | Total | $\sigma$ | Total | $\sigma$ |
| Random | 34.4 | 1.39 | 0.0 | 30 | 0.6 | 9.6 | 0.59 | 20.4 | 1.14 |
| Global Inv. | 34.2 | 1.57 | -0.13 | 28.8 | 0.75 | 12.0 | 0.9 | 16.7 | 1.2 |
| Local Inv. | 41.7 | 1.82 | 14.2 | 34.7 | 0.51 | 15.4 | 0.5 | 19.3 | 1.37 |

Table 4.3: Exp2(O world) results: Summary of performance scores from three different aggression functions.

completion of one unit of work; using the difference in time taken to complete the big and the small loops, a conversion factor was obtained and used to calculate the total work done by the team of robots (Equations 4.5). With this method, the results of different trials and different aggression functions can be compared.

$$work_{cf} = \frac{time_{big}}{time_{small}} = \frac{66}{45} = 1.46 \tag{4.4}$$

$$work = Trips_{small} + Trips_{big} \times work_{cf} \tag{4.5}$$

### 4.5.4 Simulation Results

In the simulation test a total of 20 trials of 20 minutes were run for the random, global and local investment aggression functions.

The simulation results in Table 4.3 show that while the random and global investment approaches perform similarly, the local investment approach is better. Even the number of trips completed by the robots using local investment aggression is greater than when using random or global investment. The histograms presented in Figures 4.18, 4.19 and 4.17 present information about the trips completed by each of the robots plus the team performance. It can be seen that the local investment approach enables the robot doing big loops to complete more loops than when using a random or global investment approach, and this only slightly reduces the number of trips executed by the robot doing the smaller loops. This emergent property in the system increases the amount of work performed by the team, since the completion of big loops is related to the achievement of more work.

Figure 4.20 provides two interesting observations: (i) there is virtually no time spent in panicking, which suggests that the navigation and fighting behaviours are quite robust, and

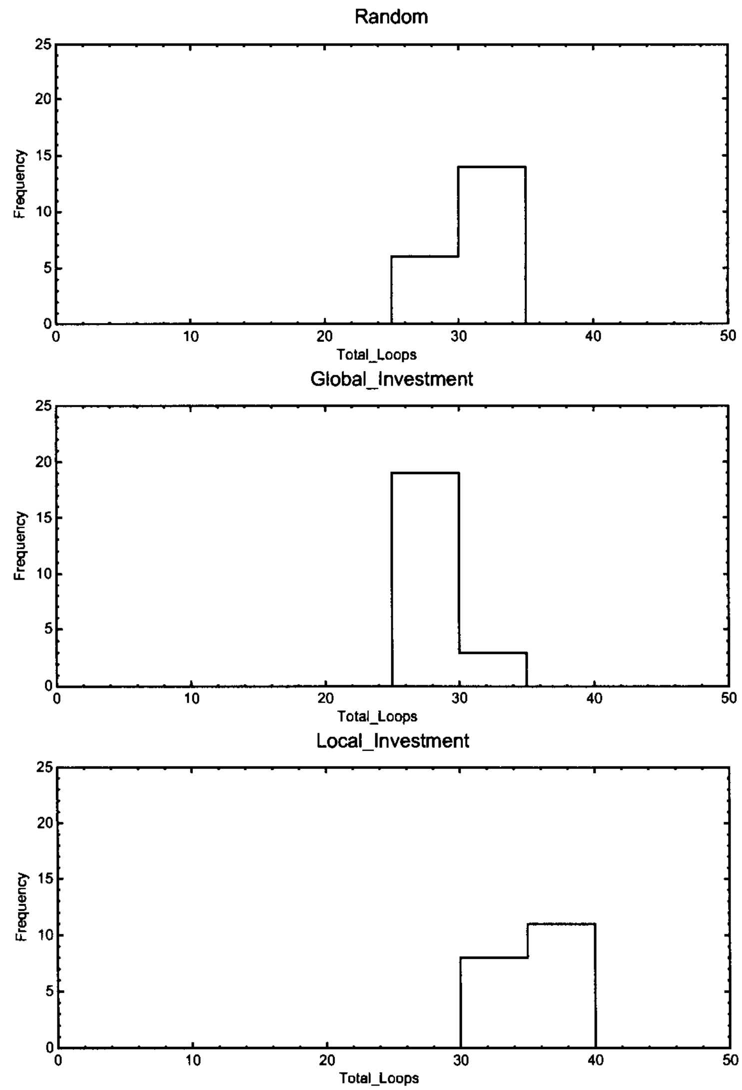


Figure 4.17: Exp2(O world) results: Histograms showing distribution of performance scores for three different controllers: random aggression, global investment and local investment aggression.

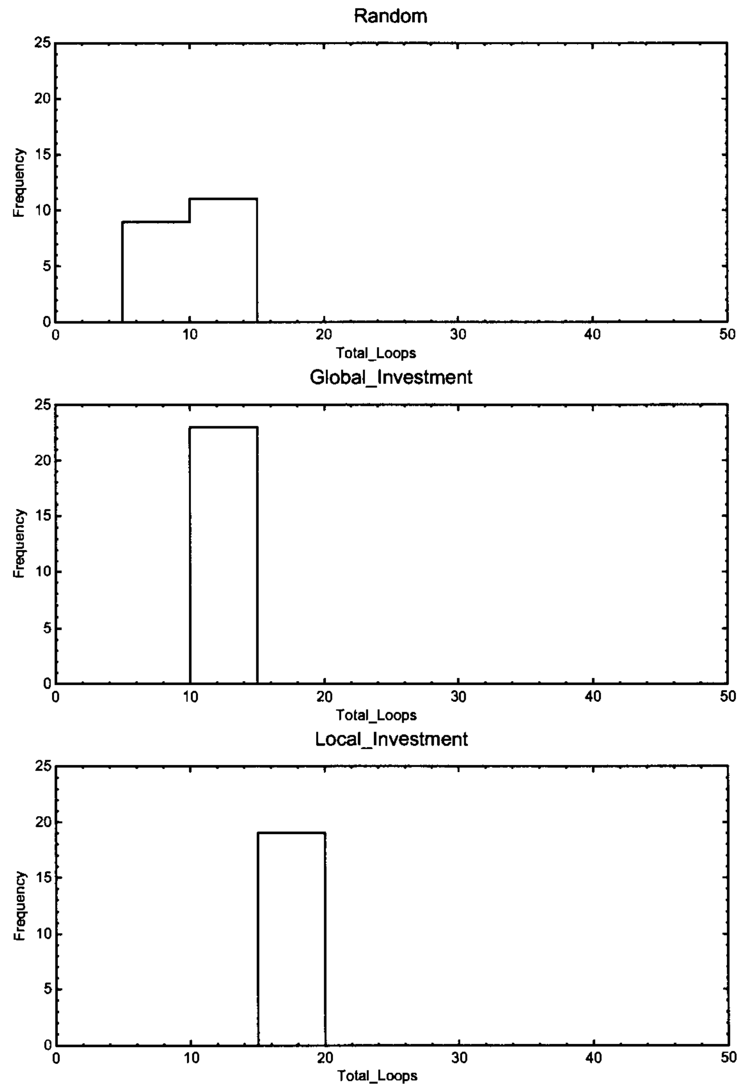


Figure 4.18: Exp2(O world) results: Histograms showing distribution of performance scores for three different controllers for Robot 1 (big loops).

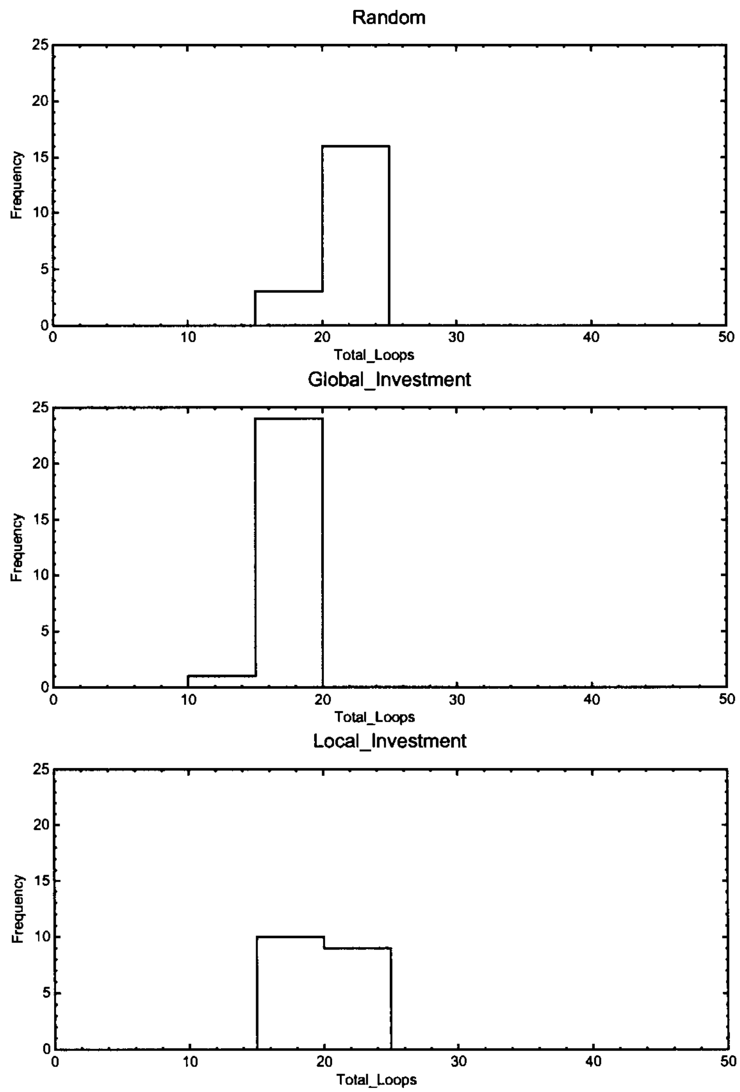


Figure 4.19: Exp2(O world) results: Histograms showing distribution of performance scores for three different controllers for Robot 2 (small loops).

| Trial Type | Total Work | Total Trips | $R_{bigloop}$ Trips | $R_{smallloop}$ Trips |
|---|---|---|---|---|
| Random 1 | 29.1 | 25 | 9 | 16 |
| Random 2 | 31.6 | 27 | 10 | 17 |
| Random 3 | 29.6 | 25 | 10 | 15 |
| Random Avg | 30.1 | 27.3 | 9.6 | 16 |
| Local Investment 1 | 39.3 | 32 | 16 | 16 |
| Local Investment 2 | 38.3 | 31 | 16 | 15 |
| Local Investment 3 | 37.9 | 31 | 15 | 16 |
| Local Investment Avg | 38.5 | 31.3 | 15.6 | 15.6 |

Table 4.4: Experiment in the real world results: Summary of performance scores for random and local aggression.

(ii) for the three aggression functions evaluated there is almost no difference in the time spent in fighting and navigation; they are quite similar in fact. In the case of the O world where there is a difference between the tasks that each of the robots achieve, big loops and small loops, the navigation time seems to not be directly related to the amount of work that the system does. This goes against the hypothesis presented in Chapter 3, however, in the work in that chapter all robots had to complete the same task (there were no different loops).

Having shown that the performance of global investment is indistinguishable from random in this environment, only local investment and random are compared in the real robot trials.

### 4.5.5 Real World Results

In this experiment a total of 3 trials of 20 minutes for the random, and local investment aggression functions were run.

Table 4.4 presents the results obtained in the real world. Due to the small sample size, no standard deviations or histograms are given for these tests. However, it can be observed that the results are similar to those obtained in simulation, and that aggression based on local investment always produces more trips and work done than random aggression.

## 4.6 Conclusion

In experiments done in the environment shown in Figure 3.2 different strategies for choosing the level of aggression of a robot were explored. However, none of the rational methods (Hierarchy, Personal Space, Inverse of Personal space) were shown to be statistically different than random. In [14] another rational method based on the concept of 'global investment' or effort put into a task was presented. This method, however, also failed to show any improvement over a random approach. In the case of global investment a new world had to be created in order to show that the method could perform better than random. In the work presented in this chapter a new method based on 'local investment' has been presented. The method was shown to improve the performance of a team of robots in the environment where all the other methods had failed and as such it is one of the main contributions of the work presented in this thesis.

In all the environments tested the local investment method performed better than or as well as all other methods. This suggests that the method is robust in a number of types of environments and, therefore, does not require special assumptions about them. Nevertheless, it is certainly possible to design a world in which the local investment method would fail compared to random or other rational methods. A hypothesis not investigated in this thesis is that worlds in which the local investment would fail may not be very realistic.

The local investment mechanism is decentralized and independent of a navigation strategy as can be seen in the experiments done in the O shaped world where a different robot controller was used and still the advantages of the local investment method were shown. Only existing sensors (sonar and laser) have been used, the robots do not require id's or any network device. Though not tested in other robot platforms, the local investment method should be usable by any robot with similar sensing capabilities to the Pioneers-3DX used in the experiments here.

Experiments done in the O shaped real world show that the local investment scheme can be successfully implemented and that the computation of aggression is trivial to do (only a counter that is incremented when in narrow sections of the world).

Simulation has been used extensively in the development and testing of the robot controllers that were used in the real Pioneer-3DX robots. In the end, the robot controller used for both, the simulation and the real world experiments, was the same. It only differed in some tuning parameters (distances to obstacles and other robots, etc). The results obtained in simulation were carried to the real world experiments. Though the numbers were not an exact match, they were similar, and, most importantly, the relations between who did better and worse were the same.

For the length of the trials the robots in simulation and the real world experiments were completely autonomous; they did not require human assistance at all. This is part of a long-term goal of having robots in the world doing real tasks in an autonomous way.

The main limitation of the local investment approach is that is not anticipatory and does not hold a global picture of the situation in the world, and so the method is short-sighted. It does not encode any information about past interactions or the probability of other robots in other places in the world. Because of this, the method may not achieve the highest performance for a team of robots doing a task. However, a requirement to include anticipatory behaviour and global pictures of the world would definitely have an impact on the complexity and scalability of any aggression approach. In defence of the local investment approach, it is simple to compute and very scalable.

In summary, a novel aggression function based on the idea of local investment has been presented. An implementation was described, and the method was shown to increase the performance of a simulated robot team at a classical transportation task and, in a simplified form, in real life with a team of two robots. Though this method has some limitations, it should be widely applicable as an interference reduction technique in mobile robot teams.

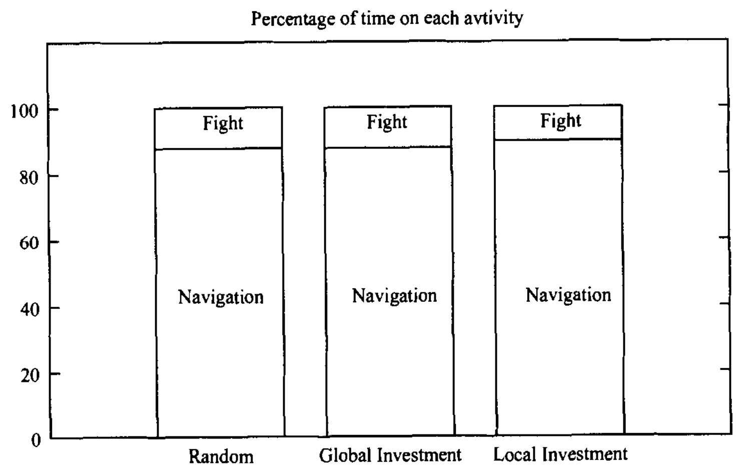


Figure 4.20: Exp2(O world) results: Proportion of time spent in each activity.

# Chapter 5

# Discussion and Future Work

## 5.1 The Environment

Robots are not only embodied but situated agents. They live in a world full of constraints that will impose its own limits to what can be achieved or not.

The local investment method proved to be better than other approaches in most of the worlds tested. Still, it may be possible to design an environment in which the scheme would not work properly.

In general, if robots are going to share space with humans it makes sense that their goals should be equally obtainable by humans. An architect will not design a building in which one thousand people have to compete to get in or out the building through a door where only one person can fit; on the contrary it is almost always the case that the environments where we live or work have been designed in a way that interference is reduced. Robots working in this type of environment would benefit from the design principles we have created for ourselves.

It may be the case that a local approach is good enough when dealing with human environments. For this to be confirmed many experiments should be performed in realistic human environments.

## 5.2 The Use of Communication in Worm Fights

In the experiments presented in this thesis it was found that one of the emergent properties of the rational schemes used was the formation of worms of robots. These worms of robots

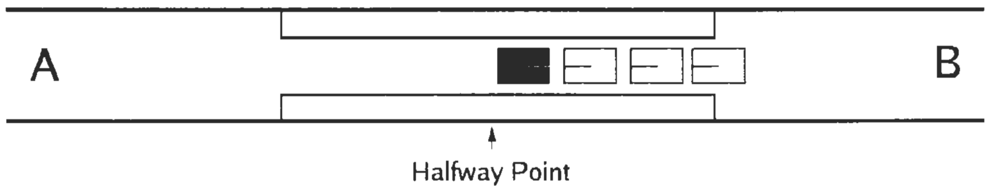


Figure 5.1: The left-most robot is travelling to the right while a worm of white robots is travelling to the left, it is more economically rational to allow the worm of robots to win the fight, however the local investment and global investment methods would make the black robot win the fight.

would usually last for some time and then break apart, mainly because of the shape of the environment and fights with robots going in opposite directions. The robots, however, were blind to the fact of being part of a worm or not and therefore incapable of using such property to set their aggression.

[24] and [45] support the idea that in some pack animals like lions and baboons, the size of the pack affects the level of aggression of the individuals. This ability requires two characteristics to happen: (i) they need to recognize members of the pack, and (ii) they need a way to read and communicate their intentions to other members of the pack.

[52] presents some interesting theories into the roles of 'own' (the hyenas' pack) and 'others' (other predators like lions, etc). Their evolution experiment showed that the most successful agents were not only the most aggressive but also the ones capable of differentiating between own and others.

Following nature's suggestions, a worm of robots could have an aggression equal to a function of the aggressions of its individuals; this would give a real advantage in fights against individual robots (Figure 5.1). To enable this type of behaviour the robots, similar to hyenas, would need to recognize the members of their worm and their aggression levels. This could be implemented in the real world by means of fiducial id's and communication capabilities like a network or some more stylized sensing capability (after all, hyenas do not have any wireless network to communicate, they do it through their sensorial inputs, smell, vision, hearing etc.).

The hypothesis is that a communication scheme in which robots in a worm can add up their aggression and make the first robot in the worm as aggressive as the pack would have a direct impact in the performance of a team of robots, either because the system saves energy or because more resources are transported in a unit of time.

## 5.3 Different Types of Robots

The aggression scheme proposed in this thesis has been tested in real robots. However, the robots were the same (Pioneer-3DX). One of the hypotheses presented in the methods used in this thesis, as well as in [57], is that the scheme could be used by robots with similar sensory-motor capabilities. This is a test that has not been done yet but should work correctly.

Another more interesting option would be to mix different types of robots as part of the same team. Previously, the utility of aggression in real animals has been discussed but animals do have one more characteristic that the experiments with the Pioneers did not have; in nature the aggression level of an animal usually is indicative of its condition and possible performance (in gathering food, or securing an area, etc). The leader of a pack of wolves is the strongest of them. In the case of the experiments presented here there is no difference in the capabilities of the robots they all are clones and as long as they have new batteries, they all perform the same way.

If robots with different configurations (speed, energy consumption, etc), have to work in a team then the complexity of defining how to set an aggression function is increased. This is probably a more realistic case in the future where different types of robots will have to coexist and work together. It also opens up the space to include humans or animals as part of the cooperative force.

## 5.4 A Theory of Mind

In a multi-robot experiment, it is not a requirement that all robots set their aggression in the same way even if they are physically similar (i.e. all of the robots are Pioneer-3DX). How would it be possible to have a team of different robots in which their aggression functions are unknown for the other robots, and yet achieve the maximum amount of work possible? Such system would need to have learning capabilities. Robots could learn the way others fight through their interactions and slowly increase the performance of the team. A small extension to such method could enable robots to anticipate the outcome of a fight and give the right of way without ever fighting, if knowing that with a high probability the battle was going to be lost. The prediction capability would improve the performance of a system from a conservation of energy perspective, and perhaps even in the amount of work done by the team. Animals do possess this type of anticipatory behaviour when fighting, an example being the case of cheetahs abandoning hunted prey to hyenas without even bothering to defend it.

The hypothesis is that a robot that can anticipate the way another robot is going to fight, will be more successful than one that does not have such ability.

## 5.5 Other forms of fight

In this thesis, two different implementations of fights were shown, the first one inherited from the work in [57] where it was named 'inverse chicken game'. This type of fight was used in all the simulation experiments except the ones in the O world, because the fight had some problems with small environments (See section 4.5.2). A second type of fight was created based on the concept of a 'staring contest'. Experiments in the real world proved the feasibility of fighting and the possibility of using either of the fighting techniques in real situations.

Another aggression communication technique explored but not used was through the use of sound. This is interesting because it also seems to be used in nature (i.e. roaring contests in red deer [16]). The loudest or highest frequency could mean the strongest in a confrontation.

It is nevertheless the case that many other abstractions of a fight could be implemented in real life. This is an open area for research that should be explored.

## 5.6 A global fight

Despite its success, the local investment scheme has the main limitation of being local and therefore blind to what is going on in the whole world. Other rational methods presented in this thesis and in [57] have the same problem. If robots had a map of the world and knew the location of themselves and all the other robots, more efficient ways of solving fights could be obtained. In defence of the local approach, it has the benefit of being almost infinitely scalable because there are no requirements for communication, id's or special sensors. A global method, on the other hand, would require communication and id's for the robots and therefore would not be feasible for large teams of robots.

Still, in situations where the number of robots is known to be small, approaches to fighting that involve global planning should perform better than any local method.

## 5.7 Territorial Subdivision

In [21] a territorial multi-robot task division was presented. The work in this thesis as well as in [57], assumed that robots were not able to pass resources from one to another. But what if that were not the case? What if robots could indeed pass resources from one to the other?

Given an environment and a number of robots working in it, it may be that to maximize the work done by a team of robots, a combination of territorial division and the fighting schema presented in this thesis is required. This is an area not yet explored and at a first glance looks like a complex problem to solve.

## 5.8 Evolution and the Theory of Games

In Maynard-Smith [40], a study into game theory and its relation to animal behaviour is presented. Most of the applications of evolutionary game theory in the book are directed towards animal contests and therefore are quite related to the problem of multi-robot systems working in the same space, as is presented in this thesis.

In the different methods for selecting the robot's aggression presented in this thesis, their goodness is supported by the execution of different experiments that measure the performance of each technique. Questions arise such as: What if one of the robots decides to cheat and always displays as a very aggressive agent? What is the best way of setting an aggression when different robots may follow different strategies? What if robots know their aggressions beforehand?, etc. It may be that these scenarios can be modeled as games, if this is the case, game theory could provide answers to these questions. In [53] game theory is used to prevent parasitic behaviour in an ad-hoc network; a node/agent will do equally well if cheating than if not cheating, therefore there is no advantage in cheating. The use of game theory is not explored in this thesis and considered out of scope. Nevertheless, it should be approached in future research.

Also in [40], it is shown that in a conflict of "Doves against Doves", a Dove will withdraw from a contest against another Dove at a random time causing each Dove to win half of its battles. However one important characteristic of the randomness of the times that each Dove waits before withdrawing from the contest is that it follows an exponential distribution. This causes the majority of conflicts to be solved quickly, while a small amount of them take longer, reducing the costs associated to the display behaviour. In the work here and in [57], the level of aggression is set randomly between a minimum and maximum values, but the distribution is linear. This suggests a new hypothesis; that the use of an exponential distribution will improve the performance of a multi-robot system that uses random aggression.

Another important characteristic of the experiments in this thesis is that the number of robots in the system is quite small; other authors have considered types of games in which the populations are finite and small [49].

In the case of the local investment method, it could be approached from a theoretical game perspective as a contest with 'variable rewards'. In this type of contest each robot inside an area of interference knows the cost of the resource. For example, if a robot has spent 10 seconds passing a narrow corridor, the cost of the resource would be the time of backing up plus the time of coming back to the location where it originally was. This can be approximated to twice the time it has already spent (20 seconds). If another robot coming in an opposite direction has only spent 5 seconds, the cost for that robots would be 10.

Another possibility for the type of problem shown in this thesis is to view every fight as a competition between an agent and the rest of the population. This type of conflict is called 'playing the field', the rules of these conflicts are different to those for one-to-one conflicts.

## 5.9 Standard Deviation in the Results

In some of the experiments presented in this thesis large standard deviations were obtained. Next is a list of the reasons why this is the case:

1. The robot controller used in the simulation experiments in Chapter 3 is not extremely robust. In defence, one could say that it is easy to have a single robot moving in an environment. It is quite a lot more complicated to have a large number of robots in an environment in which the designer intentionally wants the robots to interfere a lot with each other. This is a special case of the experiments in this thesis and generally not the case for most other robot experiments.

2. One of the goals of the work presented in this thesis was to show that a rational aggression mechanism allows a group of robots to achieve more work. It is not the goal (at least in the first part of the experiments) to develop the most robust navigation controller. This is clearly something quite complicated on its own and not the purpose of this thesis.

3. Though the starting condition of the experiments was always the same, the simulator worked in an asynchronous way. This made the results vary from trial to trial.

4. Through the execution of more and longer trials it would have been possible to reduce the standard deviation. Because of time issues this was not feasible.

Even with some large standard deviations the results presented in this thesis are statistically significant.

# Chapter 6

# Conclusion

Next is a list of original goals of this thesis followed by a section comparing the goals vs. the results obtained in the different experiments.

Goals:

1. Find a rational aggression mechanism that is decentralized, independent of a navigation strategy, makes use only of existing sensors (sonar and laser), works in heterogeneous robots systems, and is simple to compute.
2. Show that a rational aggression mechanism outperforms a random scheme in a team of robots.
3. Show that the robots using the aggression mechanism are autonomous and do not require human intervention.
4. Demonstrate that a stereotypical competition, inspired by animal behaviour, can be implemented in the real world.
5. Show that the results obtained in simulation are carried to the real world.
6. Show that a rational aggression mechanism is robust and performs well in different world configurations.

The rational aggression methods proposed in this thesis are both decentralized and independent of a navigation strategy as can be seen in the experiments done in the O shaped world where a different robot controller was used. Only existing sensors (sonar and laser) have been used, the robots did not require the use of id's or any network device. Though not tested in other robot platforms, the investment methods should be usable by any robot with similar sensing capabilities to the Pioneers-3DX. [Goal 1]

The only additional requirement for the implementation of the investment based methods, when compared to the previous best approach (random aggression), is the use of a counter that, in the case of global investment, is increased until a goal is reached, and in the case of local investment, is increased while in a narrow space. In both cases the methods are easy to compute and the overhead on the robots is minimal. [Goal 1]

In this thesis the best method found for deciding the aggression of a robot, 'local investment', has been shown to perform better than any other previous method including global investment and non-rational methods. The global investment method also performed better than random but several modifications to the world were required. This suggests that the global investment method may not be robust to real world situations. On the contrary, the local investment method proved to be good in three different environments. It is believed that the local investment method is more robust and could be used in many real situations. As discussed previously, experiments in many other environments should be done in order to be more certain about the robustness of the local investment approach. [Goals 2 and 6]

For the length of the trials, the robots in simulation and in the real world experiments were completely autonomous; they did not require assistance at all. This is part of a long-term goal of having robots in the world doing real things in an autonomous way. Also, for the first time a stylized competition inspired by animal behaviour has been implemented in real robots. The first type of fight, based on the concept of 'inverse chicken game', had spatial disavantages that required the use of a different type of fight. The 'staring contest' fight was able to cope with the limited space in the O shaped world and proved to be robust for the total length of the experiments. [Goals 3 and 4]

Simulation was used extensively in the development and testing of the robot controllers used in the real Pioneer-3DX robots. In the end, the robot controller used for both the simulation and the real world experiments was the same, it only differed in some tuning parameters (distances to obstacles and other robots, etc). The results obtained in simulation were carried to the real world experiments. Though the numbers were not an exact match, they were similar and, most importantly, the relations between who did better or worse in simulation were the same. [Goal 5]

In summary, the best rational aggression method found so far, based on the idea of local investment, has been presented in this thesis. The method was shown to increase the performance of a simulated robot team at a classical transportation task and, in a simplified form, in real life with a team of two robots. The local investment method worked well in three different environments. Also in this thesis, the first real implementation of a stylized competition to solve spatial interference problems has been presented. Though the methods and techniques presented here have some limitations, they should be widely applicable.

# Appendix A

# Brown's thesis

Brown is a co-author of [14] which is highly referenced in this thesis, especially in Chapter 3. Her thesis is complementary to my work and is being produced simultaneously as part of the Rage project [47]. Her thesis abstract is reproduced here as an indication of the scope and goals of her work.

### Using Emotion and Mood in Robots to Cope with Dynamic Environments

**Abstract of MSc thesis Simon Fraser University, 2005**

**Sarah Brown**

Emotions in humans are responses to the environment and are often very useful. For example, the fight or flight response to something frightening prepares people to defend themselves or run for self-preservation. While emotions have a direct cause, mood is more diffuse and can be thought of as a more generalized feeling which relates to how well things have been going for a person. Mood tends to affect a person's emotional response to their environment. Typically, though not always, negative moods foster negative emotions, and vice versa.

This thesis describes a controller, built upon the research in [57], which incorporates the need for energy into the transportation task. Unlike this previous work, the environment in which the robot must operate changes over time as robots and obstacles are added into and removed from the environment. In order to cope with the changing environment, the controller incorporates the notion of emotion and mood.

Emotions are triggered by specified events in the system, such as transporting an item or encountering another robot at the recharging station. These emotions are then used to not only select the most appropriate response to the event, such as the level of aggression in a fight, but contribute to the mood of the robot.

The purpose of mood is to provide a long term notion of how difficult the environment has been for the robot. Mood is used to then affect how emotions are triggered and to what degree, much like in people. Mood is also used to regulate transitioning between the transportation task and the need for energy. A bad mood indicates that the environment has been more challenging than normal and so the robot should recharge more often to prevent the batteries from running out due to difficulties in reaching a charging station.

A standard controller with a static threshold for recharging is compared against the same controller, but with emotion and mood elements in a simulated environment and evaluated with respect to how many items are transported.

# Appendix B

# Zhang's thesis

Zhang is a co-author of [14] which is highly referenced in this thesis, especially in Chapter 3. His thesis is complementary to my work and is being produced simultaneously as part of the Rage project [47]. His thesis abstract is reproduced here as an indication of the scope and goals of his work.

### A reinforcement learning method for choosing the best aggression

**Abstract of MSc thesis Simon Fraser University, 2005**

**Carl Zhang**

A team of decentralized robots performing a transportation task in a shared space may suffer from spatial interference, which can severely reduce the effectiveness of the team performance. In [57], a stereotyped competition was applied to break the deadlock when two robots had a face-to-face confrontation in a narrow space that allows only one robot to pass. However, regarding the improvement of overall team performance, none of the methods demonstrated in [57] for selecting a robots aggression level performed better than selecting aggression at random. In my research, I am trying to find out the correlation between the aggression of a robot and the current state of its external environment, modeled by the robots sensor readings and total time it has invested in the transportation task. A recurrent network is used. The network takes a robots sensor readings as inputs and outputs the robots aggression. Feedback connections of the recurrent network make it possible for the robot to have a short-term memory of its previous states. Weights of the edges that connect input nodes to output node can be adjusted when the state of the external environment changes. Reinforcement learning is applied, and it allows a robot to learn to choose the best aggression selection strategy during transportation task via trial-and-error. Also, an environment that is more realistic and complex than the one in [57] is used.

# Appendix C

# T–Tests: Distribution of t

Table C.1: Distribution of t (two tailed)

| | Probability | | | | | |
|---|---|---|---|---|---|---|
| d.f. | 0.5 | 0.1 | 0.05 | 0.02 | 0.01 | 0.001 |
| 1 | 1.000 | 6.314 | 12.706 | 31.821 | 63.657 | 636.619 |
| 5 | 0.727 | 2.015 | 2.571 | 3.365 | 4.032 | 6.859 |
| 10 | 0.7 | 1.812 | 2.228 | 2.764 | 3.169 | 4.587 |
| 15 | 0.691 | 1.753 | 2.131 | 2.602 | 2.947 | 4.073 |
| 20 | 0.687 | 1.725 | 2.086 | 2.528 | 2.845 | 3.85 |
| 25 | 0.684 | 1.708 | 2.06 | 2.485 | 2.787 | 3.725 |
| 30 | 0.683 | 1.697 | 2.042 | 2.457 | 2.75 | 3.646 |
| 40 | 0.681 | 1.684 | 2.021 | 2.423 | 2.704 | 3.551 |
| 60 | 0.679 | 1.671 | 2 | 2.39 | 2.66 | 3.46 |
| 120 | 0.677 | 1.658 | 1.980 | 2.358 | 2.617 | 3.373 |

# Appendix D

# Player/Stage

All the experiments performed in this thesis are done using the Player/Stage robot development and simulation system [23].

Player is a commonly-used server and abstraction protocol [56] that connects a user-defined control program to the sensors and the actuators on a mobile robot. Stage is a robot simulator that provides multiple virtual robot devices to Player.

The Stage models approximate ActiveMedia Pioneer-3DX robots. Their dimensions are 44cm long and 33cm wide, and they are equipped with front and rear sonar rings, and a SICK laser range finder.

Player and Stage are freely available under the GPL from:
`http://playerstage.sourceforge.net`.

# Bibliography


[1] R. C. Arkin and T. Balch. Cooperative multiagent robotic systems. In D. Kortenkamp, R. P. Bonasso, and R. Murphy, editors, *Artificial Intelligence and Mobile Robots*. MIT/AAAI Press, Cambridge, MA, 1998.

[2] Angelo Arleo and Wulfram Gerstner. Modeling rodent head-direction cells and place cells for spatial learning in bio-mimetic robotics. In *FROM ANIMALS TO ANIMATS 6: Sixth International Conference on Simulation of Adaptive Behavior*, pages 236–245. MIT Press, 2000.

[3] Robert Axelrod. *The Evolution of Cooperation.* Basic Books, New York, NY, 1984.

[4] Randall D. Beer. A dynamical systems perspective on agent-environment interaction. *Artificial Intelligence*, 72:173–215, 1995.

[5] Randall D. Beer. Dynamical approaches to cognitive science. *Trends in Cognitive Sciences*, pages 91–99, 2000.

[6] Randall D. Beer. Autopoiesis and cognition in the game of life. *Artificial Life*, pages 309–326, 2004.

[7] Randall D. Beer, Roger D. Quinn, Hillel J. Chiel, and Roy E. Ritzmann. Biologically inspired approaches to robotics: what can we learn from insects? *Communications of the ACM*, 40(3):30–38, 1997.

[8] Peter Bonasso, David Kortenkamp, and Robin Murphy. *Mobile Robots and Artificial Intelligence*, chapter Mobile Robots: A Proving Ground for Artificial Intelligence, pages 3–18. AAAI Press, 1998.

[9] C. Breazeal. Emotion and sociable humanoid robots. *Human-Computer Studies*, 59:119–155, 2003.

[10] C. Breazeal and B. Scassellati. How to build robots that make friends and influence people. In *Proc. IROS99, Kyonjiu, Korea.*, 1999.

[11] R. A. Brooks. A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation*, 2:14–23, 1986.

[12] R. A. Brooks. Elephants don't play chess. *Robotics and Autonomous Systems*, pages 3–15, 1990.

[13] R. A. Brooks. Intelligence without representation. *Artificial Intelligence Journal*, 47:139–159, 1991.

[14] Sarah Brown, Mauricio Zuluaga, Yinan Zhang, and Richard T. Vaughan. Rational aggressive behaviour reduces interference in a mobile robot team. *The 12th International Conference on Advanced Robotics ICAR2005*, 2005.

[15] A. Clark and J. Toribio. Doing without representing? *Synthese*, 101:401–431, 1994.

[16] T. H. Clutton-Brocknext, S. D. Albon, R. M. Gibson, and F. E. Guinness. The logical stag: Adaptive aspects of fighting in red deer (cervus elaphus l.). *Animal Behavior*, 27:211–225, 1979.

[17] D.C. Dennett. Why not the whole iguana? *Behavioral and Brian Sciences*, 1:103–104, 1978.

[18] D.C. Dennett. Cognitive wheels: The frame problem of ai. *Minds, Machines and Evolution*, pages 129–151, 1984.

[19] L.M. Dill and A.H.G. Fraser. The worm returns: hiding behavior of a tube-dwelling marine polychaete, serpula vermicularis. *Behavioral Ecology*, 8:186–193, 1997.

[20] Dario Floreano and Stefano Nolfi. Adaptive behaviour in competing co-evolving species. In *4th European Conference on Artificial Life*, pages 378–387. MIT Press, 1997.

[21] M. S. Fontán and M. J Matarić. Territorial multi-robot task division. *IEEE Transactions on Robotics and Automation*, 14(5):815–822, 1998.

[22] Erann Gat. *Three-layer architectures*, pages 195–210. MIT Press, Cambridge, MA, USA, 1998.

[23] Brian P. Gerkey, Richard T. Vaughan, Kasper Stoy, Andrew Howard, Gaurav S. Sukhatme, and Maja J. Matarić. Most valuable player: A robot device server for distributed control. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1226–1231, 2001. (Also appears in Proceedings of the Second International Workshop on Infrastructure for Agents, MAS, and Scalable MAS at Autonomous Agents 2001, Montreal, Canada, May 29, 2001).

[24] J. Grinnell, C. Packer, and AE Pusey. Cooperation in male lions: kinship, reciprocity or mutualism? *Animal Behavior*, 49:95–105, 1995.

[25] O.E. Holland and C. Melhuish. Stigmergy, self-organization, and sorting in collective robotics. *Artificial Life*, pages 173–202, 1999.

[26] SRI International, March,12 2005. `http://www.sri.com/about/timeline/shakey.html`.

[27] iRobot, March,12 2005. `http://www.isr.com/consumer/`.

[28] N. Jakobi. Half-baked, ad-hoc and noisy: Minimal simulations for evolutionary robotics. In P. Husbands and I. Harvey, editors, *Proceedings of the Fourth European Conference on Artificial Life: ECAL97*. MIT Press, 1997.

[29] N. Jakobi, P. Husbands, and I. Harvey. Noise and the reality gap: The use of simulation in evolutionary robotics. *Lecture Notes in Computer Science*, 929, 1995.

[30] Fred Keijzer. *Representation and Behaviour*. The MIT Press, Cambridge, Massachusetts, London, England, 2001.

[31] Kevin Mitchell Kenneth Prestwich. *Game Theory*. Basic Books, College of the Holy Cross, Worcester, MA USA 01610, 1999.

[32] J. R. Krebs and R. Dawkins. Animal signals: mind reading and manipulation. In J. R. Krebs and N. B. Davies, editors, *Behavioural Ecology: an evolutionary approach*, pages 380–402. Blackwell, 1984.

[33] Autonomy Lab, March,12 2005. `http://www.cs.sfu.ca/research/groups/autonomy/`.

[34] J. Lettvin, H. Maturana, W. McCullogh, and W. Pitts. What the frog's eye tells the frog's brain. In *Proceedings of the IRE, ABiALS 2004*, volume 47:11, pages 1940–1951, 1959.

[35] Henrik Hautop Lund, Barbara Webb, and John Hallam. A robot attracted to the cricket species gryllus bimaculatus. In *4th European Conference on Artificial Life*, pages 246–255. MIT Press, 1997.

[36] J.H. Marden and J.K. Waage. Escalated damselfly territorial contests are energetic wars of attrition. *Animal Behavior*, 39:954–959, 1990.

[37] Northeastern University Marine Science Center, April,4 2005. `http://www.onr.navy.mil/media/gallery_category.asp?Cat=general&curpage=3`.

[38] H. Maturana. The organization of the living. *International Journal of Man-Machine Studies*, pages 7/3:313–332, 1975.

[39] H. Maturana and F. J. Varela. *Autopoiesis and cognition: the realisation of the living*. Boston: Reidel, 1980.

[40] John Maynard-Smith. *Evolution and the Theory of Games*. Cambridge University Press, 1982.

[41] D. McFarland and T. Bosser. *Intelligent Behaviour in Animals and Robots.* MIT Press, 1993.

[42] David McFarland. *Animal Behaviour.* Longman Scientific & Technical, 1985.

[43] R. Möller. Perception through anticipation. a behavior-based approach to visual perception. In A. Riegler, M. Peschl, and A. von Stein, editors, *Understanding Representation in the Cognitive Sciences*, pages 169–176. Plenum Academic / Kluwer Publishers., 1999.

[44] N. Nilsson. Shakey the robot. *Technical Note 325, SRI International*, 1984.

[45] Ronald Noe. A model of coalition formation among male baboons with fighting ability as the crucial parameter. *Animal Behavior*, 47:211–213, 1994.

[46] Rolf Pfeifer. Building fungus eaters: Design principles of autonomous agents. In *FROM ANIMALS TO ANIMATS 4: Fourth International Conference on Simulation of Adaptive Behavior, Cambridge, MA*, pages 3–12. MIT Press, 1996.

[47] Autonomy Lab RAGE Project, March,12 2005. `http://www.cs.sfu.ca/research/groups/autonomy/projects/rage/`.

[48] R.E.Fikes and N.J.Nilsson. Strips: A new approach to the application of theorem proving to problem solving. *Artificial Intelligence*, 2:189–208, 1971.

[49] J. G. Riley. Evolutionary equilibrium strategies. *Journal of Theoretical Biology*, 76:109–123, 1979.

[50] Atila robot at MIT, March,12 2005. `http://www.ai.mit.edu/projects/hannibal/hannibal.html`.

[51] Genghis robot at MIT, March,12 2005. `http://www.ai.mit.edu/projects/genghis/genghis.html`.

[52] Matthias Scheutz and Paul Schermerhorn. The more radical, the better: Investigating the utility of aggression in the competition among different agent kinds. In *Proc. of SAB 2004, Los Angeles, CA.* MIT Press, 2004.

[53] V. Srinivasan, P. Nuggehalli, C. Chiasserini, and R. Rao. Cooperation in wireless ad hoc networks, 2003.

[54] F. G. Varela. *Principles of biological autonomy.* New York: Elsevier, 1979.

[55] Richard T. Vaughan. *Experiments in Animal-Interactive Robotics.* PhD thesis, University of Oxford, 1999.

[56] Richard T. Vaughan, Brian P. Gerkey, and Andrew Howard. On device abstractions for portable, reusable robot code. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2121–2427, Las Vegas, Nevada, U.S.A, Oct 2003. (Also Technical Report CRES-03-009).

[57] Richard T. Vaughan, Kasper Stoy, Gaurav S. Sukhatme, and Maja J Mataric. Go ahead, make my day: robot conflict resolution by aggressive competition. In *Proc. Int. Conf. Simulation of Adaptive Behaviour, Paris, France*. MIT Press, 2000.

[58] B. Webb. What does robotics offer animal behaviour? *Animal Behaviour*, 60:545–558, 2000.

[59] R.C. Ydenberg and L.M. Dill. The economics of fleeing from predators. *Advances in the Study of Behaviour*, 16:229–249, 1986.

[60] Yinan Zhang, Mauricio Zuluaga, Sarah Brown, and Richard T. Vaughan. Aggression proportional to task investment can improve group efficiency. *Last-Minute-Results Poster at the International Conference on Simulation of Adaptive Behaviour (SAB'04)*, 2004.

[61] Mauricio Zuluaga and Richard T. Vaughan. Reducing spatial interference in robot teams by local-investment aggression. *IEEE/RSJ International Conference on Intelligent Robots and Systems IROS2005*, 2005.