The practical art of moving physical objects

Citation for published version (APA):

Penders, J. S. J. H. (1999). *The practical art of moving physical objects*. [Doctoral Thesis, Maastricht University]. Universiteit Maastricht. <https://doi.org/10.26481/dis.19990611jp>

Document status and date:

Published: 01/01/1999

DOI:

[10.26481/dis.19990611jp](https://doi.org/10.26481/dis.19990611jp)

Document Version:

Publisher's PDF, also known as Version of record

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.umlib.nl/taverne-license

Take down policy

If you believe that this document breaches copyright please contact us at:

repository@maastrichtuniversity.nl

providing details and we will investigate your claim.

THE PRACTICAL ART OF MOVING PHYSICAL OBJECTS

by Verkeering van de ...

... ..

...

...

...

...

...

...

MOVING PHYSICAL OBJECTS THE PRACTICAL ART OF

ISBN 90-72125-64-9

Dissertation Series No. 99-4

© 1999 J.S.J.H. Penders

Cover design: Alexander Brodsky

Copyright reserved. Subject to the exceptions provided for by law, no part of this publication may be reproduced and/or published in print, by photocopying, on microfilm or in any other way without the written consent of the copyright owner; the same applies to whole or partial adaptations. The copyright owner retains the sole right to collect from third parties fees payable in respect of copying and/or take legal or other action from this purpose.

THE PRACTICAL ART OF MOVING PHYSICAL OBJECTS

PROEFSCHRIFT

ter verkrijging van de graad van doctor aan
de Universiteit Maastricht,

op gezag van de Rector Magnificus, Prof.dr. A.C. Nieuwenhuijzen Kruseman
volgens het besluit van het College van Decanen,
in het openbaar te verdedigen
op vrijdag 11 juni 1999 om 14.00 uur

door

Jacobus Sybrandus Johannes Hendrikus Penders

Promotor: Prof.dr. H.J. van den Herik

Co-promotor: Dr. P.J. Braspenning

Leden van de beoordelingscommissie:

Prof.dr. P.T.W. Hudson (voorzitter)

Prof.dr. A. de Bruin (Erasmus Universiteit Rotterdam)

Prof.dr. E. Flerackers (Limburgs Universitair Centrum)

Prof.dr.ir. A.J.M. Vermunt (Katholieke Universiteit Brabant)

Prof.dr. H. Visser

Voor mijn Ouders,

Vader, helaas heeft de afronding van dit werkstuk
voor jou te lang geduurd.

Preface

The supply problem is as old as human society, and thus the art of moving physical objects is just as ancient. In this thesis an attempt is made to uncover some of its inherent mechanisms. However, as is to be expected, this attempt is limited and restricted to a partial study of this inclusive art.

The roots of this thesis can be traced back to the time I was a student at the University of Amsterdam. I combined my job, in the robotics group at KPN Research, with finishing my master's degree in mathematical logic. In particular it was my teacher Frank Veltman who encouraged me to study a team of robots from a logical point of view. Together with Bart van der Mark computer simulations were made to evaluate the behaviour of our robot teams. Later on I also became involved in studies of logistics. Robert van der Wel and myself contacted IML (Dortmund) and Eltag Bailey (Genoa), and the IDEAS consortium was formed. Of course many more people were involved in the IDEAS project, but I would like to mention in particular Maurizio DelCanto and Torsten Schreiber. Both of them had an open mind for my ideas.

Having these two streams of research available, my coach Peter Braspenning encouraged me to bring them together in a thesis. Peter always grasped some consistency, even though the pieces of text which I showed him were little more than collections of notes. Finally, my supervisor Jaap van den Herik advised me to arrange my texts into a composition of problems and answers, spread over sections with clear cut headings.

So much for the main process. I had inspiring discussions with many friends and colleagues on each of the subjects. It is not possible to mention all of them here so I restrict myself to those who were at certain moments decisive. I remember my discussions on logics with Frank Veltman and Johan van Benthem and it was Jos Vermunt who pushed me forward in logistics. Lyuba Alboul pointed out several pitfalls in geometry. While Marga Jager introduced and guided me through Aristotle's theory of perception. Besides the scientific discussions my training in the

martial arts also contributed to the development of my point of view and for this I am grateful to my teacher Rene van Amersfoort and the other students.

The preparation of the manuscript of this thesis took quite some time. I received support from my company KPN Research and in particular from Jan Hein Donker to whom I owe the layout and who managed to maintain consistency between the many versions. Also Marjan Bolle kept an eye on the progress and took over the work when necessary. Martien van Gaalen and Rudi Hietbrink prepared the drawings. With regard to the English language I thank Ray Fairclough for reading and correcting nearly the whole manuscript. Furthermore, Carol Waddington and Aisling Murnane each advised on several chapters. Since then I have made changes to the text and any mistakes in the present version are my own. Last but not least, I wish to thank Lyuba for creating the right ambience for me to be able to work easily and naturally.

Contents

1	Introduction	1
1.1	Agents, Activities and Actors	1
1.1.1	Agents and Actors	2
1.1.2	Activities	3
1.2	Supply Chains in the Past	3
1.2.1	The Silk Route	3
1.2.2	Means and Infrastructure	4
1.2.3	Organisation and Actors	6
1.2.4	Trade and Communication	6
1.3	The Essentials of a Supply Chain	8
1.3.1	Many Pieces, Many Actors	8
1.3.2	The Main Constituents	8
1.4	Two Problem Statements	10
1.5	The Global View	11
1.5.1	Division of Labour	11
1.5.2	The Study of Logistics	13
1.5.3	Unravelling a Chain	14
1.6	The Local View	16
1.6.1	Autonomous Performance	17
1.6.2	Task Specification	18
1.6.3	Know-How	19
1.6.4	Conflict Resolution	20
1.7	Lines Through the Thesis	21
2	Organisations and Supply Chains	23
2.1	The Historical Perspective	23
2.2	Organisational Models	24
2.2.1	The Silken Thread	24
2.2.2	Authority-Based Organisation	25
2.2.3	Contract Net	27

2.3	From the History of Warfare	29
2.3.1	Food Supplies	30
2.3.2	Autonomy in the Army	32
2.3.3	The Concert of Actions	33
2.3.4	The Military Supply Chain	34
2.4	From the History of Mail Services	36
2.4.1	The Geographical Orientation	37
2.4.2	Historical Notes	39
2.4.3	Bilateral Mail Connections	40
2.4.4	A Royal Service	43
2.4.5	A Contract Network	45
2.4.6	Towards Modern Times	47
2.5	Conclusions from the Examples	49
2.5.1	Silken Thread	49
2.5.2	Authority-Based Organisation	50
2.5.3	Contract Net	51
2.5.4	Models and Present Supply Chains	51
3	The Fabric of a Supply Chain	53
3.1	Organising Supplies Activities	53
3.2	A Structure of Operations	54
3.2.1	The Framework of Conditions	54
3.2.2	Relation to the Organisational Models	57
3.2.3	Resource Management	58
3.3	Towards a Formalism	61
3.3.1	Passive Elements	61
3.3.2	A Formal Language	62
3.3.3	Decomposing and Combining Operations	64
3.3.4	Composing a Chain	68
3.4	A pan-European parcel distribution	69
3.4.1	The practical problem	70
3.4.2	Available Operations	71
3.4.3	Rearranging the Chain	74
3.4.4	Redesign of Services	76
3.4.5	A Hub in the Service	80
3.5	To Conclude with the Supply Chain	81
3.5.1	Other approaches	81
3.5.2	Summary on the Structure of Operations	83
3.5.3	Problem Statement 1	84

4	Robots and Motor-Schemas	85
4.1	Limitations in Robot Design	85
4.2	Agents on a Cargo Terminal	86
4.2.1	Sorting	87
4.2.2	Robot-Agents	87
4.3	Sensory-Based Robotics	87
4.3.1	Motor-Schemas	88
4.3.2	Control Architecture	89
4.3.3	Implementation of the Robot	93
4.3.4	Sensing and the Robot Design	97
4.4	Sensor Basis	99
4.4.1	Definition of the Basis	100
4.4.2	Cognitive Basis	102
4.5	Constructing Sensor Views	105
4.5.1	Surveying	105
4.5.2	Sensors divided over Motor-Schemas	106
4.5.3	Sensor Data Fusion	107
4.6	The Robot Designer	111
4.6.1	The Frame of Reference	111
4.6.2	Design Directions	112
4.7	Conclusions from Robot Design	113
4.7.1	Problem Statement 2	113
5	Spatial Interaction of Robots	115
5.1	Spatial Features	115
5.1.1	Ontology	116
5.1.2	No Communication	116
5.1.3	Autonomous Agents	117
5.1.4	Intentions of Agents	117
5.2	Team Behaviour	118
5.2.1	Distributed Control and Autonomy	118
5.2.2	Different Intentions	119
5.2.3	Congenial Teams	120
5.2.4	Newcomer in a Team	122
5.2.5	Other Team Designs	124
5.2.6	Reviewing Team Behaviour	127
5.3	Interaction: Collisions and Conflicts	127
5.3.1	Spatial and Temporal Preliminaries	128
5.3.2	Interaction Patterns	130
5.3.3	Reviewing Collisions and Conflicts	135
5.3.4	Intentions Defined	135
5.4	Obstacle Avoidance	137

5.4.1	Flee-Opposite Avoidance	137
5.4.2	Potential-Field Avoidance	140
5.4.3	Manoeuvring-Board Avoidance	141
5.4.4	Reviewing Obstacle Avoidance	143
5.5	Flee-Opposite Interaction	144
5.6	Potential-Field Interaction	145
5.6.1	Basics	146
5.6.2	A Robot Team	148
5.6.3	Balanced Conflicts	150
5.6.4	Nearly-Balanced Conflicts	152
5.6.5	Far-Away Goals	153
5.6.6	Reviewing the Potential-Field Procedure	155
5.7	Conclusions on Interaction	156
5.7.1	Problem Statement 2	157
6	Motor-Schemas and Cognition	159
6.1	Intentions and Action	159
6.1.1	Intentions of Actions	161
6.1.2	Preview of Sections	162
6.2	Philosophical Notions	163
6.2.1	Objects	163
6.2.2	Consciousness	164
6.2.3	Abstract Reasoning	165
6.3	Robots and Cognition	167
6.3.1	Sensor Views and Motor-Schemas	168
6.3.2	Cognition and Internal States	169
6.3.3	Actions and Intentions	169
6.4	Actions, Intentions, and Imagery	174
6.4.1	Aristotle: Imagination, Intellect and Desire	174
6.4.2	Motor-Image	176
6.4.3	Imagery	177
6.5	Revisiting Cognitive Notions	180
6.5.1	Cognitive Basis	180
6.5.2	Awareness and Actions	182
6.6	Acting in a Physical World	186
6.6.1	Beyond Solipsism	186
6.6.2	Communication and Niche	187
6.7	Conclusions on Cognition	188
6.7.1	The Problem Statements	189

7	Conclusions	191
	Appendix: Potential-Field Procedure	195
	Summary	201
	Samenvatting	207
	References	213
	Other references	222
	Curriculum Vitae	223

1 Introduction

This study examines *supply chains* in society, we have called it a study in the practical art of moving physical objects. In this chapter we present the contours of this art and outline our research. We discuss some historical phenomena and distinguish two views on the subject: the global and the local view. They are captured into two problem statements, which determine the lines to be followed in the subsequent chapters.

1.1 Agents, Activities and Actors

The human race lives in scattered communities all over the planet. The fact that this is well-known by the members of those communities, shows that human communication links connect the communities. However, besides communication links, there also exist extended networks of transport links. This makes it possible for people living in arctic regions to eat oranges and tropical fruits which have been grown in different climate zones.

Not only humans, but also various other biological species live in communities. Ants and termites, for instance, live in communities of a million members. These communities are a dazzling concentration of activities of ant *agents*. At first sight, not much of an ordering can be recognised. However, ants are able to build impressive and ingenious home bases. To do so requires a high degree of specialisation amongst the workers in the community. Workers inside the home base must be provided with supplies by other workers, not only with food but also with construction materials. Hence, there must be some ordering of their activities: i.e., they form supply chains.

Our study is restricted to human society. In human supply chains many agents are also active, and likewise must have order. Regarding these chains, the achievements of modern technology – fast communication and the huge and powerful means of transport – attract much attention. But, as an analogy to the ant societies,

our research faces the question: *how can the activities of so many agents come to form a supply chain?*

1.1.1 Agents and Actors

A supply chain is a concatenation of actions of performing agents. The agents may be either human or artificial. Particular instances of artificial agents are robots. Agents performing within a supply chain are called *actors*. The terms agent and actor will be used as follows: an agent is a natural or artificial entity that is able to perform, an actor is an entity participating in some overall activity. Actors, for instance, perform a task in a supply chain. Agents can become actors, but not all actors are agents; companies are actors but not agents.

When a human being and his¹ performance are discussed, the emphasis is on his role as actor: on *why* he is required to perform a certain task. The assumption is that a human knows how (or easily learns how) to execute a task: he is an agent and it is described how to become an actor. When artificial agents are involved, most attention goes to *how* they perform and how to make (sophisticated) agents out of them.

The difference in approach is understandable. In science fiction, robots might be depicted as a nearly mature species of self-sufficient creatures. However, such visions are far from the modest aims of robot designers. Robots cannot adapt their behaviour as easily as humans can. What is more, they can only with difficulty perform an action which humans consider simple; for instance taking a match out of the matchbox and strike it. For a robot this task poses many problems, e.g., finding the orientation of the matchbox, finding the right end of the match, finding the striking-edge of the matchbox. A robot can repeat a task with an amazing regularity, for instance, turn a thousand screws a day with a constant speed without showing any sign of fatigue. Each type of robot has a limited domain of application. This domain is called the *niche* for which the robot is fit. So, robots have their own particular niche, which differs substantially from that of a human actor.

In a supply chain actions need to connect to one another. What is essential is that the supplies flow through the chain. The agents can perform in their particular niches. Most niches involve human agents with specialised skills, others might be filled by robots or artificial agents. Arranging a supply chain consists in directing the supply flow and connecting the niches. In this thesis, we compare niches and study how agents can be applied and become actors.

¹ In cases where a choice in gender is possible, we use the masculine, as far as the author is concerned this is only for reasons of readability and does not express any position in a scientific gender discussion.

1.1.2 Activities

We study the linking of activities and the inherent structure of a supply chain. The history of human society has many examples of humans organising their labour. However, the basic structures of these organisations are seldom explained. Presumably this is due to the fact that humans learn easily and understand almost instantaneously how to participate as an actor. In any case, whenever tasks are divided among agents, each agent must know how to perform physically and must understand how to perform as an actor. But what does it exactly mean to perform? Studies in robotics have shed some light on what, in this respect, is presumed when a task is delegated.

Obviously, robotics focuses on the physical performance of a robot. However, in the field of robotics and artificial intelligence there is an increasing tendency to have teams of (autonomous) agents co-operate. The effort of a team is successful if the team can accomplish a task as part of a larger chain. Hence, the study of supply chains provides a context to evaluate team work. In a supply chain, there must exist some basic ordering principle that forces the flow in the appropriate direction. As soon as we know and understand this principle, the appropriate position of the niches of our agents also become clear.

The task of a robot needs to be specified unequivocally. Indeed, it is a major problem to make a robot adapt its actions to (slightly) changed circumstances. Contrary to robots, human agents are flexible. They have imagination and fantasy. However, if they exhibit too much fantasy when executing an action, the chain might not proceed in the intended direction. In brief, it is not at all easy to define what it means to perform physically and to perform as an actor.

1.2 Supply Chains in the Past

The phenomenon of a supply chain is probably as old as human society. Supply chains have emerged from the social process of division and specialisation of labour; when consumer and producer become separated, a chain of processes to move goods from the one to the other is needed. We discuss a few examples from ancient history to outline the basic characteristics of a supply chain.

1.2.1 The Silk Route

For several thousand years there existed a trade route for silk from China to the Middle East and Europe. This route had a length of nearly 6,000 km. It passed through several large and dangerous deserts. The most dangerous one was the Taklamakan (in translation: the land from which nobody returns). The Taklamakan had a length of 1,000 km, and was notorious for its "moving hills". At the west

end of this desert there were the Pamirs, i.e., mountains with an altitude of 7,000 metres. The Persians called them "*the roof of the world*" [Franck and Brownstone, 1986].

The route is called the *Silk Route*. Silk was exported from China to the Roman Empire. History shows peaks and sharp falls of the trade which often coincided with the rise and fall of empires. The period from 50 B.C. to A.D. 200 was prosperous for the silk trade. Roughly, the route passed through four major territories. On the eastern end the Chinese empire under the Han-dynasty (206 B.C.–A.D. 231), and on the western end the Roman Empire. Adjacent to the Roman Empire, was the Parthian Empire (covering Persia, 247 B.C.–A.D. 224). The territory, connecting Parthia with China was from 100 B.C. ruled by the people which the Chinese called Yüeh-Chih. The major town was Bactria (nowadays called Balkh). The Yüeh-Chih first formed five princedoms. From A.D. 50 till the second half of the 3rd century it was united in the Kushan Empire [Nikonoriv, 1997]. The Kushan Empire not only controlled a considerable part of the Silk Route, but also of the so-called Indian Grand Road. The latter route, headed by Bactria, ran southwards to India.

The most striking fact about the silk trade is that there were hardly any direct contacts between the Romans and the Chinese. A visit of a Chinese delegation to Bactria is recorded as early as 138 B.C. [Franck and Brownstone, 1986]. Another Chinese embassy visited the Parthians in 97 A.D. Of course, as neighbours, Romans and Parthians had much intercourse, sometimes friendly, sometimes violent. However, the only recorded direct contact between Romans and Chinese was in 166 A.D. It is worthwhile to note that for centuries the Parthians succeeded in deceiving the Chinese about the Roman Empire and in frustrating Roman searches for intelligence. Thus they prevented any direct overland contacts. The control over the area and the trade brought the Parthian and the Kushan Empires great wealth.

So far this brief discussion of the Silk Route in ancient history. For trade purposes, almost one thousand years later Marco Polo (1271–1295) and the Polo family travelled along approximately the same route to China. However, rather soon after Marco Polo's time the trade routes overland were replaced by sea routes. Because of the decline of the trade on some (and later all) routes, many of the cities along the route disappeared; in the Taklamakan desert they were literally covered by sand.

1.2.2 Means and Infrastructure

The Silk Route is an example of a supply chain. It is obvious that a supply chain requires actors and agents to perform, but its existence also relies on a large set of preconditions or conditions *sine qua non*. Law and order are required, and means for transport and travel. Means include not only a system of roads and tracks but

also food supplies, shelter, and people (workers) maintaining the roads and providing services.

The trade on the Silk Route proceeded over several alternative routes, each of which consisted of a long chain of cities and oases. Without the water and food supplies provided by the cities and oases, no transport was possible, certainly not in the deserts. The local societies provided the means and the manpower to maintain the infrastructure. We remark that in private trade it is often not recognised that the trade is fully dependent on enabling conditions and accompanying efforts provided by local societies. The history of the silk trade is illustrative: the local societies shared in the advantages of the trade, and some of the cities grew to prosperity. However, at the decline of the trade many cities disappeared.

The requirements of a supply chain are well recognised in the military field. Consequently, several of the terms used stem from this field. An important term is *infrastructure* which is according to Webster's dictionary [1990]: "*the whole system of bases, services, training establishments etc. required for the use of troops in military operations // the basic framework of any organisation.*" The infrastructure includes many of the preconditions mentioned above.

Besides the infrastructure, the military field much appreciated the value of local means: the army needs supplies, preferably obtained from local sources. In times of war, often new supply chains have to be established and maintained by the armies themselves. The enormous effort demanded is quite strikingly stated by the Chinese general Sun Tzu. He noted in 490 B.C. when advocating that an army should forage from the enemy: "*One cartload of the enemy's provisions is equivalent to twenty of one's own.*" [Sun Tzu, B.C. 490, p 21].

In the time of Sun Tzu the supplies of the army consisted largely of live cattle. It is not difficult to imagine the costs of maintaining (feeding) the cattle until the ultimate moment of consumption. The solution was to keep the army moving and loot the areas it was passing through. The supply problem became manifest in a siege. Sun Tzu, well aware of the problem, says: "*the worst policy of all is to besiege walled cities*" [Sun Tzu, B.C. 490, p. 23].

It is not surprising, that the military were the first to regard the supply flow and chain as an object for study. Traces are found in ancient history, in Chinese history as is clear from the work of Sun Tzu cited above, but also in ancient Rome [Piedie, 1994]. The modern science is headed *logistics*. Its founder Jomini [1855] defined logistics as: "*the practical art of moving armies*". Our subject is the moving of objects and is not particularly focusing on armies. Thus, as a variation on Jomini's definition above, we have called our subject the (science or) *art of moving physical objects*.

1.2.3 Organisation and Actors

We aim to shed light on the phenomenon of a supply chain. We discussed the silk trade as an example, but what was the essence of what we nowadays call the Silk Route? Though this question has a much broader scope than our research (the Silk Route, was primarily an exchange route between cultures; trade was additional and only a result of the exchange²) we restrict the discussion to the field of logistics, i.e., logistics in its original meaning of moving supplies in supply chains.

To allow the supplies to flow in a supply chain, an ordering of activities is required. In an army, all actors involved are under a single line of command. It is often assumed that this provides optimal conditions for a good organisation. Plans can be worked out by a staff and are executed by issuing commands to subordinates. Analogously, in civil applications there is a strong tendency to think of a required organisational structure in terms of ownership and the corresponding authority and delegation of authority. For example, multinational companies consist of holdings, divisions, country organisations, departments and so on. However, neither a line of command nor an ownership structure says anything about how the actual supply chain (which remains the ultimate aim) is arranged. It often turns out that the supply chain shows deficiencies, although the command structure of the army or the company is well-organised.

A strict organisation of command and control in the chain does not guarantee a smooth flow of supplies. Moreover, supply chains can work without any command structure. The silk trade was based on the activities of individual merchants who were going from one city to another with their merchandise. There was no actor who designed an overall plan of how to collect silk in China, and to transport it to the West and distribute it there. As discussed above, for centuries there was no direct contact between the Roman and Chinese Empires, let alone that business plans covering the whole trade route were made. The same is true for other products. For instance, tea, coffee, and spices were sold and distributed over the Eurasian continent for centuries (on the Spice Route), without being monopolised by any single (commercial) organisation. Hence, a supply chain can exist quite independently of whatever kind of (ownership) relationships exist amongst the actors.

1.2.4 Trade and Communication

Supply chains are often associated with *trade*. Communication is important as well, but communication and communication links are not a prerequisite for a

² Buddhism was imported to China along this route. Also Syrian jugglers travelled from city to city along the route. Some ended in China, where they were much admired by the Chinese audience.

supply chain. Neither a supply chain nor trade depends on communication. Trade, and subsequently a regular flow of supplies, might precede regular communication.

It is supposed that the main routes over the Eurasian continent, the Silk Route(s) among them, were first carved by nomadic peoples. In prehistoric times, the continent gradually became divided into several settled areas in well-watered river valleys and oases; the areas were divided by steppes. In the steppe, nomadic peoples wandered with their herds. The nomads, however, followed regular patterns and migration routes, which could span 600–750 km [Encyclopaedia Britannica]. Goods loaded on their herd animals were exchanged, for instance, for pieces of craftsmanship [Franck and Brownstone, 1986]. In fact, the nomadic peoples established long-distance supply chains. Later, nomads often served as caravan personnel.

Communication and communication links do not always precede trade. In the so-called *silent trade* [Franck and Brownstone, 1986] only a minimum of communication is required. Silent trade works as follows. At some place (agreed by tradition) one trading party lays out goods and leaves the place. The other party lays out a pile of goods, which it thinks to be equivalent, and leaves the spot. The first party returns and adds or subtracts from its own pile and so on. This kind of trade presupposes only a mutually agreed protocol and requires no direct communication. Similar patterns of exchange on the Silk Route, are reported by Chinese and Roman resources [Franck and Brownstone, 1986].

An even more direct is the so-called *pirate trade* [Franck and Brownstone, 1986]. A tribe breaks into the camp of another tribe, takes away the objects of its interest and leaves behind objects which (from its point of view) represent an equal value. In this case the protocol is forced by one party upon the other. And no communication is present at all.

Trade can emerge without communication. And, if in the historical examples there was communication, it was not much faster than the forwarding of goods themselves. The merchant or his messenger, would have to accompany the goods on their way. In the silk trade this was obvious, but also in the Middle Ages of Western Europe this was usually the case, for instance, in the Hanseatic League [cf. Krans, 1963]. Dedicated communication links existed in the form of mail services. Letter mail was for a long time the fastest communication channel. However, a regular mail service is not something which can be set up overnight. It is quite an effort to set up the required organisation. It will only emerge when already a sufficient need exists for a permanent exchange of information. And this means that there must be a sufficient (public) interest in communicating, for instance, for political purposes or because of emerging trade. Communication links emerge from trade rather than vice versa.

1.3 The Essentials of a Supply Chain

Below, we aim to uncover the constituents of the *art of moving physical objects*. Its essence is *moving* the objects. However, keeping the objects moving is not simple. A supply chain is a subtle linking of relatively small physical activities. If only one link fails, there is no supply. The art also consists of properly arranging the process pieces.

1.3.1 Many Pieces, Many Actors

Supply chains can be quite large, and consist of large structures distributed over many actors. Moreover, large numbers of items flow through them. Some old-established chains range world-wide. Examples are the supply chains of silk, tea, coffee, spices, and diamonds but also the chains of the mail services. Though it is not usual terminology, we also speak of the mail supply chains. The complete supply chain from a producer to an end consumer is often neither known nor traceable. Consequently, discussions must be restricted to a certain part of the supply chain. In these cases we should speak of a partial supply chain, but the addition 'partial' will be left out, unless confusion may arise.

Recently established and large supply chains are found with single branded articles which can be bought and sold all over the world. Examples are world-wide brands of lemonade, sportswear, televisions, and personal computers. The brand-name articles are produced in a factory and are to be distributed all over the world. They require an extended distribution chain which starts at the factory and ends at the consumer. Organising a large supply chain goes beyond the borders of a single company. The trades in silk, tea, coffee and diamonds serve as an illustration. In these trades, many agents and actors are involved, each performing some activity or service, independently. In this respect a supply chain for a modern single branded product is not different. And just as in the image of an ant colony we come to our central question: *how can the activities of so many agents come to form a supply chain?*

1.3.2 The Main Constituents

A supply chain is embedded in its environment. In studying supply processes, Vermunt [1996] distinguishes three systems in the environment: the *physical system*, the *organisational system* and the *information system*. The physical system comprises the means and physical infrastructure. Agents can be regarded as means and belong to the physical system. The organisational system includes the framework of social structures that determine how and which parts of the physical system are divided among the members of the society. It encompasses the social and organisational part of the infrastructure. Some terms characterising the social system are: ownership, permissions and contracts for access, application, and use. Social institutions, such as public authorities, agencies, companies etc. are actors

and belong to the organisational system. Trade is an organisational phenomenon too. The information system basically consists of communication channels of whatever kind. It provides actual information on conditions in the supply chain and the feedback from the process. (The information is among others, often used as a basis for charging fees.)

Besides the three systems, Vermunt [1996] distinguishes three supply (chain) activities: *transport*, *storage*, and *transformation*. The meaning of transport and storage is assumed to be clear. Transformation means that goods are processed and have their external or internal characteristics changed. In any case, the activities change the four goods-related parameters of: time, place, quality and quantity. Roughly speaking, transport and storage change the place and time co-ordinates, transformation changes quality and quantity.

In a supply chain goods move, the chain does not change the quality of the goods. In the idealised case, quantities remain unchanged, in the sense that the total quantity of goods that goes into the supply chain goes out as well. We mainly deal with transport and distribution related storage. Transformations are considered only as far as they are necessary for transport. The agents in a supply chain only re-order goods in time and space, their actions concern transport and storage.

Figure 1.1: The basis of a supply chain.

A supply chain links existing (local) physical and organisational systems. We emphasise that in general, these systems exist for their own sake; they are certainly not dedicated to a particular supply chain. In the case of the Silk Route, there was a close relationship between the cities and oases in the deserts and the trade on the route. But this does not mean that the cities solely existed for the sake of trade. In our interpretation of Vermunt [1996], a supply chain links existing systems and supplies flow through the interface of these systems. Figure 1.1 suggests the sup-

ply chain flowing through this interface. On the basis of the examples given above we notice that the information system is not always so important as Vermunt [1996] and also Figure 1.1 might suggest.

1.4 Two Problem Statements

The subject of our study is the practical art of moving physical objects. Moving an object directly is not much of a problem; a move as part of a bigger process is quite different. We referred earlier to the metaphor of an ant colony and its supply machinery, and briefly discussed the historical Silk Route, which was a large chain spanning the whole Eurasian Continent.

A supply chain is always spread out over space and time and has to bridge occurring gaps. Consequently supply chains are often distributed over many agents and actors. In the examples given, no single actor is able to carry out the whole supply process on his own. Simply, no actor exists that both covers the world and penetrates sufficiently deep in local societies. A division of labour is inevitable. Moreover, in general, the actors in a supply chain are relatively independent and may have quite different backgrounds. With these characteristics, a supply chain distinguishes from any other industrial process and we can indeed speak of a *practical art of moving objects*.

In a supply chain, goods are handed over from one actor to another actor. We have argued above that communication plays a subordinate role. In the old days, communication means were limited. Nowadays, modern techniques for telecommunication and computer applications allow fast and extensive communication over long distances. The time needed for forwarding information is negligible compared to the speed of forwarding physical objects or goods. Extensive communication might precede the real exchange of goods, without delaying the flow of the goods. Thus, it seems that the art of moving objects deserves a revision and a re-evaluation in the light of modern techniques. Moreover, the question arises how to use and take advantage of advanced information exchange. We examine the inherent structure and coherence of a supply chain, to provide a combined solution for two major problems formulated below in two problem statements.

The first and most prominent problem is how to make the goods flow through a supply chain:

Problem statement 1. *How can the activities of the actors in a (large) supply chain be arranged such that a (goal)directed flow of objects results?*

The focus in the second problem is on the agents needed as resources to perform in the supply chain. We formulate this as follows:

Problem statement 2. *How can a set of resources be made to perform as an actor within a supply chain?*

The first problem deals with linking activities into a chain, and hence we are considering the chain from a *global* point of view. In problem statement 2 we look in detail at the individual actor. The view point of the individual actor is called the *local* point of view. Since a supply chain requires means provided by resources, actors are in charge of the resources. The division of labour must be such that the required resources can be applied and actors are indeed able and willing to let them perform.

Designing a supply chain requires a strategy which combines a solution to both problems. However, combining them provokes some tension. The requirements of the global view interferes with the autonomy of the individual actors (the local view). The discrepancy between the global and local view is explicitly studied in the field of robotics, for instance by Noreils [1992]. In Chapter 4 we discuss the design of autonomous robots. The basic points to be addressed in robot design are interesting for their own sake, but they also shed light on the interplay between arranging operations and delegating tasks to the resources.

In the remaining sections of this chapter we discuss our basic choices on how to approach the ordering of a supply chain, and how to fit in the individual actors. We further specify the problem statements and clarify the basic points for discussion in the following chapters. By doing so we also provide a coherent overview of the topics discussed in the next chapters.

1.5 The Global View

From the global point of view we look at the structure of a supply chain. Below, we analyse how the structure is established, how actors fit into it and make the goods flow.

1.5.1 Division of Labour

The term *actors* is used to denote the whole range from companies, sub-organisations throughout individual workers. The kind of (logistic) actors we are arriving at are: airline companies, railways, express delivery services, postal services, and what-so-ever large or small companies involved in processes of forwarding goods. Furthermore we mention actors such as the postman or the newspaper boy, who finally hand over the item to the consumer at the end of the chain. They all have in common that each actor has a share in realising the same process, that is, the supply of certain goods.

Dimensions

Since, in a supply chain the work is distributed over many actors, there exists a *division of labour*. The labour is divided among actors, but the obtained results have to be linked. We distinguish three dimensions in the division of labour and use them to analyse the linking in a supply chain. The dimensions are:

- 1 task decomposition;
- 2 coherence;
- 3 resource assignment.

(ad 1) The task decomposition is a splitting up of the overall task such that actors are able to execute the sub-tasks. This presupposes some familiarity with the actors. The task decomposition anticipates the competencies of the available actors. If the task is new and has never occurred before, there will be no actor who is experienced. This dilemma is a basic problem when trying to set up any supply chain.

(ad 2) After the sub-tasks are allocated to the actors, some mechanism is needed to maintain coherence in the execution. *Coherence* means that the results of all sub-tasks link together to solve the original request. However, the goal of the supply chain might not be fixed in advance or may even be unknown, as we have seen in the silk trade. We say that a process has *basic coherence* whenever any output is obtained. Ideally speaking, a supply chain does not change quantities of goods (what is fed into it goes out as well). A supply chain is a chain if and only if it has basic coherence. No one would have spoken of a Silk Route if silk had not arrived somewhere. However, at that time nobody had an overview of the route. No one knew both: where the silk did go into the chain and where it did come out.

When the goal of a chain is fixed in advance, it is possible to judge whether the chain is coherent with respect to this goal: if the goal is not achieved the process is not coherent (in this respect). Given a pre-set goal, the notion of coherence can be further developed. Bond and Gasser [1988], for instance, define coherence as a measure for: "*how well the system behaves as a unit along some dimension of evaluation.*"

(ad 3) *Resource assignment* deals with the means to execute the process. The term resource includes the human agents (workers) who have to do the job. Often a resource is defined as "*an object used during execution of actions*" [Pelavin and Allen, 1986]. In such contexts one often speaks of resource allocation. However, as Pollack *et al.* [1987] pointed out, human agents and robots are resource bounded, unable to perform arbitrarily large tasks. Thus one also has to take account of a certain scarcity here. In the context of a division of labour, we speak of assigning a resource to a task, i.e., a resource is put at the disposal of an actor.

Resources, and in particular agents, have to perform the tasks. The role of an actor consists of incorporating the resources into the process. Somebody or some instance has to select and combine the appropriate resources, and assign (sub-)tasks to them. In general, an actor will prepare the resources to perform a task.

1.5.2 The Study of Logistics

We have not entitled our study a study in logistics. The present-day science of logistics emphasises finding optimal combinations of the three systems: the physical, the organisational and the information system, within supply chains. The US council of Logistic Management defines logistics as: *"(...) the process of planning, implementing and controlling the efficient, cost-effective flow and storage of raw material, in-process inventory, finished goods, and related information from point of origin to point of consumption for the purpose of confirming to customer requirements"* cited in [Van Goor *et al.*, 1990]. Note the central role of the cost-effective criterion. Here we find terms such as purchase logistics, which aims at minimising inventory costs within a company. Besides costs, the definition above puts a strong emphasis on control. This is even more explicit in the following definition of logistic management. *"Logistic management stands for complete control of the goods flow and the information flow connected to it"* [Aertsen *et al.*, 1996]. Note that these definitions differ quite strikingly from the one given by Jomini [1855] and cited above in Subsection 1.2.2.

The cost-effective criterion is essential to the subsistence of any supply chain. The overland routes for the silk trade were replaced by sea routes because of the costs. However, it is difficult (if not impossible) to estimate costs in terms of sums of money. To evaluate costs, a basic overview of the supply chain is a prerequisite. Some supply chains do not serve a previously fixed goal. The historic silk trade was a case in which no goal was defined in advance. Thus, directly focusing on costs does not provide any insight into what the process really is, that is studied and optimised. The same is true for the control of a chain.

Because of the division of labour it is essential in supply chains that physical activities are linked appropriately. If not linked appropriately, there is no supply. By putting the emphasis on information and control one can easily ignore the genesis of a supply chain and overlook the mutual dependencies among the physical, organisational, and information systems. At worst, such an approach might culminate in attempts to improve the process by only extending control over it. In Chapter 2, when discussing warfare, we encounter disastrous examples. Further, these examples will show that there exists a basic dependency among the three systems. First, it is evident that without the means and the physical infrastructure there will not be any process at all. Second, the means have to be ordered. We state that an organisational framework in which actors can participate is a prerequisite for any information exchange or control. Attempts to control the supply chain are senseless if the framework does not already exist.

The direction taken in modern logistic science is understandable. The physical system, the means and physical infrastructure do not change so quickly. It has an inherent inertia. The means also do not change much in time. Moreover it is quite expensive to change the infrastructure. The organisational system is based on society, and consists of procedures and habits. At first sight it seems quite easy to change them. But as any practitioner will affirm, its resistance and inertia seem to surpass even that of the physical system.

When the goal of a chain is defined in advance, a command and control structure that spans the chain is possible. In such cases we can indeed turn to questions as to how to optimise and control the chain. Manufacturing industries are good examples. The primary process might, in the extreme case, consist in a connection of huge chemical installations, and here the emphasis on control is quite appropriate.

New Directions

The science of logistics nowadays takes several new directions. Modern society is facing limitations on natural resources and space, and aims at environmental protection. This has stimulated the study of logistics to explore new directions. Quite interesting approaches have been developed under the heading of *city logistics* [Wolters, 1995] and *traffic logistics*, in which the aim is to maintain an optimal traffic flow in the limited space within the city centre. The objective is to optimise traffic flow, but only socially acceptable means and measures can be applied. These studies have to re-evaluate the relationships between the physical and organisational systems. In the case of supply chains, the social restrictions are rarely made explicit, but they are nonetheless present, because otherwise supply chains cannot endure. How else can so many actors effectively co-operate in a single process?

1.5.3 Unravelling a Chain

A supply chain deals with a division of labour, and to answer problem statement 1 we must reveal its structure and coherence. We introduce two important concepts: *operation* and *co-ordination*. An operation denotes an activity, i.e., a task. By co-ordination operations are linked.

Operations

The concept operation stems from the military field. It is described as: "*any movement or series of movements made in carrying out strategic military plans*" [Webster's, 1990]. Henceforth, we consider operations as activities directed to realise the goal for which the supply chain is established.

The operations in a supply chain are actions needed to perform the moving (or handling) of goods. An operation deals with two types of entities:

- *passive elements*: the items that have to be worked on (moved);
- *actors*: those who are in charge to perform.

The passive elements are in particular goods. The flow of the passive elements is the *real process* [Hartman and Roos, 1988] of a supply chain. Actors use resources as the means to perform. Actors prepare the required resources, and let them carry out the operations. Thus, we define operations as:

tasks that are determined by start and end slots, and are executable by the resources (of actors).

The start and end slots mentioned are related to the passive elements. The slots specify the co-ordinates (time and place) of the passive elements. An operation is defined by conditions: some occur at the start and others have to be obtained at the end.

In Chapter 3 we describe a supply chain by a *framework of conditions*. The framework states the co-ordinates of operations, that is, when and where to execute them. The framework of conditions shows the operations linked into a goal-directed structure. This is how we obtain a view on coherence. The framework of conditions describes how the goods flow through the process, like coins falling down inside a coffee machine. The framework of conditions shows the route the coins follow.

Operations are described by input slots and output slots. Nothing is said about how to execute operations nor about how to find the resources. The aspect of resource assignment and resource allocation are covered by what we define in Chapter 3 as *resource management*.

Co-ordination

Supplies flow through a chain only if there is coherence in the chain. Preferably, the execution of each operation makes the goods flow towards the destination. In practice we see that the goods flow in many directions, but seldom straightforwardly to the destination. Coherence has to be achieved via co-ordination actions. Co-ordination is defined as follows:

co-ordination consists of arranging points in time and places for a series of successive actions.

Co-ordination differs distinctly from *co-operation*. First, co-ordination does not require co-operation. Even antagonists can act in co-ordination, as we have seen above in the example of pirate trade. We cannot speak of co-operation in pirate trade, but a basic kind of co-ordination is sought for: the pirates need to go out to meet the other party.

Co-ordination has to do with the *interaction* of actors. Interaction results in group or system behaviour, which goes beyond the behaviour of the individual actor. Co-ordination is a characteristic on the level of group behaviour. A supply chain is typically the final result of collective behaviour, and consists of co-ordinated interactions of many actors. The reasons the individual agents have for acting – to be co-operative or on the contrary, antagonistic – do not count toward the end result. Interaction will be the underlying concept of Chapter 5, where we discuss teams of robots.

Our notion of co-ordination is rather elementary, planning is based on it. In a supply chain planning is used to ensure that, for instance, the carrier does not leave before the goods have arrived. Planning can improve mutual predictability of actors. Actors participating in a clearly ordered supply chain can predict each others' functional role and behaviour. Thus we obtain for instance the broader notion of co-ordination given by Bond and Gasser [1988]: "*Effective co-ordination implies some degree of mutual predictability and lack of conflict.*"

1.6 The Local View

In the previous section we discussed how to arrange operations into a (large) chain. In this global view operations are described by conditions only. In this section we examine problem statement 2: how to make a set of resources perform as an actor; that is the local view. When a task or operation is delegated to an actor it is presumed that the actor knows what he and his resources are supposed to do.

When given a bundle of letters, the postman understands what he is supposed to do, which letters to deliver where. The point is obvious but not self-evident: when replacing the postman with another person, the latter will at first not be able to perform, or performs only poorly.

We use the distinction between the global and local point of view to discuss anew the three dimensions in the division of labour distinguished in Section 1.5. The dimensions have a local counterpart at the point where the actor and his resources are concerned.³

³ On the local level the notions of *actor* and *agent* almost coincide. The actor is the entity that participates in a chain, the agent is the entity able to perform, as defined before. However, in practice, on the lowest level the performing human is at the same time actor as well as agent.

Table 1.1: Views on a division of labour.

Global aims	Local consequences
task decomposition	task specification
coherence	know-how
resource assignment	conflict resolution

A task is part of a structure as it is restricted and delineated. This is laid down in its specification. To achieve coherence, the actors need to know what results are expected. Resources are assigned, but the actor also needs to know how to apply them and he needs to resolve conflicts when resources are used for multiple tasks. Actors are also assumed to be self-employed. Before discussing the local consequences of the division of labour, we first clarify our concept of autonomous performance.

1.6.1 Autonomous Performance

In the field of AI one finds several attempts to define the notion of *intelligent autonomous agent* (see for instance [Braspenning, 1997]). These attempts often aim at distinguishing intelligent autonomous agents from less intelligent agents. In the moving of objects, performing is in itself so complex that our focus is not to delineate the intelligence of the agent, but rather to reduce the required level of detail in specifying a task. In practice, not every detail can be provided from the global point view, and we let the agent do as much as possible. We try to enlarge the scope within which an agent performs autonomously. Thus, we are treating autonomy as relative to the context of application.

We have argued for co-ordination of time and place parameters, and fixed a task between start and end points. Assuming that the circumstances are normal, we can define our concept of autonomous performance:

an agent performs autonomously if it can accomplish a task, determined by the co-ordinates of start and end, without being externally guided.

Within the art of moving objects it is rather natural that actors perform autonomously. Supply chains are based on autonomous performances. Consider for instance a postman. He participates in the (huge) framework of a mail forwarding chain (we examine this in Chapter 2). In the morning he packs his bag with letters and sets off to deliver them. His task is described by place and time slots. Places include the office from where he sets out and the mailboxes addressed on the letters. Time is determined by when he has to arrive at the office, and maybe there is a restriction on when the letters must be delivered. But how the postman actually

executes his task, and, for instance, repairs the flat tyre of his bicycle or car, is left to him. He has to execute his task autonomously. The relative autonomy of the postman complements the global description of an operation. In general, the less the actor is allowed to perform autonomously, the more detailed the description of an operation has to be.

The postman's performance might be curtailed by secondary factors, for instance, that he has to wear a uniform, or he has to go by foot, bicycle, or car. However, for the task itself, or the result of the supply chain of which he is the last link, these matters are not essential. We do not examine these secondary factors any further: they are not really important to performing the supply process.

The interplay between performing autonomously and the supply chain can be further illustrated by another example in the postman's domain given by Rosenschein and Zlotkin [1994]. Two postmen, each having a bundle of letters, start to negotiate an exchange of letters so that each will have a shorter route for delivery. These postmen have autonomy to change and re-arrange the parameters of their respective tasks. Autonomy of the actors in this sense affects the description of the operations.

Incidentally, actual practice (in the Dutch postal service) falls somewhere in between the examples given. The bundles of letters given to the postmen are pre-sorted so that each bundle contains only letters which are to be delivered in one of the predetermined routes. The postman himself has to arrange the letters according to how he actually goes down the route. Very often, he has to combine bundles from different origins.

1.6.2 Task Specification

The delegation of a task requires that the actor understands what the task is. When consensus is achieved, the operations provide the basis for action planning. We argued that operations are complemented by the autonomous performance of the actors. The range of autonomy determines how detailed the begin- and end-points of an operation need to be described. Nevertheless, the task specification has to be clear on what exactly has to be changed, where the task starts, and how to obtain the end situation. Clarity alludes to the point that the instructions must be understandable for the actor. We touch here the rather difficult matter of communication; in particular, the question of whether the communicated message indeed indicates to the receiver at the local level what the sender on the global level intended.

Task delegation is a point which shows a huge contrast between human beings and robots. Human beings with roughly the same background can understand one another rather easily. Task delegation and communication about a task with a robot are (still) problematic. But, we observe that also in cases where the background of

human actors differs too much, the communication is difficult. Since we are only dealing with the moving of physical objects, the problem of task delegation and communication can be simplified to: the message has to trigger the receiver to perform a certain action. The point is to make the receiving actor execute the instruction.

When we have insight into the spectrum of actions which an actor can perform we can predict and specify his behaviour. In Chapter 2 we discuss examples from a human context which show that supply chains presume some consensus between the local actors. In Chapter 3, we define the operations and show how they provide a basis for planning. However, a supply chain might work well with exchange protocols, as in silent trade or pirate trade. Verbal communication is not essential. In Chapter 4 we discuss a robot design, thus setting the stage for Chapter 5, where we study multi-robot teams, and these simple creatures rather easily interact with each other. In Chapter 6 we discuss the prerequisites for interaction and communication in a philosophical perspective.

1.6.3 Know-How

To execute a task, an actor needs to know how to execute it. A technical term for the latter is that of *know-how* [Singh, 1994]. The term is originally from Ryle [1949] who distinguished *knowing-how* from *knowing-that*. *Knowing-that* roughly coincides with knowing facts, e.g., knowing geography. *Know-how* comprises knowledge of how to execute a task. To be able to perform as requested, the actor and the agent need to have a pattern of actions at their disposal. In a supply chain, tasks even ones as simple as packing a box with goods, require a prescribed set of actions.

If an agent itself is assumed to deal with the details of executing the task (autonomy), it needs to be able to deal with all the peculiarities which it encounters in its real environment. Dealing with a real life environment is problematic in robotics. In Chapter 4 we discuss robotics, the discussion focuses in particular on the process of information collection.⁴ We introduce the notion of a *cognitive basis* to describe what is knowable (understandable) by an agent, that is, what information it can pick up from the environment. The analysis is resumed in Chapter 6 where we explore cognition and action-driven behaviour of humans. We develop an epistemic framework within which we can relate and compare the know-how knowl-

⁴ To the reader's information I have to note that I owe the line of thinking, which underlies the analysis of Chapter 4, to the book of Dretske [1984], *Knowledge and the Flow of Information*. Ever since I first read the book, its mechanistic picture never left me, though I am inclined to change the book's title into "*Knowledge and the Flow of Data*".

edge of different actors. Non-verbal cognitive notions based on actions are defined: they can be shared among quite different agents.

1.6.4 Conflict Resolution

Once tasks are specified and the actors have know-how available, the third point (see Table 1.1) arises: while executing tasks conflicts might occur, for instance when two tasks require the same resource. We distinguish conflicts occurring among tasks allocated to one actor from conflicts that occur amongst actors. The conflicts of the first type have to be solved by the actor. Conflicts amongst actors, in general, concern shared resources and general solutions to such conflicts are not easily obtained.

Besides being considered as problems, conflicts can be viewed as a source for interaction and may give rise to collective behaviour. Supply chains originate from the interaction among autonomously performing actors, they result as a form of collective behaviour. Examples can be found in silent trade and pirate trade. Silent trade provides a trading protocol which requires no direct communication. Pirate trade is also a protocol, though the borders between violence and compulsion on the one hand and voluntariness on the other hand are vague. In Chapter 3, we specify the context of the supply chain and set up more elaborated protocols.

Previously, we have made reference to ant colonies. A puzzling point in studying societies (either natural ones or artificial ones) is how patterns of collective behaviour can be traced back to the behaviour of individual actors [Mataric, 1995]. In an ant colony a supply chain emerges, but it does not seem that ants have an explicit notion of working in a supply chain. But in the human context, supply chains also emerge, even via pirate trade, while no one is explicitly aware of the supply chain.

We aim at a systematic approach to interaction. Interaction emerges when actors try to resolve a conflict. However, the general subject of interaction is very broad, and touches many different topics. We study interaction in Chapter 5, but restrict it exclusively to spatial interaction (as part of the art of moving). We develop an interaction model to study interaction in a multi-robot system. The robots occasionally meet each other and need to pass each other. The reactive behaviour of avoiding each other leads to team behaviour. We see that the resulting team behaviour is not simply a summation of the individual actions. It turns out that even in a context of non-social actors, some social behaviour is generated. The studies in Chapter 5 provide insights into how characteristics of collective behaviour (one of which is a supply chain) emerge in the interaction of autonomous actors.

1.7 Lines Through the Thesis

This introduction indicates that the art of moving objects fosters discussions on a distinct set of domains: organisation theory, robotics, artificial intelligence, and philosophy. We first discuss the organisation of a supply chain, the global view, in the Chapters 2 and 3. Next we explore the local view. We use the domains of robotics and Artificial Intelligence, in the Chapters 4 and 5, to deepen our insight into co-ordination and interaction. In the last chapter, we draw conclusions for the domain of philosophy and cognitive science.

As guidelines to the reader we make the following suggestions.

For the general reader: Chapter 1 introduces the problem statements to be discussed. The problem of organising a supply chain is analysed in Chapter 2 in historical examples of warfare (2.2) and mail services (2.3). The first sections of Chapter 3 describe in general terms the method we propose for analysing and designing a supply chain. Chapter 4 starts with a discussion of a robot's niche (4.1) and how a robot might fit into an overall framework (4.2). In 4.3 the design of a robot is discussed. In Chapter 5 we discuss spatial interaction in multi-robot team. General characteristics are discussed in 5.2. In Chapter 6 we continue with a philosophical discussion. Section 6.1 provides the general lines which are further specified in the other sections. Several examples are given throughout the chapter.

For readers specifically interested in robotics we advise the following plan of reading: Chapter 4 provides the embedding of a robot application in a supply chain followed by a discussion on robotics design. Section 4.2 specifies the application of robots, and 4.3 discusses the design of a robot. Sections 4.4 and 4.5 deal with sensors, after that we go into the role of the robot designer (4.6). In Chapter 5, the spatial aspects of robot behaviour are discussed. Section 5.2 deals with team applications. In Section 5.4 we discuss several specific obstacle avoidance procedures. Sections 6.1 and 6.2 might also be of interest to readers with a preference for robotics.

Readers with specific interest for Intelligent Agents and Distributed AI are advised as follows. Chapter 2 gives practical examples of applying protocols, they might be considered as examples for multi-agent protocols. For setting the context of robot intelligence, a quick look at Sections 4.1 and 4.3 might be useful. Distributed AI is dealt with in Chapter 5. In Section 5.2 the discussion has a somewhat informal setting, but in Section 5.3 a formal framework is developed. It is applied in Sections 5.4, 5.5, and 5.6 to specific multi-robot systems.

For readers interested primarily in the related topics of *philosophy* and *cognitive science* Chapter 1 provides an introduction, in combination with Section 4.3. The practical problems of robotics and artificial perception are first dealt with in Sections 4.4, and 4.5. Section 4.6 discusses the role of the designer and is also relevant. The discussion is resumed in Chapter 6.

In Figure 1.2 the above overview is summarised. Moreover, the distinction between the global view (Chapters 2 and 3) and the local view (Chapters 4, 5, and 6) is shown.

Figure 1.2: Schematic overview of the chapters.

2 Organisations and Supply Chains

In Chapter 1 we formulated two problem statements in relation to the study of supply chains. The present chapter and the next one are dedicated to the first problem statement, which was: *how can the activities of the actors in a (large) supply chain be arranged such that a (goal) directed flow of objects results?* For problem statement 1, we look at the supply chain from a global point of view. In this chapter we introduce three organisational models and discuss historical instances of these models in the domains of warfare and mail services.

2.1 The Historical Perspective

We start this chapter with a brief discussion of three organisational models: a silken thread, an authority-based organisation and a contract net. We compare the models by discussing examples of historical supply chains. The examples are taken from the military domain and the domain of mail service. In both domains, the supply chains incorporate rather large numbers of actors and resources. We analyse how the chains have been arranged in these domains.

Maintenance of a supply chain requires enormous effort. This becomes clear when discussing the warfare domain. We encounter the severe constraints which the supply problem imposes on the power of the army. In warfare our two problem statements are interwoven, as was expressed by Sun Tzu [490 B.C., p. 87]: *"The enlightened ruler lays his plans well ahead; the good general cultivates his resources"*. In the warfare discussion we will occasionally, contrary to what has been stated above, briefly discuss the second problem statement as well.

The second domain which we discuss is the organisation of mail services. The mail services are based on rather particular chains, which are still in use. We briefly compare the histories of the French postal administration "La Poste" and the mail services in the Low Countries before the Napoleonic periode (1795). The

French system was developed under the central authority of the king; in the Low Countries the service emerged amongst the main (relatively autonomous) cities.

2.2 Organisational Models

In this section we discuss three major organisational models and evaluate them. In the most basic supply chains, there does not exist a pre-determined *goal* for the chain, nor any overview of where the chain has to lead to. This was in particular true for the silk trade. On the occasion we call its organisational structure a *silken thread*. In many modern supply chains, the begin and end of the chain are known and are fixed in advance. A rough overview of the chain is mostly available. For such a context, we distinguish two more organisational structures, one based on a common authority, the other on securing a maximum of independence. We call the first the *authority-based organisation*. It occurs in a chain in which all resources are incorporated within a single hierarchical institution. Consequently, the division of labour is based on a delegated authority. The second structure is called a *contract net*, and is orthogonal to this. It is a chain in which actors co-operate on the basis of a case to case agreement or “contract”.

When evaluating these three models, we first consider the goal of the chain. In case the goal of the chain is unknown we find the silken thread; in case the goal is known we can have an authority-based organisation or a contract net. We further evaluate the models along the three dimensions in a division of labour which are distinguished in Chapter 1: *task decomposition*, *coherence* and *resource assignment*. In addition to the three dimensions, we also look at the *autonomy* of the actors.

2.2.1 The Silken Thread

In Chapter 1 we have discussed the silk trade and the silk route. The silk was produced in China and eventually sold in the Roman Empire. Characteristic of the silk trade was that it was not known which starting and ending points the chain connected. The supply chain for silk showed a primitive (in the sense of basic) organisational structure.

The silk trade was based on individual actors (merchants, sometimes cities) that bought/exchanged silk for other materials. The Chinese exchanged silk (but also rhubarb) for glass and stones (specially jade). The jade mines were along the route, and the silk was forwarded along the route. In the east, silk was exchanged for jade, in the western part of the route it was exchanged for various different products. (Note that there was no currency commonly agreed upon!) Silk was sent *en route*, but where it would end, no one had the slightest idea. The organisational

structure, if we can speak of any, is depicted in Figure 2.1. It is a chain of actors linked together. Each actor made his resources perform, but the actors did not particularly intend to form a chain, they were just linked to each other. The chain just happened to be generated.

Figure 2.1: The silken thread. Tasks are ordered according to incidental contacts amongst the actors.

The chain just was generated, but little system can be found concerning the division of labour. The division of labour emerged historically (maybe politically) and was oriented on geographical conditions. An actor with “power” over the (scarce) resources could claim the labour (and its profit). The Chinese kept the recipe for silk production secret for a thousand years. Moreover, as said in the previous chapter, the Persians dominated and controlled the part of the trade route that connected the Pamirs to the West, not least for the trade profits. The existence of the chain and the coherence in execution was indirectly guided by these trade profits. In the long run, if the trade had not been profitable, it would have dried up. At any rate, the silken thread can hardly be called an organisational structure. Certain actors happened to dominate (small) parts of the chain. Task decomposition is enforced by the resources these actors have available. Coherence is indirectly guided by trade profits.

2.2.2 Authority-Based Organisation

A real and fairly common organisational structure is one in which actors are put together into groups or teams, headed by a leader. Groups can be thought of as consisting of only several human agents, but they might just as well be complete companies. Group members are subordinate to a group leader, where the authority of the group leader is the major binding factor of the group. The type and range of the authority can vary considerably. Generally, the organisation is completely permeated by a hierarchical structure of delegated authority. Figure 2.2 shows a hierarchical team structure on top of a supply chain. Tasks are allocated group-

wise to team members that are subordinate to a common team leader. At the top of the hierarchy is the leader, whose authority spans the whole task.

Figure 2.2: An authority-based organisation. Every task subsumes under a single node in the hierarchical structure.

We briefly review the three dimensions in the division of labour. An authority-based organisation provides a basis for task decomposition. The teams are subordinate to a single authority and task decomposition follows the lines of the authority: this is inherent to the hierarchy of the authority. Every actor has a role or function assigned to it, otherwise he/it would not be member of the team. Moreover, the role of each actor within the group is unique. Identical roles for two different actors are generally speaking not allowed, unless one has explicit reasons to do so (double capacity for instance). Thus, a task is straightforwardly decomposed along the lines of the hierarchy. Moreover, access to resources is ordered along the same structure. An actor has a fixed set of resources assigned to it. Thus, assigning tasks to an actor results in direct availability of resources. Within an authority-based organisation, actors, resources and tasks are regularly linked. Moreover, when the organisation is appropriate to a task, a coherent blueprint for solving the task is provided by the hierarchical structure.

The rationale behind the hierarchical structure is that the levels of greater authority are presumed to be associated with a broader viewpoint on the task at hand. The most pronounced authority-based organisation is found in armies. The military view is expressed by Thorpe [1917, p. 40]: “*The central location [read: authority, J.P.] gives a general view of all parts in their proper inter-relations and reveals what should be required of each part. Such general view could not be had*

by any of the parts." The structure of hierarchy provides the overview of the supply chain, which is so obviously absent in a silken thread. However, as is obvious from classical organisation theory – see for example [Botter, 1967] – the group is founded as a juridical entity and ownership of the organisation is the origin of authority. However this does not guarantee that the line of authority indeed coincides with a more general view.

Each actor at each level is provided with a delineated span of control and access to resources. Moreover each actor has limited responsibilities. On the one hand, the authority backs the individual, in the sense that every aspect that lies beyond the competence of a single actor can be shifted to the authority (with the broader view). On the other hand, the broad view is a putative privilege at the higher levels. Thus, the effectiveness of an authority based organisation "(...) *may be increased or decreased depending on the accuracy of the global view provided by the higher authoritative node*" [Bond and Gasser, 1988].

A certain task might require resources divided over different actors. At execution time, co-ordination of execution has to be obtained via the same lines. Co-ordination of such tasks requires an extra hierarchical step, as indicated in Figure 2.2. However, the line of hierarchy is usually not considered to intervene at execution time. Thus, co-operation among actors on the same level, but subordinate to different branches in the hierarchy, is very difficult to obtain. Within organisations that have a clearly structured hierarchy of authority and division of tasks, co-ordination of execution easily turns into a struggle of competence between the teams or team leaders. Obviously, such struggles do not contribute to a smooth execution of the process. Another consequence is that the lower levels will accept only tasks that will be covered by its own competence, and avoid tasks that surpass it. The potential creativity which the actors in the organisation may have, or that might be generated amongst co-operating actors (synergy), will not be used and is lost.

To deal with these consequences, classical organisation theory focuses on the assembly of the teams. Teams might be either *product-oriented* or *function-oriented* [Botter, 1967]. Product-oriented means that the scope of activities of the team covers the whole process which a single product has to go through. In a function-oriented team, specialists of a certain type (or certain types) are brought together to carry out a sophisticated process. A supply chain is typically geographically spread, and it is impractical to work with a product-oriented team. A functional orientation of teams seems forced by the geographic nature of the enterprise.

2.2.3 Contract Net

The authority-based organisation can be juxtaposed with the *contract net*. (The metaphor stems from distributed artificial intelligence see [Bond and Gasser, 1988] for further references.) In this case the actors are totally independent of one

another and they meet only for the single occasion of executing the particular job at hand. It starts as follows. A certain actor, called the *task manager*, breaks down a task into sub-tasks. Subsequently, the task manager broadcasts a so-called *task announcement message*, such that actors can make a bid for the sub-tasks. The sub-tasks will be allocated to the best (in some respect) bidder. When the bid is awarded, the bidding actor becomes a *contractor* to the manager. Figure 2.3 shows a contract net on top of a supply chain.

The contract net metaphor is oriented towards selecting and involving the actors with the appropriate resources available. The actors are supposed to make a bid. Comparison of the bids enables one to choose the most appropriate actor. In the contract net, neither a method nor some rule for *task decomposition* is provided. It is presumed that the task manager just does it. This requires him to have a fairly good insight into what actors he can expect to make a bid, otherwise no (serious) bids will be received. Note that in an authority-based organisation, actors are categorised as a matter of course; in the contract net such a basis is missing. The contract net is intended to provide a solution for “optimal” resource deployment, however, its success depends on the manager only.

If a good decomposition is made, then, as in the authority-based organisation, a blueprint for a solution is available. However, as the (sub-)tasks are allocated to independent actors, there is no recipe for obtaining *coherence* in execution, nor is there an authority that can indeed bring the results together. On a task for task basis, actors have to settle matters of coherence; a prescription for achieving coherence is missing.

Figure 2.3: A contract net.

2.3 From the History of Warfare

The first domain in which we discuss supply chains is warfare. We make a brief excursion into history. The discussion focuses on the aspect of moving the army, getting the battlefield machinery on the move and keeping the army supplied. Marching and supplying the army is rather distinct from the real battlefield operations. In times of war armies have to establish and maintain supply chains. In fact the supply chain is often decisive.

An army is generally considered to be the paradigm of an authority-based organisation. In an army the span of the authority reaches very far, and the resources in particular are fully subsumed under the commanding authority. Resources, including the soldiers, can be applied whenever thought necessary.

We note that in the military supply chain, the distinction between passive elements and resources is not sharp. Of course, the supplies in terms of goods remain passive elements, but resources, in particular the soldiers, have to be moved (and move themselves) as well. The soldiers are organised in units. Each of the units is an actor. The formation of the units, such as divisions, battalions and platoons, is directly connected to considerations concerning the supply problem. The units are the actors for whom the movements are planned, as well as the actors that have to establish the supply chain. The formation of units is decided along the hierarchical lines. Designing a supply chain for an army gives some insight into how authority is delegated downwards.

The military supply chain is central in our discussion. We will illustrate the discussion with the tactics of Napoleon Bonaparte, the inventor of modern war. The discussion is particularly based on the works of Jomini [1855]. He first defined the term *logistics* to cover the whole subject of "*the practical art of moving armies*." Jomini was a staff officer under Napoleon. Later he changed service to the Russian army (with Napoleon's approval, it seems [Ogilvie, 1990]). Incidentally, it is said that the famous book entitled "*The Art of War*" of Chinese general Sun Tzu [490 B.C.] was Napoleon's secret weapon [Clavell, 1981, p. 11]. It is even said that all doctrines, recommendations, tactical solutions etcetera developed in the West can be traced back to Sun Tzu [Ogilvie, 1990].

Whatever the case, our study has received much inspiration from military experts. The end of the Cold War, Perestroika and the opening of the Iron Curtain have considerably changed the objectives of the defence forces. Until then, the main objective was a sustained defence of a large area. Now, the main objective is peace keeping performed by relatively fast intervention in crisis areas. Thus, in military logistics, the aims changed from establishing a solid, long-lasting supply chain, to defining minimal units carrying supplies for short periods. Such units can move

fast, but they must be backed by reliable supply structures. Establishing the supply chains for fast moving army units challenges conventional ideas on the topic.⁵

2.3.1 Food Supplies

Below, we sketch the problem of food supply in order to understand the complexity of preparing a military campaign. The French "Grand Armee" that served Napoleon consisted of several armies, each of which included about 100,000 soldiers. Napoleon started the Russian Campaign with an army of 453,000 soldiers [Encyclopaedia Britannica]. Maintaining these armies posed a huge supply problem. To begin, we go a little further back into history, to the transition from the sixteenth to the seventeenth century. This period was characterised by a spectacular change in warfare, in particular, an immense growth in the sizes of armies. In 1567 the Duke of Alba impressed by taking along an army of about 10,000 men to suppress the revolt in the Netherlands. Only one century later, in 1672, Louis XIV mobilised 120,000 men to deal with the Dutch [van Creveld, 1977]. What makes this spectacular is not so much the fact that it was possible to mobilise so many soldiers, but that it was possible to maintain and supply such huge armies. To understand the supply problem we discuss two major aspects: the limitations on food production and the constraints of transport.

Agricultural production at the time was strongly dependent on a natural equilibrium, which imposed strict limits on the total production. Among other factors, the so-called *seed ratio* was important [Slicher van Bath, 1960]. A ratio of 1 to 3 was not exceptional in those times, meaning that one third of the harvest was needed for seeding next year. If overconsumption occurred, the subsequent regeneration of production could take years. A key factor was the interdependency of the local population and the production of food. The population was needed to plant, grow and harvest the food, but the population needed to be fed as well and hence consumed a considerable part of the seed. The introduction and distribution of the potato, in the seventeenth century, relieved the supply and food problem. The same area of farmland grown with potatoes could feed two or three times as many people [Slicher van Bath, 1960]. Nevertheless, local food resources remained restricted by production limits.

These facts are important for the supply of an army. The size of armies increased by a factor of 10, and the need for food supplies grew equivalently. Van Creveld [1977] provides some calculations. An army of 100,000 men, in ten days, marching through an area of a hundred miles long and ten wide, consumes about 1/10 of

⁵ A consequence of the change of objectives is that modern military aims show similarity to those in the Roman Empire. Roman military forces, in particular the navy, were relatively small and also designed for quick interventions (see [Peddie, 1994]).

locally grown food. From this it is easy to estimate how quickly a delay in marching leads to local food shortage. Assuming that there is only one harvest per year, military operations, and in particular sieges, could easily exhaust a large area of its food supplies.

The large armies of the seventeenth century had to face huge (food) supply problems, not only because of the enormous numbers of soldiers, but also because the tactics of war consisted mainly of sieges. We recall Sun Tzu's citation in Chapter 1, that sieges were the worst of all military policies. In a siege, the army inevitably has to be supplied over distances.

Transport of supplies also faced severe limitations. In Roman times, when moving grain over a distance of 180 km by horse power the costs of transportation equalled the value of the grain. In the seventeenth century, due to technical improvements the distance covered was three times more [de Graaf, 1996]. However, limitations on supply distances remained. To understand the difficulty we consider the following facts, "(...) a team of 6 horses consumes as much food as 30-40 men, the men together can carry more on their backs than the 6 horses can pull (...)" [van der Laan, 1994]. Oxen move much slower, and consume only half as much food as a horse [Slicher van Bath, 1960]. Nevertheless, for a long time, the cheapest transport has been by ship⁶, and we must admit that this still is the case.

The end of the sixteenth and beginning of seventeenth century covers the Eighty Years' War in the Netherlands. The campaigns of Maurice of Nassau elucidate the interplay between siege war and transport. At that time the area was called *The Low Countries*. The reason for war was the revolt of the Dutch against Spain. Even more than at present, the Low Countries were intersected by an abundance of water-courses. The main access to the cities and their most important mutual connections were the water-courses. Maurice of Nassau, commander in chief of the Dutch armies, showed himself adept at utilising the water-courses to move rapidly around with his armies and their supplies. He often managed to meet and besiege the Spanish, before they could make themselves ready for defence. Water transport was essential for Maurice. Flanders was occupied by the Spanish as well, but the area does not have such a dense network of water-courses. Maurice was never very successful in Flanders.

Most of Europe is not provided with so many water-courses, and hence the large armies in their overland campaigns were confronted with serious supply problems. Van Creveld concludes that the famous marches of the age (specifically the one of

⁶ Within the Roman Empire water routes were the primary (cheapest and for bulk the fastest) connections within the Empire. The military campaigns of the Romans followed the major water-courses, as the Roman invasion of France demonstrates, see [Peddie, 1994]

Prince Eugene in 1701) in fact were “flights forward” [van Creveld, 1977, p. 39]. As supplies could not be moved towards the armies, the armies had to move towards their supplies. The armies were doomed to be on the move, for no other reason than to forage.

At this point, we close our overview of the background of the supply problem. By now it will be clear that large armies, modern as well as Napoleon’s Grand Armée, face considerable supply problems.

2.3.2 Autonomy in the Army

At the beginning of this chapter we stated that we would focus mainly on the first of our problem statements, that is *how to arrange the actions* of actors to obtain a supply flow. We discuss warfare only in this section. We deviate from the line of discourse to make some remarks concerning the second problem statement. It concerns how to *make a set of resources perform as an actor*, in the military perspective. In the Subsection 2.3.3 we return to the first problem statement or the “concert of actions” as the topic in the military field is called.

An army is the paradigm of an authority-based organisation. At the top of the hierarchical pyramid is the commander-in-chief, or whatever his designation. A great deal may be said on the subjects of authority and discipline, but it is not so relevant for our studies. Of interest however is how the army is subdivided into units (and actors). Primarily, armies are organised according to the battlefield operations and requirements. In sub-dividing an army the supply problem has to be taken into account. We will first look at the subdivision of the army into units, and at the connections to the supply problem.

In the Napoleonic era the army was mainly divided into cavalry (horseman) and infantry. The distinction represented a distinction in moving speed. Cavalry, with its horse power, was fast in moving but as we have seen above, moving by horse power is expensive in terms of food consumption and requires careful planning of supplies. For the slower infantry, (slower but cheaper) oxen were often used. Cavalry and infantry were distinguished not only for the battlefield operations, because they required different logistics as well.

The *autonomy* is defined as the independence of units. In the army it is determined by their degree of self support. Modern armies are divided into brigades, each of which incorporates all elements of the whole army. Operationally a brigade is independent (or autonomous). Its size is about 2,500 men. Its logistics’ independence is fifteen days. For this period it has ammunition, food, etc. available. The smallest unit is a platoon, which consists of 40 men. A platoon is self-supporting for several hours only [van der Laan, 1994].

So much for the supply characteristics of the units. If we look at how the units in the modern army can be applied and what the scope of their operations is, the constraints set by the supply problem become even more striking. It takes a considerable period of time to move an army unit. This period is much longer than the period for which the application of the unit in the battlefield is predictable. For the tactics of war, one can predict the behaviour of a brigade within a range of 6–30 hours. However, the time needed to arrange the supplies and execute the moving of the brigade and its supplies is 36–48 hours. For a whole army the span of application ranges between 48 hours to eight days: moving the army takes between two weeks to three months [van der Laan, 1994].

The inertia of moving the army is due to the fact that the “fighting machinery” is considerably smaller than the “supply machinery” behind it. In military terms this is expressed by the *tail-to-teeth ratio*. In the Vietnam war the tail-to-teeth-ratio of the American army was 3.5 to 1 [van der Laan, 1994].

2.3.3 The Concert of Actions

Having discussed the formation of units and their supply characteristics, we examine Napoleon's style of warfare and how he set up the supply chain for the Grand Armée on campaign. The discussion addresses our first problem statement: *how to arrange the actions of actors* (or units).

Napoleon invented a new style of warfare. Contrary to the siege tactics of the previous era, he introduced the movements of armies as a tactical principle. His plans of war consisted of *planning* the movements of the armies to obtain a concentration of forces at particular (weak) points at the front battlefields. “(...) *a few examples should be given as illustrations of the great importance of a good system of logistics. One of these is the wonderful concentration of the French army in the plains of Gera in 1806; another is the entrance of the army upon the campaign of 1815. In each of these cases Napoleon possessed the ability to make such arrangements that his columns, starting from points widely separated, were concentrated with wonderful precision upon the decisive point of the zone of operations; (...) The emperor was his own chief staff officer. (...) Turning his dividers about from point to point on the map, he decided in a moment the number of marches necessary for each of his columns to arrive at the desired point by a certain day; then, placing pins in the new positions, and bearing in mind the rate of marching that he must assign to each column, and the hour of its setting out, he dictated those instructions which are alone enough to make any man famous.*” [Jomini, 1855, p. 263–264].

The planning approach might seem at first sight a simple arrangement and ordering of operations. However, as we have seen, marching over such long distances as Napoleon ordered, required an immense supply organisation.

The mobility and the range of operation of the fighting machinery is considerably limited by its supply tail. Thus, distance and space can be used as a means to defend. Moving the supposed battlefield over long distances into unknown terrain in which no supplies can be stacked in advance exhausts the supply tail. This tactic was applied by general Barclay (the Russian commander-in-chief) against Napoleon. Jomini [1855, p. 30]: *"Ever since Ivan the Terrible brought nationalism to Russia, the military leaders of that nation have been with maintaining the integrity of the country's far-flung borders (...) the Russian military mind became accustomed to think in terms of "war in space," realised that territory could be traded for time, and saw that the endless steppes of Russia offered a military barrier more effective than any fortification constructed by man."* Not only Napoleon but Hitler as well, failed in this invasion! This underlines the heavy burden of establishing and maintaining a supply chain.

According to Van Creveld [1977] Napoleon was always aware of the risk that the supply chain could fall apart. As a remedy, he managed to keep the army on the move most of the time. The ability to keep the army on the move is supposed to be the genius of Napoleon. We conclude that he fought his battles more with logistic ingenuity than with any other skill, but also his major defeat (in the Russian campaign) was due to logistic factors.

2.3.4 The Military Supply Chain

Concerning our *problem statement 1*, we have seen that Napoleon extensively planned the movements of the diverse armies and that the planning indeed consisted of co-ordination in time and place. In the citations above, it is only said that Napoleon made appropriate arrangements. According to Jomini the arrangements show the genius of the general; illustratively, he calls it the *concert of action*.

On the operational level the planning required a detailed arrangement of the supply chain. We have calculated the enormous supply needs of a huge army. Napoleon used to send out staff officers to arrange (and stack) supplies, in advance of the arrival of the armies. Incidentally, Jomini himself, was an officer of such a department ("Maréchal de Logis"). Napoleon's methods were supported by a perfect personal understanding, which worked well in the campaign to Gera (1806). The Russian campaign (of Napoleon) in 1812 showed however that the practical limitations of the supply chain undermine the planning approach.

For a modern army, planning is still crucial. The period for which the tactical battlefield planning is valid is many times shorter than the period needed to prepare the movements of the army with its supply tail. The planner needs a thorough understanding of what will be achieved after the (sub)operations are executed. In the next chapter we develop a systematic approach, in which (similar to Napoleon's intuitive approach) each task (or part of an operation) is described explicitly in

terms of times, places, and states-of-affairs that will be achieved as a result of executing the task.

In Subsection 2.3.2, we saw that the formation of units in the army is not only based on battlefield operations but on supply needs as well. The supply independence determines the range of autonomy of the actors or units. In the army, the resources are all available, that is to say, they are within the span of control of the commander-in-chief. Once, the actors have all the means available to execute their part of the operation, it is their experience in executing the operations, that provides the necessary basis for success. We have seen an illustration in Maurice of Nassau and his army, who were accustomed to moving their supplies by water-courses but were not used to travelling over land.

Authority

In warfare, the plans to confront the adversary are made on an ad hoc basis. It is the task of the commander-in-chief to decide on the plan of campaign. He *decomposes* the plans into *tasks* and provides the *coherence* in the chain of actions. This implies that the authority must hold an accurate view of the actual situation and must be well informed about it. The Grand Armee, passing through the desolate regions of Belorussia, had to rely on supply from base. By June 1812, the logistical co-ordination was lagging so far behind and the real situation differed so much that "(...) *the Emperor's orders, (...), were becoming impracticable of execution*" [Thorpe, 1917, p. 16].

Our hypothesis is that one has to find a balance between, on the one hand, the plans and planning of the authority, and on the other hand, the autonomy granted to the execution level, to the units or actors. Thorpe [1917, p. 40] argues: "*Centralisation in control and decentralisation in execution rest on the same principle; namely: that the operation is left to the functionary, or functionaries, who is, or are, in the best situation to estimate the operation.*" The actors should be authorised to act autonomously, within the plan laid down by the authority. Authority exercised over long distances is not sufficient to direct the army, as the tragedy of Napoleon's Russian campaign showed.

Information and Communication

In Chapter 1 we argued that the information system is not the primary constituent of a supply chain. Communication generally appears to be a difficult subject in warfare. It is not common to share the full plans and framework with the whole army. Sun Tzu [490 B.C.] advises: "(...) *confront your soldiers with the deed itself, never let them know your design*" [p. 81]. "(...) *he [the general J.P.] drives his man this way and that, and none knows whither he is going*" [p. 77]. But, if so, how are the actors directed such that they execute appropriately the tasks which are incorporated in the plans?

Again we encounter the importance of a good logistical system, now as the basis for effective communication. Thorpe [1917, p. 16] writes: "*Logistics furnishing the service of information and discharging the duties of communication, enables the tactical formations to be made promptly.*" Co-ordination is a prerequisite for internal communication and provides the necessary transparency which the units need to understand the order. The genius of the commander-in-chief and his staff has to be proven by the performance of the supply chain. Having achieved a well oiled supply train, global orders suffice to have the units execute their tasks. Only then it is appropriate to state that "*It will then be enough to indicate to them the point to be reached and the object to be attained, the route to be pursued and the hour at which they will be expected to be in position*" [Jomini, 1855, p. 260–262]. And the commander-in-chief can dictate "*(...) those instructions which are alone enough to make any man famous.*" [Jomini, 1855, p. 263–264]⁷.

2.4 From the History of Mail Services

In this section we discuss examples of supply chains found in the history of mail services⁸. In mail services as well as in warfare, many actors are performing. However, contrary to warfare, the aim of a mail service is not to get actors on the move, but to get the passive elements, the mail items, flowing.

The flow of letters and the performing actors are well distinguishable, although obviously there is no flow without actors. The metaphor of the coins in the coffee machine, given in Chapter 1, applies rather well for mail services. This makes the mail service domain particularly interesting for our study; the problem of *how to arrange the chain* can be treated rather separately and independently from the aspect of how to organise actors with the appropriate resources (problem statement 2).

Below we sketch the basic operations of a mail service network. Generally a (long) history precedes the generation of a mail service network. A service network is not something about which one today decides and that will work tomorrow. We provide three examples. They are interesting from a historical perspective, but what is more, they are examples of the three organisational structures, the

⁷ It is interesting to note that the Romans seem to have been well aware of the importance of the military supply organisation. Josephus, discussing the Jewish wars (65 AD) states: "*(...) They regard success due to luck as less desirable than a planned but unsuccessful stroke.(...)*", cited in [Peddie, 1994, p. xi].

⁸ Major parts of this text are from [Penders, 1995c].

silken thread, the authority-based organisation and the contract net, distinguished in Section 2.2.

Photograph 1. Mail exchange, translations: left text “Scene at the mail office. Print dating form the seventeenth century.”; right upper corner “From the history of the domestic mail service”; right lower corner “USSR Post”.

2.4.1 The Geographical Orientation

Mail services have a typical “hard-core” in a geographically oriented set of operations that guides the flow of letters through the system. The basic activity expected from a mail service is that it can *pick up* a letter or parcel *anywhere* in its “rayon” and *deliver* it anywhere else. Any mail service comprises the following three functions: *acceptance*, *transport* and *delivery*. This is true even of a point-to-point message service. However, most of the current mail services operate in a network structure.

For instance, currently in the Netherlands the postal company has divided the country hierarchically into main regions, sub-regions, districts etcetera [Bruijs, 1980]. The smallest area is that of the postman’s round. Moreover, every region has a “processing unit” which ranges, depending on the level, from a *sorting centre* to the postman.

In a network structure the three functions can be combined into three network operations: *collection*, *exchange* and *distribution* [Damen, 1994] (we give an overview in Figure 2.4). Collection aims at assembling larger sets of letters for transport, distribution is the opposite: splitting of the sets. Two processing units are connected by an exchange function, which generally involves transport. Damen [1994] distinguishes (process internal) levels of transport units (TUs): in collection TUs are packed together into a higher level TU, in distribution they are unpacked. Disregarding any organisational model, the process always starts with *acceptance* and ends with *delivery*.

Figure 2.4: The basic structure of a mail service.

The roles of actors are clearly delineated in the network structure and its geographical layout. A letter is picked up by the postman (collector) and forwarded to the closest higher level processing unit (actor), which in turn forwards it to the closest higher one and so on. The choice of which actor is next in processing is simply determined by a comparison of the address on the letter and the geographical layout of the network. As a result, the route of each item is predictable from the very beginning. In particular, while the process is running, every receiving actor knows where to forward the mail piece: no particular instructions are needed. No communication is necessary about how to proceed with and forward the mail

items. The necessary prerequisites are provided by the network structure and protocols.

2.4.2 Historical Notes

We discuss three historical examples of mail services, situated in north-western Europe, in particular France and Holland. Before going into the specific examples, we try to give an idea of the background of the period.

As is known, at the end of the Middle Ages the power of the cities (starting in Italy and Flanders) grew with the emergence of trade, at the expense of the old feudal system. The cities, as we will see, became major actors in setting up (public) mail service chains.

In the sixteenth century, the major power in Europe was the Holy Roman Empire later the Habsburg/Spanish empire. The empire included Spain, much of Italy, the Low Countries and Portugal as well as an impressive array of colonies. Though the Low Countries⁹ were formally part of the empire, the imperial administration was not very effective, especially not in Holland. At first, the area was not of much interest, and thus at the start of their economic growth, the Dutch cities were left alone. Later on, the troubling political affairs and revolt which arose in the Low Countries were considered by the Imperial court as a minor local problem. The emperor focused his attention on international relations with the Ottoman Sultanate, England and France [Parker 1979, p. 19].

In the middle of the sixteenth century, Antwerp was Europe's major trading city. Antwerp participated in the revolt of the Dutch cities. However, in 1585 Antwerp was recaptured by the Spanish and separated from the Northern Low Countries. With the recapturing it was believed that the major pillar of the Dutch revolt had fallen. However, the consequence was that many merchants moved from Antwerp to Amsterdam and thus trade shifted as well.¹⁰ At the end of the sixteenth century Amsterdam became the world's trading capital.

⁹ The Low Countries include Flanders and Holland. The names Holland and Flanders strictly designate smaller areas than the Northern and the Southern Low Countries. We will use both names loosely, when convenient and no confusion is expected.

¹⁰ The population of Antwerp fell from 84,000 in 1582 to 42,000 in 1589 [Parker, 1979].

2.4.3 Bilateral Mail Connections

Within the Habsburg/Spanish Empire there existed a mail service owned by the (originally Italian) family *Taxis* (or Thurn and Taxis). The family was granted a monopoly on mail services by the emperor. We will not discuss the Taxis service extensively, but as it influenced the mail services of Holland, Taxis will be mentioned every now and then. Similar to many other imperial arrangements, the Taxis monopoly was not effective in Northern Low Countries (Holland etcetera) but Taxis did operate and had an effective monopoly in the southern Low Countries (Flanders etcetera). In the next pages we discuss the mail services which developed in the Kingdom of France and in Holland. The Taxis company geographically separated these mail services.

City Couriers

The history of the (public) mail services is connected to the relative independence and growth of the cities. Historically, the use of mail to communicate emerged from the ability to write and read. Thus, it is not surprising that the first regular mail services developed within the cities. In general, the cities entrusted the mail services to privileged professional couriers. The first written references to such privileges are rather old, e.g. Strasbourg tenth or eleventh century [Lauffer, 1954, p. 24], Cologne from fourteenth century, Amsterdam at least since 1395 [van Overvoorde, 1902, p. 76].

Cologne developed a long standing tradition of city couriers. The Taxis service was present in Cologne but could not beat out the city couriers. The Dutch cities played a role here. They were dominant in trade and in general preferred to do (mail) business with others than Taxis. Thus, it was not until 1751 that Taxis regulated all mail services in Cologne.

Intercity Connections

With the growth of trade, the need to exchange mail between the cities grew as well. We focus on the city of Amsterdam. In its glorious period Amsterdam maintained quite a few bilateral (long-distance) mail connections. Figure 2.5 gives an overview of the major (long-distance) connections to and from Amsterdam. The most important ones were: Antwerp-Amsterdam, Cologne-Amsterdam and Hamburg-Amsterdam. We have drawn some onward connections and two neighbouring connections: Antwerp-Cologne and Cologne-Hamburg. These connections were for most of the time, but not always, operated by Taxis.

Figure 2.5: Major bilateral mail connections of Amsterdam.

The importance of the mail connections is shown in Table 2.1. The city of Amsterdam had specialised offices for single destinations, one office for Antwerp, one for Cologne and one for Hamburg. Further it had two offices for the inland (Holland) connections.

Table 2.1: Yield from the mail forwarding in 1752 [van Overvoorde, 1902].

From	Yield in Dutch florins	To	Yield in Dutch florins
Amsterdam	383.600 (total)	Antwerp	100.000
		Cologne	62.000
		Hamburg	100.000
		England	80.000
		other destinations and inland	41.600
The Hague	103.000 (total)		
Rotterdam	64.000 (total)		
Other towns	77.200 (total)		

To indicate long-distance speed in mail services we provide the following table from around 1650 [based on Ten Brink, 1969 and Van Overvoorde, 1902]:

Table 2.2: Delivery times for long-distance mail connections.

Connection	delivery time in days	notes
Amsterdam-Antwerp	4	
Amsterdam-Brussels	4	
Amsterdam-Paris	11–12	
Amsterdam-Vienna	18	
Amsterdam-Genoa	36	
Amsterdam-England	7	
	4	from 1660 onwards
Amsterdam-Hamburg	6–7	
	3.5	from 1634 onwards
Amsterdam-Petersburg	38	
Antwerp-Hamburg	4	(Taxis, 1660)

As mentioned, at the time the cities were rather powerful and sovereign. This was in particular true of Amsterdam and Hamburg. As a consequence, the long-distance connections were, in a certain sense, silken threads. Local actors, not only the cities but Taxis as well, connected their services to their neighbours to form the threads.

The bi-directional Amsterdam-Hamburg connection was an important communication line and rather profitable as Table 2.1 shows. Moreover, this connection was a dominant link in many other mail threads and is illustrative of the phenome-

non of a silken thread. The cities of Hamburg and Amsterdam granted each other monopolies on the onward connections. Hamburg connected not only to Northern Germany, Sweden, Denmark, but also to Poland and Russia. Amsterdam was the intermediate to connect to England (until 1794) and to France [van Overvoorde, 1902, p. 184–186; ten Brink, 1969].

In England, foreign mail exchange was monopolised from 1568 onwards [van Overvoorde, 1902, p. 45] thus England had a single link to the continent. For some time the connection from Amsterdam to England had to pass through Antwerp. Spain still considered the Low Countries as its dependency, which implied that England formally could not have direct contacts with the city of Amsterdam. Only in 1668 was a direct contract between Amsterdam and England signed [van Overvoorde, 1902, p. 250].

The dominance of the Amsterdam-Hamburg link was mainly based on the mutual agreement between two local (but powerful) actors. In the given circumstances it could defy even the Taxis company which was the Spanish emperor's favourite. We have noted the endurance of city couriers in Cologne. Moreover, the Hamburg-Amsterdam connection was an important and direct competitor for the Taxis system; Taxis also forwarded mail from England via Antwerp and Cologne to Hamburg.

2.4.4 A Royal Service

The French mail service *La Poste* is the oldest governmental postal system of modern western Europe. It was established in 1464, in order to secure message exchange between the King in Paris and the regions which he ruled. The original structure of "La Poste" was straightforward. Starting in Paris, the connections spread out to the regions, as depicted in Figure 2.6. The routes were established by a system of livery stables, at distances of four hours travel. The livery stables were to have four or five fresh horses ready.¹¹ The couriers (called postillons when on horse back) carrying the mail were sent along these routes and changed horses at the stables. The system was intended to deliver letters from the King in Paris and return responses to Paris.

The original organisation of *La Poste* is centralised and *authority-based*. The postillons were civil servants, entrusted with responsibilities which surpassed the basic functions of acceptance and delivery; they sometimes had to control the content of the letters as well. The flavour of an army-like organisation is distinctly present. Generally a message was assigned to a courier, who personally had to

¹¹ A similar system of livery stables was in use for the imperial courier service in the Roman Empire [Peddie, 1994]. These couriers passed messages mostly verbally.

deliver the message. Originally the service was intended for government mail only and the couriers were forbidden to carry other mail items.

A typical detail: there existed no confidentiality of the mail within La Poste, on the contrary, sometimes orders were issued to check the contents of the letters. In 1495 for instance, the courier had to check whether the contents of the mail opposed the Council of Basel.

At the start, La Poste was not a real network. The three main functions found within the system were: *accept*, *transport* and *deliver*. *Collection* and *distribution* processes were limited to combining mail items for one journey of a courier, the courier himself had to travel the whole route. Within the structure there was no real exchange or sorting done. The nodes in Figure 2.6 represent exactly what [Damen, 1994] calls merging and splitting.

Figure 2.6: La Poste, France (in the early stage, 1464).

Gradually La Poste accepted more private mail, but until the beginning of the seventeenth century, the service for private mail was considered a favour to the public [van Overvoorde, 1902]. To give an idea of the size of the system: in 1630 it connected twenty cities with (post) offices.

Besides La Poste, there existed in France (and Western Europe) a mail system, for more or less public use, organised amongst the students of the University of Paris. This mail system gradually disappeared (and was purchased by La Poste later, 1643). In 1728, all city couriers except for La Poste were forbidden, and a full monopoly for La Poste was established.

2.4.5 A Contract Network

In the "golden age" (seventeenth century) of the Northern Low Countries, the country consisted of a federation of seven republics. However, there was no strong central government (quite different from France). The real executive power in the country was exercised by the cities, the most important ones were Amsterdam, Leiden, The Hague, Rotterdam, Utrecht and Gouda. Each of the cities had its local couriers. Parallel to the trade interests, the cities maintained bilateral daily mail services.

As noted before, Holland was intersected by an abundance of water-courses. Transport of merchandise mainly utilised the water-courses. Moreover, the water-courses were the major connections between the cities. The connections over land were rather poor. Nevertheless, travel over land was faster. The mail exchange was executed first by couriers on foot. After the Eighty Years' War (1648) the foot couriers were replaced by horseman, the latter were called *postiljons*. The use of carriages was forbidden until 1756, because of the lower travelling speed.

It is known, in particular from the contracts on the connection Amsterdam-Antwerp, that the courier had to deliver the mail item himself, i.e., the courier from the one city had to deliver the mail items personally in the destination city. The service fees were paid by the receiver, to ensure proper delivery. Thus both couriers on a bilateral connection had to travel the entire distance.

Though only a detail, a remarkable difference between La Poste and the mail services of the Low Countries was confidentiality. Because of business sensitivity in the latter case, confidentiality of mail was a paramount concern, and couriers were even forbidden to show addresses of mail items to others.

Daily Exchange

For reasons of economy, it was desirable that the couriers exchanged mail half way. So special arrangements to divide or exchange the fees had to be made, and one had to define a real exchange operation. An impression of the practice is given by Photograph 1. The cities of Utrecht and The Hague took in 1590 [van Overvoorde 1902, p. 465] the initiative to arrange a mail exchange office in Alphen aan de Rijn, a small town approximately in the centre of Holland. The office in Alphen grew out into the exchange office serving five daily (nightly) routes from and to: Amsterdam, The Hague (via Leiden), Gouda, Rotterdam and Utrecht. The full system already existed before 1677 [van Overvoorde, 1902, p. 272].

These five routes formed the kernel of a geographical structure of operations, with a real exchange between them, see Figure 2.7. The structure functioned without there being a central ruler or authority. Note that despite its strong position, Amsterdam could not impose its will in the area of "national" mail services. The cities together defined a network of mail links, connections and an exchange procedure.

The execution of the service was granted as a privilege to so-called “postmasters”. The postmasters in turn outsourced the work to sub-contractors. Contracts lasted usually for a three-months period only, and prolongation depended on quality (and price) of delivery [ten Brink, 1969].

Figure 2.7: Mail connections in the Low Countries (1677).

The established mail service network is an example of a contract net. The cities as a kind of third party, a *gremium*, defined the service network. For the execution they appointed postmasters (task managers). The postmasters in turn let actors make bids on the sub-tasks and arranged sub-contracts.

Continuation

The Dutch cities formed a gremium that enabled a (literal) contract net for daily mail service. The task managers were appointed by the individual cities. Later on, as the central authority in The Hague grew, the role of the gremium was transferred to the central authority. In fact it was a transfer of rights from the cities to the state of Holland. It started with the installation of the *Statenpost* in 1745. This transfer is by some authors called a *centralisation* of the mail service. However, it was not a centralisation in the sense that La Poste in France was centralised; performance of the service remained with sub-contractors. Amsterdam only transferred its rights in 1748 [van Overvoorde, 1902, p. 94]. The city's hesitation being due not only to the loss of income, but above all "*to lose control over such an important influence on commerce and trade.*" Typically, Amsterdam retained the monopoly on the foreign mail connection Hamburg-England.

The decentralised (or distributed) service operation in the Low Countries remained: the cities were too keen to retain their independence. Nevertheless the service functioned quite well, and attracted the attention of Catherine the Great of Russia, who asked her ambassador in The Hague in 1782 for information [ten Brink, 1950, p. 12].

Besides the mail network discussed, there did exist in the Low Countries a mail system similarly organised as La Poste. It was called "Staatspost". It was an administration of 22 couriers and four postillions, serving only the parliament and departments of state in the city of the Hague and its surroundings [ten Brink, 1950, p. 13]. This system never developed into a wide ranging mail service.

2.4.6 Towards Modern Times

We have discussed three historical examples of mail services, the bilateral Amsterdam-Hamburg connection, La Poste and the mail contract net. Before resuming our supply chain study in the next subsection, we briefly overview the developments in the mail service domain till recent times.

During the annexation of the Low Countries by France in the Napoleonic era, the mail service was – after French model – made into a state monopoly in the year 1807 [van Overvoorde, 1900]. The initial intention was to upgrade the service, in particular for the smaller towns and rural areas. However, in effect it was regarded as a profitable source of state income [ten Brink, 1950; Hogesteger, 1989].

The official mail service was rather expensive, and did not really cover all areas of the country. As a consequence, there co-existed official and unofficial circuits of mail transport, in particular by shippers. (Note that water carriage was the prevailing mode of transport inside the country.) Moreover, the official couriers, the postiljons, often served the official circuit as well as unofficial ones. The unoffi-

cial circuits had their exchange offices ("sorting centres") at inns at the edge of the cities [Hogesteegeer, 1989].

Whether the public was really satisfied with the mail services in the federal republic of the Low Countries (before 1794) is not known. However, complaints about the monopolised service (after 1807) are abundantly present in Dutch literature of that period. People attached stones to their letters in order to bypass the monopoly, thus a novelist writes: "(...) *nothing so dreadful than that on Saturday afternoon he (cousin Nurks) sent me by coach a stone with a letter attached (...)*" [Hildebrand, 1839].

Modern time

Most postal services nowadays are organised into a central administration (similar to the French model). The reasons for centralising were diverse, but centralisation is not a requirement originating from the basic activity of a mail pick-up and delivery service. Rather recently several European countries (the Netherlands, England and Germany) have privatised their mail companies. The European Union has not yet decided on a common policy. A proposal supported by the Netherlands, Sweden and Finland, for liberalisation of the postal monopolies in the European countries has met with strong resistance, particularly from the French, German and English lobbies [NRC Handelsblad, 8-9-1997].

In 1989 in the Netherlands the mail company was privatised and granted a (partial) monopoly. The State of the Netherlands, however, retains its role of gremium, as the institution that issues the monopoly. The national Dutch mail company is called PTT Post. This company acquired an international courier service and express service called TNT in 1996. TNT is one of several private companies that gradually penetrated (Western) Europe with long-distance express services for mail-like items. These companies operate on the edge of the services of the national and monopolised mail companies. With the acquisition of the TNT company, PTT Post gained control of long distance connections covering the whole of Europe and most parts of the other continents. We cannot help but draw a comparison with the role of the former Taxis company (in the seventeenth century).

The mail service domain is changing, not only at the international level but also at the national level. In the Netherlands initiatives follow one another rapidly. In 1995, a private person declared his intent to start a mail service network for weekly delivery throughout the country. He thought he could build his network on the basis of (2500) locally operating actors. Each of them had to manage delivery in his own region. Exchange of the mail would take place at a central sorting centre, transport to and from which being provided by a haulier. Payment of the local actors would be on the basis of the number of items delivered [Algemeen Dagblad, 25 July 1995]. Since this announcement, not much has been heard about it, and in 1997 it went bankrupt [Het Financieële Dagblad, 13-8-1997]. This initiative shows

that it is not difficult to obtain a desk top design of a service. However, in such a large scale network it is quite difficult to have actors perform as foreseen in the design. The comment of the former director: "*we were too small.*"

City couriers established the first regular mail services. Still nowadays, urban areas are of most interest. An energy company has decided to have its mail delivered by some "underemployed" (in Dutch: overcomplete) employees. They will service only urban regions, mailings for the rural regions will be left to the postal company [Het Nederlands Dagblad, 24-4-1997]. Also several local communities did set up their own local mail services, using employees from their Sociale Werkplaatsen (social welfare workshops) [de Volkskrant, 24-12-1997]. On a last note, in the nineteen seventies, in particular in the city of Amsterdam, the new profession of city courier on motorcycle developed. At present, most of these companies have been acquired by the national PTT Post company [de Volkskrant, 26-8-1995].

2.5 Conclusions from the Examples

In Chapter 1 we stated that a supply chain runs through the interface of the physical, the organisational, and the informational sub-systems (Figure 1-1). We distinguished three dimensions in a division of labour: task decomposition, coherence and resource assignment and defined autonomy of the actors. We argued that a division of labour has to be on a par with the independence (or autonomy) of the actors involved. In this chapter we discussed three organisational models, a silken thread, the authority-based organisation and the contract net. Which model applied also depended on the presence of a pre-defined goal. The warfare and mail service domains provided examples of the organisational models.

Below we compare the organisational models and draw conclusions from the given examples, by investigating: the goal, task decomposition, coherence, resource assignment¹² and autonomy. Table 2.3 provides an overview.

2.5.1 Silken Thread

The first model to investigate is what we have called *silken thread*. The silk trade was organised as a silken thread, and so were the bilateral mail connections. A silken thread is not designed (for a particular goal), the goal or final destination of the chain is open. In a silken thread, task decomposition emerges as a pattern of

¹² Note that resource assignment is not the same as resource management, we discuss the relation between these two notions in the next chapter, subsection 3.2.3.

actions that happens to have a result. Task decomposition is determined by those actors that happen to get involved and that have the resource available. Typically, autonomous actors link themselves into the chain. Resource assignment is accidental: it depends on which resources happen to be accessible for the actor. Thus, task decomposition is rather opaque, and depends on regular (local) practise. Actors act fully on their own account, they are autonomous. A consequence is that actors cannot be substituted for each other. If a certain actor is missing the chain is cut off. It is not exceptional that the actors strive to maintain opaqueness. The outcome of the chain is not determined at the start, and one cannot speak of real coherence in execution. A silken thread might emerge within silent trade and pirate trade. All that is required in a silken thread is that each actor relates to his neighbours. The nature of this relation is not important.

A prerequisite for performing in a silken thread supply chain is that the actors have to organise the background of the chain themselves. The background consists (see Chapter 1) of the physical, organisational and information systems. We discussed the example of the Hamburg–Amsterdam mail connection, and pointed out how these cities arranged their roles in the chain.

2.5.2 Authority-Based Organisation

In general, the goals of an *authority-based* organisation are fixed; they are determined at the highest hierarchical level. The organisation is “designed” to obtain certain results. In the case of warfare this is clear, but also La Poste served to obtain fixed goals. La Poste was initiated by the central authority, and its operational structure resembled the lines of authority. The exchange unit in the structure was located at the central authority’s base in Paris.

In an authority-based organisation the actors are subject to a hierarchy, which provides the required ordering of the background. The hierarchical structure provides an embedding of the chain in a mix of the physical, organisational, and information system. An authority-based organisation “guarantees” a fixed outcome to a known task.

The hierarchy provides an inherent structure for task decomposition, moreover resource assignment, and coherence of execution are fixed. The hierarchical structure presupposes that the higher levels have a more global view. One can say that the actors transfer their autonomy to the leader, and in return they obtain a well ordered place in the supply chain. However, the authority-based organisational structure has a particular tendency to reduce autonomy of actors in this sense.

In the warfare examples it became clear how important the accuracy of the global view of the leaders is. Moreover, the less one leaves to the autonomy of the actors, the higher are the requirements for the global view. To prevent overload of the

higher levels, it is in our view better that the authority-based organisation grants autonomy to the actors (delegated authority).

2.5.3 Contract Net

The bidding process is the essence of the *contract net*. The task manager has a central role. He must have an overview the process, because he has to provide a task decomposition. This implies that the goals are fixed. By initiating the bidding process, the task manager explicitly asks for actors with certain skills (and of course, certain resources). All the potential skills of the actors can be applied. In this sense the contract net is open and preserves the independence of the actors.

The Low Countries' mail exchange network was an example of a contract net. The way the network was set up is illustrative. The profession of courier was fairly well-known and already in operation, but the resources were distributed over many independent parties. The cities united in a gremium which defined procedures to exchange mail and fees. The cities negotiated on how to put them into an overall structure and delineated exactly what was expected of each party.

The contract net does not offer such a solid *coherence* as is inherent to the authority-based organisation. Coherence has to be arranged or enforced by the task manager, but the task manager and the contracting actors are in charge for this particular task only. The actors are basically independent and autonomous. Before the bidding process starts, much work has to be done by the gremium. Without a gremium in the background it might become impossible to obtain coherence. In the latter case coherence has to be achieved on the basis of much case-to-case communication.

Table 2.3: Overview of organisational models.

	Goal	Task decomposition	Coherence	Resource assignment	Autonomy of actors
Silken thread	open	from actor to actor	accidental	arbitrary	basic
Authority based	fixed by mission	fixed in organisation	fixed	fixed	transferred to leader
Contract net	given with task	provided by gremium and task manager	accidental	negotiated	(partly) preserved

2.5.4 Models and Present Supply Chains

We have seen supply chains of the mail services in which all three organisational models appeared. Practical factors determine which model indeed occurs. Each model has its pros and cons.

A silken thread results when actors happen to link to each other. The existence of the thread depends on each of the actors individually. Silken threads occur in situations where the actors are forced or allowed to order the background.

The authority-based organisation is developed from the power of certain actors. Its advantage is that it makes sure that actors and their resources are available. Moreover, they can be urged to act. Supply chains, however, can cover very long distances and authority exercised over long geographical distances easily vanishes. One of the causes, as we have seen in warfare, is that the viewpoint of the authority differs too much from the actual situation and hence an order can become impossible to execute.

A contract net in a supply chain presupposes that the chain is clearly ordered and transparent. The contract net is intended to involve the actors with the appropriate resources. In the background some gremium is required that establishes "standards". The gremium, for instance, has to delineate what a bidding process is, what a contract is et cetera. It has to make this clear for the actors, so that they can determine what their own competence is. Moreover, the task manager needs standards to be able to communicate the tender to the actors. Given that there is a gremium, actors can compete in a bidding process as defined in the contract net metaphor.

Though the silken thread might seem obsolete, it is generally the natural (evolutionary) origin for many types of trade. Moreover, in many branches it still exists. If the chain, for what ever reason, cannot be directly organised (as a contract net is organised), *free-market* mechanisms have to fill the gap in co-ordination. A free-market by definition is not ordered, and occurs typically in opaque silken thread supply chains. In this respect, our point of view coincides with that of Malone [1986], who considers *markets* as co-ordination structures or mechanisms. Whereas free markets are associated with silken thread supply chains, the concept of a market with *perfect competition* or a *perfect market* is associated with the context of supply chains organised similarly to a contract net.

3 The Fabric of a Supply Chain

In this chapter, we further explore the phenomenon of a supply chain and our problem statement 1: *how can the activities of the actors in a (large) supply chain be arranged such that a (goal) directed flow of objects results?* As stated before, we view the chain from the global perspective. Surveying examples of supply chains in the previous chapter, we have seen that they consist of a web of activities. The fabric of the chain is the subject of this chapter. We introduce an analysis method, and develop it into a design approach which is applied to a real case. In the real case the framework of the chain is re-defined and thus a re-design of its services is obtained.

3.1 Organising Supplies Activities

The supply problem is inherent to the division of labour as we argued in Chapter 1. Supply chains do not exist for their own sake, neither do they exist independent of the society. The aim of a supply chain is determined by the needs of the society (even in the case of warfare). Evaluation of the effectiveness of a supply chain is bound to the conditions and qualities of the society which it has to serve. Throughout history the evaluation criteria have changed. In present society, transportation speed has increased considerably. The transportation rate – that is the transported volume multiplied by the distance covered (ton-km) – increased even more. But, the speed of communication has increased most of all; it has made a revolutionary improvement. Communication speed leaves the transportation speed for goods far behind. This forces us to reconsider supply chains.

In this chapter we introduce our notion of a *structure of operations*. In Section 3.2 we discuss the basic points and argue that this notion filters out aspects which are common to all supply chains. Section 3.3 deals with the details of a structure of operations; moreover a methodology is developed. In Section 3.4 we apply the notion and its accompanying design method to a case of setting up a pan-European piece-goods distribution network.

Before starting our discussion, we would like to mention that the study in this chapter was in particular inspired by the further development of the European Union, and the new demands arising for the supply processes.¹³ The ideal of the EU is that the countries gradually integrate. However each country has developed its own variant on Western Society. It will be obvious that integration of processes, which are so much interwoven with local conditions, raises its own problems.

3.2 A Structure of Operations

A supply chain is scattered over space and time. Within the theatre of war (as the military like to call their business) armies have to move from scene to scene, and time of arrival is essential. In the Netherlands, a letter mailed in one city should nowadays arrive the day after in the other city. A supply chain consists in linked activities of geographically spread actors. This point is basic in our notion of a *structure of operations*.

Hoogeweegen [1997, p.17] compares several definitions of supply chains. His final definition, which is also close to our approach is: "*the supply chain is the chain linking each element of the production and supply process of products and services from suppliers to the end customer.*" This definition shows where to find supply, and besides production chains it includes also services (for instance mail services). However, it does not state of what the chain itself consists, what makes the chain a chain.

3.2.1 The Framework of Conditions

In a supply chain at a certain point in time, a certain *state-of-affairs* occurs: an operation starts and after some time another *state-of-affairs* results. An *operation* is an activity that changes some pre-conditions into certain post-conditions. More technically, an operation is a *transition* of the *start-state* into the *end-state*. Figure 3.1 represents this schematically.

In Chapter 1 we distinguished *passive elements* – for instance letters – from actors. When an actor has to execute an operation, the passive elements have to be available at the appropriate time and at the appropriate place. After executing the operation, the passive elements have changed place – that is, certain post-conditions are obtained. An operation (or action) works on the passive elements and is con-

¹³ This work has partly been carried out in the IDEAS project, sponsored by the Commission of the European Communities in the Brite-EuRam programme.

ceived of as nothing but the changing of some pre-conditions (e.g., the co-ordinates of the passive elements) into post-conditions (different co-ordinates are obtained).

Figure 3.1: An operation with start-state and end-state.

In general, in a supply chain several operations occur, they are collected in the *set of operations*. The post-conditions of one operation can be matched to the pre-conditions of another one; i.e., we can link operations. By going from state to state, a chain of operations results. We call the chain in which operations are arranged and ordered by a linking of pre- and post-conditions, a *framework of conditions*.

A supply chain can be modelled in a framework of conditions. The framework of conditions describes the pre-conditions starting at the point of departure. Going through several intermediate states, it eventually arrives at the post-conditions. When the proper post-conditions are obtained, the goal is achieved. The framework provides an overview of the chain, and explicitly describes the co-ordination aspect for the flow of the passive elements. The framework of conditions and the set of operations – together the structure of operations – provide a schema for dividing labour in a supply chain. Figure 3.2 gives an impression, the chain is represented on the bottom line, decomposed into operations. Each operation is assigned to an actor. The connections between the operations are described in the framework of conditions; coherence is obtained when the execution of the operations satisfies the framework of conditions. Thus, the framework serves for task decomposition as well as for obtaining coherence. However, how operations are assigned to the actors is left open.

Figure 3.2: A Structure of Operations.

A framework of conditions is present in any supply chain. In the warfare domain we encountered the planning approach of Napoleon. His plans of campaign consisted of an (implicit) enumeration of foreseen and necessary operations, and an (explicit) ordering of the operations into a plan or framework. In the domain of mail services, the basic operations of picking up a letter at one location and delivering it at a certain point of time at another location were performed by couriers. We discussed different examples of how these operations were arranged into a framework of a mail service network. In the system of La Poste the operations were linked straightforwardly to the central node in Paris. The cities in the Low Countries designed an elaborate framework of conditions within which the operations of the couriers were ordered.

The social and political changes in Russia, which started around 1990, are illustrative of the importance of the framework of conditions. In the winter of 1991/92 the food supplies to the capital Moscow were very poor. While the major political changes were in progress in 1989 and 1990 there was no problem in supplies; however in 1991 supplies became problematic. Roughly-speaking, food (passive elements) and all resources were available as in the years before, but the effectiveness of the supply system was considerably reduced. Our (speculative) diagnosis is that the framework of conditions which used to be provided by the central authority had fallen apart. This was because of the collapse, or because of the failing social recognition of the authority or whatever other reason. In any case, the society still had to "agree" upon a new framework of conditions. In the winter of 1991/92 a widely-accepted framework was missing. Fortunately, society can change fairly quickly, and one year later the urgency of supply problem had vanished!

We leave this example with our speculative diagnosis. Notwithstanding that, this example and the historical examples of Chapter 2, show that a supply chain consists of properly linked activities. Also a supply chain reflects the ordering of the background of the chain. In different (historical) backgrounds different structures emerged. The ordering of the background is important, and we have to take it into account. However, how the background ordering is achieved is not the subject of our study.

3.2.2 Relation to the Organisational Models

In our structure of operations we have a set of operations. The operations are arranged and ordered into a framework of conditions. Below, we relate the structure of operations to the organisational models, discussed in Chapter 2; moreover, we proceed with the overview given in Section 2.5 (cf. Table 2.3) and go through the same points: determination of the goal, task decomposition, coherence, resource assignment and autonomy, see Table 3.1.

The structure of operations provides an overview of the supply chain, and implies that the *goal* of the (partial) supply chain is known. It makes explicit the broader view which in an authority-based organisation is associated with the higher levels of authority and, furthermore, it provides a basis for distributing the process over relatively independent actors.

Table 3.1: A structure of operations.

	Goal	Task decomposition	Coherence	Resource assignment	Autonomy of actors
Silken thread	open	from actor to actor	accidental	arbitrary	basic
Authority-based	fixed	fixed	fixed	fixed	transferred
Contract net	given	by manager	accidental	negotiated	retained
Structure of operations	of given	by the set of operations	by the framework of conditions		

Task decomposition and *coherence* touch upon the essence of the flow in a supply chain. Determining the set of operations is making a task decomposition. The framework of conditions provides an overview of how actions are connected into a supply chain. It is a prescript for achieving coherence. The framework of conditions is developed in combination with the division of labour. With this approach, task decomposition and coherence of execution are treated as each other's complement.

The structure of operations focuses on establishing the flow of passive elements in a supply chain; the structure is rather neutral as to how the actors are mutually related. As a consequence, *resource assignment*, and also the point of *autonomy* for the actors, remain open. Thus the structure of operations can be combined with any organisational model. Nevertheless, resources need to be assigned to operations and a further explanation of resource assignment is given in Subsection 3.2.3.

In a *silken thread* chain no actor knows or dominates the whole chain. A survey of the chain will provide a framework of conditions and an overview of the chain. However, the difficult point is to clarify which operations are executed by the actors. As we have seen in the mail connection Amsterdam–Hamburg, actors have an interest in being closed and masking their activities.

For an *authority-based* organisation the set of operations is given. That is to say, the actors and their performance are well ordered in the hierarchical structure. The hierarchical structure determines the set of operations. The framework of conditions provides an explicit overview of the supply chain. Thus, a comparison of both shows whether the present organisation suffices to maintain the supply chain.

In a *contract net* the task manager determines the set of operations and thus task decomposition. However, the contract net has no inherent prescript for achieving coherence. The cities in the Low Countries attempted to change the silken threads in the mail services into a contract net. The decomposition was at hand, but it was rather cumbersome and it took several years to establish and agree on a framework of conditions; this was not only due to the fact that an overall authority was missing.

In general, implementing a *structure of operations* is a non-trivial exercise. Task decomposition results from practice and habit. It has to be guided by the presumed competence of the actors. In warfare the commander-in-chief (the authority) decides on a plan or framework, but as we have seen, that does not mean that a mere proclamation makes it work. Having a set of operations, the framework of conditions provides a prescription for achieving coherence and generates a transparent description of the supply chain. The Low Countries mail exchange network showed that, within the framework, the work could be carried out by whatever actor with the appropriate resource available. This explains the reluctance of actors in a silken thread to be open.

3.2.3 Resource Management

The performance of a supply chain depends on the resources and on whether they will be available at the time they are needed. The passive elements will flow most smoothly if the resources are available on request. The actors participating in a supply chain are supposed to provide and arrange the resources, one way or an-

other they have to manage resources. *Resource management* concerns application and maintenance of the resources. The structure of operations simply assumes that actors and their resources are (at the proper time) available. But there is quite a difference between managing resources and maintaining the goods flow. We call the latter activity *chain management*, to oppose it with resource management. The role of *chain manager* is present in the contract net: the task manager makes a start for managing the flow by ordering tasks. However, the role of the task manager is limited to preparing the bidding process, but as we observed in Chapter 2, this role must be extended to the operational phase.

The resource manager is responsible for the resources. He takes care of his resources. The resource manager is acquainted with how the resources are best applied to carry out the tasks assigned to them. He knows the niche of the active resources. He knows the peculiarities of the workload of a resource, knows whether resources are available, or that they are overloaded with activities. The resource manager is the one who can really choose among available resources. A prerequisite is, of course, that he has some "authority" to decide whether a resource will be used or not.

Combining chain management and resource management is a point that does not have a single once-and-forever solution. From one case to another, one has to find an appropriate balance between them. Society determines the overall aim of a supply chain, and values its products and services. The value attached to the chain's products is important for determining a balance between resource management and chain management. We distinguish three attitudes towards resource management, we call them: *disposition management*, *workload management* and *service management*.

Disposition Management

A chain might be such that the *availability of its services* is of primary importance for society. In general, public services have their primary focus on ensuring availability of the services, which directly implies that resources need to be available. Thus resource management aims to have the resources at its *disposition*.

The warfare domain provides clear examples. Real fighting (happily) occurs seldom and for relatively short periods. Nevertheless, the availability of the resources at the appropriate time is decisive. Armies are provided with penalties (of which the utmost is the death penalty) to consolidate availability of the resources. The French La Poste was organised on a similar principle. La Poste was a government mail service, the postillons were civil servants. Moreover it was prescribed how many horses needed to be kept available at the livery stables. Putting the resources under central authority guaranteed the disposition of the required resources.

Disposition management is appropriate for chains where the “gain” of a single hit overrules all (previously made) costs. This holds true for warfare, but also in public order (police) or medical aid. Small scale examples are found in contracts for servicing a technical installation. For instance, for elevators or computer networks disposition of (emergency) personnel is guaranteed 24 hours a day.

Workload Management

Not all chains are valued to the extent that disposition management can be afforded. In these chains another combination between resource management and maintaining the flow has to be found.

The involvement of actors in a chain might be guided by a system in which actors are rewarded for the *period* for which the resources are active: a *wages and salary* reward system. Guided by this system, the resource manager will try to bundle operations from different chains to obtain full workloads for his resources. Here we meet a typical *workload management* attitude. Workload management gives priority to the resources, the resource manager determines the time points when resources are available to serve. This might frustrate a smooth and fast flow in a supply chain.

A practical example is that of a transport company sampling freight to obtain a full load. (One might as well think of a medical specialist who schedules appointments with patients so that the waiting room is always occupied.)

Service Management

In a supply chain, effectiveness is determined by the operational output: do sufficient passive elements flow quickly in the proper direction. This is focused on in *service management*, and is directed by a *service rewarding* principle. In a service contract system, the reward for the actors is related to how many elements flowing through the supply chain are served by them, that is, how effective the service of the actor is.

In its most basic form, rewards for services appear in trade. In Chapter 1 we discussed the silk route. Merchants are rewarded for service, they offer their merchandise at several places, and are rewarded only if the service is appreciated.

We illustrate the difference between workload management and service management with one more example. Imagine a booking office for meeting rooms. You are applying to book a room. The officer looks at the schedule and tells you that all rooms are booked. Next time you apply at another officer, he studies the schedule and tells you that all rooms are booked, but asks you to wait a second. He makes several calls and then tells you that half an hour later than requested you can have a room. The first officer aims just for a full schedule (workload) for the

rooms as one would do in workload management. The second officer tries to serve as many operations as is possible, conform the service management principle.

We have sketched different attitudes towards resource management and discussed how they combine with chain management. The point cannot be settled a priori. There remains some tension, but that is not to say that they are contradictory interests. The structure of operations of the chain manager filters out the regularities in the supply chain. In turn, this provides the insight on which to base efficiency in applying resources. The resource manager can judge efficiency, i.e., whether a certain task could have been accomplished with less (use of) resources. Moreover, actors can become specialists thanks to frequent recurrences of similar operations.

The chain manager as well as the resource manager have to weigh the "costs" of keeping resources available against the "profit" expected from eventually carrying out the task. We will not explore resource management further. The important point is that resource management is adequately taken care of, i.e., that ultimately tasks are assigned to resources and that the resources are available at the scheduled times.

Actors are autonomous in how they execute an operation and how they assign resources. Assigning an operation to an actor, means that the actor has to accept the start-state as his input and that he is expected to "return" the end-state. Penders [1994c] calls the start-state the *to-accept* state and the end-state the *transfer-to* state. This terminology is in particular appealing in a contract net, where goods are transferred from one party to another within the conditions of the contract. The terms to-accept states and transfer-to states indicate transition of responsibilities for an operation.

3.3 Towards a Formalism

We have introduced the structure of operations to analyse the fabric of a supply chain. Moreover, the role of the chain manager has been opposed to that of the resource manager. It will be clear that it is the interest of the chain manager to obtain a clear view of the supply chain. In this section we introduce a formalism to describe the structure of operations.

3.3.1 Passive Elements

We model a supply chain by a set of operations and a framework of conditions. The general descriptive scheme of an operation is:

- (i) move object *Ob* from location *p* to location *q*, at time *t* the *Ob* is at location *p*, at time *t+n* it must be at location *q*.

The framework of conditions consists of a concatenation of operations, described as in (i). To comprehend an operation and a framework of conditions it needs to be clear what *Ob* is. Of course, the *Ob* in (i) is a passive element.

In the warfare domain *Ob* is an army, a brigade, battalion et cetera. However, the operation *move Ob* is not transparent. It involves the moving of the army, the soldiers, as well as the supplies. Moreover, while the army moves the supplies are consumed, thus it does not suffice to describe an operation only by start-states and end-states. Therefore, we split the chain into *two layers*. One layer represents the chain of movement of the armies, in which layer the armies are the passive elements. In the second layer the supplies are the passive elements, and might be moved by the army itself or be provided from other sources. Supplies are consumed by executing operations of the first layer. We have discussed (in Subsection 2.3.1) that the campaigns of the large seventeenth century armies were flights forward. These armies moved between sources of food supplies. Though we will not further elaborate on warfare, our method is also applicable on warfare.

In the mail services *Ob* is a parcel or letter, and from the description of the start and end-states it is evident in what the operation consists. However, in the mail services domain the above mentioned layers can be recognised as well. In the primary layer letters flow through. But, the couriers and horses need supplies as well. In general, the food supplies and so forth are provided either by local communities or by a distinct network. For instance, in the La Poste mail service, the second layer consisted of a system of livery stables which was provided and maintained by the government. As the layers are quite well distinguished in mail services, we can abstract from the second layer. We continue our discussion with supply chains in which – similar to the mail services – the passive elements are not consumed (by the chain).

3.3.2 A Formal Language

We apply a formalism to obtain a clear description of start-states and end-states. A formal approach called *situated automata theory* is introduced in [Rosenschein, 1985; Rosenschein and Kaelbling, 1986 and 1995]. In this approach a robot and its environment are conceived of as a (sometimes complex) automaton, analogous to a Mealy automaton [Lewis and Papadimitiou, 1981]. The transitions of the automaton represent how the robot and the environment change from one state to another.

We will extend situated automata theory and apply it to the large scale problem of ordering a supply chain. It will clarify how to describe operations and compile a framework with them. To explain the approach we start concentrating on how situated automata theory applies to a robot.

The automaton, $\langle S, V, N, \text{init}, \gamma, \alpha \rangle$ consists of:

- S : the set of states,
- V : the set of outputs (percepts for the agent),
- N : set of actions (operations),
- init : initial state,
- γ : $S \times N \times S$, is a relation where $\gamma(s_1, v, s_2)$ holds if it is possible for the world to make a transition from state s_1 to state s_2 when action v is generated by the agent,
- α : $S \rightarrow V$, a function mapping S to V (stimuli or sensor data).

The robot itself is part of the context (or environment) and its states are included in the states of the environment (S). In particular this is true for the robot's inputs (sensor stimuli) and its outputs (actions). For details, see Rosenschein and Kaelbling [1986 and 1995]. The actions transfer the robot and the environment from one state (of which the inputs are part) to another, in which the actions have (partly) been executed.

On the basis of this formalism several robot programming languages are developed. A direct application is shown in [Rosenschein and Kaelbling, 1986]. Independently, but in the same spirit, the language described in [Lyons and Arbib, 1989] was developed. In [Penders, 1989] the rather opaque operation of error recovery is analysed in combination with an architecture for a complex robot system.

In situated automata theory the actions of the robot are directly related to the description of the state in which they are applied. This feature is particularly useful in a description of a supply chain. It forces one to define which steps should be made in the supply process, and simultaneously which resources are needed. However, in order to apply situated automata theory to the art of moving objects, we have to reinterpret its symbols. This is done below.

S : the set of states, contains the start and the end states. The state-of-affairs at the start and the end of the operation has to be described. However, it is not necessary that this description is *complete*. Complete means that it describes all there is to be described about the state of affairs. The description of the start- and end-state should be detailed with respect to the passive elements: their *time* and *place* coordinates. Besides, some of their transport characteristics (as far as relevant for forwarding them) may be described. Similar to scheme (i) in Subsection 3.3.1, it includes answers to such questions as: what is the object *Ob*? where is it now? where has it to be moved to? and when is it supposed to be there?

V : the set of outputs, is conceived of as a set of characteristics of start-states which are inputs to the robot. When the characteristics $v \in V$ are observed in a state, the operation can be executed.

N : the set of actions, it is evident that this set coincides with the set of operations.

$\gamma : S \times N \times S$, consists of all valid transitions from a start-state to an end-state. It has to cover all operations available for the process.

$\alpha : S \rightarrow V$, represents the sensors of a robot or the possible observations an agent (and actor) can make of a certain state. The robot has certain sensors with which it observes the state s , this is what function α captures. Certain observations trigger an action γ . In this chapter we do not describe the function α . At this point there is a considerable difference between robots and natural beings or organisations of human beings. Sensors for robots will be the main topic discussed in the next chapter (Chapter 4); in Chapter 6 we will discuss human perception.

In order to analyse the framework of conditions more elaborately, one can set up further formal definitions (definitional prescription) of γ . The definitions have to delineate which passive elements from the start-state have to be moved, and what their position will be in the end-state. These definitions are useful for setting up a computer simulation of the system. Such a computer simulation and operation program for a supply chain are described in [Schreiber, 1996]. Whether the actual outputs will exactly reflect the simulated outputs can never be guaranteed: it is a risk inherent to performing any action. This is the practical stance we will take.

3.3.3 Decomposing and Combining Operations

Below, we set up a framework by describing start- and end-states. The implementation of an operation remains open. We do not describe how an actor has to execute an operation. More strongly stated, we explicitly avoid this and leave resource management open, as discussed in 3.2.3.

Decomposing Operations

An operation might be rather extended. In order to get a clearer view of it we can decompose an operation into (a series of) *consecutive* sub-operations and/or *concurrently* executable sub-operations. A decomposition into consecutive sub-operations requires that one is able to identify intermediate states which separate the sub-operations. In Figure 3.3, operation 1 is split into three consecutive suboperations, 1.1, 1.2 and 1.3. This is necessary to preserve coherence in the framework of conditions. (In this way a context is also created for the autonomy of the actors in executing the operation.)

Figure 3.3: Decomposing an operation into consecutive sub-operations.

To decompose an operation into *concurrently* executable sub-operations, the start-state and end-state need to be separable into disjoint sub-states. This means in particular that there must be a partition of the set of passive elements corresponding to the decomposition. Two sub-operations cannot share a passive element, as this violates the requirement for independent execution. In Figure 3.4 a set of six passive elements is divided over two concurrent operations. In case a, the operations 1.1 and 1.2 process distinct sets, case b shows a violation, the same passive element (number 4) has to be processed by both (concurrent) operations.

Figure 3.4: Concurrent decomposition.

In general, it will be obvious how to divide the passive elements over the operations. We have distinguished resource management from task management in Subsection 3.2.3. By doing so we have bypassed a tricky problem, namely, the concurrent decomposition of an operation in which (by accident) the same resource is shared by two sub-operations.

As an example we consider an army. In order to speed up moving, it is decided that the army will be moved by train. The soldiers and their arms have to be moved, as well as the supplies. As the army occupies all the train wagons, no resources are left for moving the supplies.

There does not exist a systematic approach to avoid such (unforeseen) sharing of resources. Nor is it always clear whether execution of one operation might interfere (unforeseen) with some other. In the AI literature this type of conflict is headed under the so-called *frame problem* [Brown, 1987]. Simply stated it is the problem that the execution of one action has effects in the context/background which are not explicitly taken into account. A real problem arises when a certain action does interfere with some other actions. The huge amount of literature on the subject indicates the difficulty of dealing with it, for a brief explanation and overview we refer to [Toth, 1995]. In the practice of moving physical objects such conflicts soon become clear. Moreover, before execution time, the chain manager, the one who starts the decomposition of the operation, has to find consensus with the resource manager. In this interaction the problems pop up, and have to be dealt with.

Combining Operations

We now turn to composing new operations out of existing ones. By combining and arranging operations we can compose new chains. We give a straightforward example in the mail service domain. Assume we have available: an actor to pick up a letter at p and bring it to r ; another actor has to pick up a letter at q and bring it to r too, a third one should do so from r to s , and a fourth one from s to t , one from s to u . By combining the operations we obtain a small network, see Figure 3.5.

Figure 3.5: Some operations of actors in a mail service.

To treat the composition approach systematically, assume we have two operations x and y . As before, operations are marked by a start- and an end-state, which occur at certain points in time. The interval between the time points is needed for execution. First, we look at the possible relationships between the two operations. For this we reinterpret a schema given by Allen [1984] for time intervals. There are thirteen possible relationships. In Figure 3.6, seven of them are drawn, and six others can be generated by inversion of x and y . In the right column our reinterpretations are given.

Figure 3.6: Overview of relationships among operations, seven relationships are drawn, six others generate by inversion of x and y.

In the *art of moving objects* the spatial co-ordinates of the objects (goods) are as important as the temporal ones. We interpret (and here we deviate considerably from Allen) the start-state and end-state as (partial) *states-of-affairs* that occur at certain places (that is, at a certain small geographical region). In our structure of operations we can deal only with start- and end-states. From our point of view, only those relationships are interesting in which two states meet in time as well as in place. Thus, the relationships applicable are *x equals y*, and *x meets y*; they are

listed in the right-hand column of Figure 3.6, we encountered these relationships in the decompositions (Figure 3.3 and Figure 3.4).

Composing a structure of operations typically consists of finding *meets* to obtain a series of operations. Resource management is served by finding *equals*, because when start- and end-state meet, the resource manager can look whether the operation can be carried out by one resource only. In a structure of operations, the *x starts y* and *x finishes y* relationships are useful as well. In the remaining relationships, *before*, *overlaps* and *during*, no relationships among start- and end-states can be found. Thus we cannot order them into any chain. When, in the process of composing a structure of operations a *during* relationship is encountered, it may be worth to look again at the decomposition, in order to identify intermediate states.

3.3.4 Composing a Chain

With the four relationships, selected in Figure 3.6, we can structure the mail service example, see Figure 3.7. By combining the operations, we achieve a small but rather rich structure. We have services from *p* to *t* and from *q* to *r*. But also from *p* to *u* and from *q* to *t*; moreover all intermediates, as from *q* to *s*; *r* to *u*, and so on.

Figure 3.7: A typical composition for a structure of operations.

We have discussed prescriptions for composing and decomposing a supply chain. In redesigning a chain we alternately compose and decompose in an iterative process. The main requirement is that: in the start-state of any operation all passive elements have to be available. The states are specified by time and place coordinates. As a consequence we can expect that at the end-state all outputs are available.

The specifications ensure that operations are detached and thus every actor can execute its operations independent of others. The requirement of autonomy in execution forces clarity in describing an operation. The main advantage in aiming at

independence and autonomy of actors, is that once a decomposition is made and the actors are available one can compose a new structure of operations. The new structure does not necessarily have to aim for the same goals as the original one.

By determining a structure of operations, that is a set of operations and a framework of conditions, we obtain all prerequisites for planning. In the section on warfare we cited Jomini, where he describes how Napoleon used to plan his (successful) campaigns. In the AI literature such type of planning is called *Action planning* or *state-based planning* [Pelavin and Allen, 1986]. It consists of scheduling the execution of activities as in our structure of conditions. Figure 3.8 shows the principle of action planning. The planning consists of matching the post-conditions or end-state which result from action n to the preconditions or start-state of action $n+1$. In Section 3.5.1 we further discuss planning approaches in AI.

Figure 3.8: Action Planning is to match end-state n with start-state $n+1$.

3.4 A pan-European parcel distribution

In the previous section we discussed the structure of operations and its formal description. In this section we apply the structure of operations to the practical case of a pan-European distribution service. The aim is to create an open service network. It involves local service suppliers of piece goods.¹⁴ In this service, small volume goods (parcels) are sent over long distances. For reasons of economy, the parcels have to pass through the hands of many actors. The actors are willing to co-operate with one another, but somehow the combined services do not work well. The background for the service is well-ordered and set.

¹⁴ For further description of the business case from which the study originated, we refer to [Penders, 1995a; 1995b], and the series of IDEAS deliverables mentioned there.

3.4.1 The practical problem

The practical problem to be solved concerns how to *connect* the independent services and to *embed* them into an *overall service framework*. The case is analogous to the mail service example of the Low Countries. From the start many actors are available, each of them is fully equipped and experienced to perform. However, each one also has his own view on what exactly his service covers. The supply chain has to be directed so that the contributions of each actor indeed add to the supply process; in fact *coherence* has to be sought for. It is by and large a case where the chain has to be led through the niches in which each of the actors is already performing.

The actors are not subsumed under a common authority. Thus, when a request for a pan-European service arrives, a mechanism similar to the *contract net* has to come into being. However, as discussed above, the contract net provides no procedure neither for task decomposition, nor for guiding coherence. Incidentally, it might happen that the services offered by different suppliers fit at once. But experience teaches that mostly they do not. In fact, what has to be developed is less formal than a recipe for co-operation. Each actor can act as before, only a framework for co-ordination is needed. However, it is important that all parties agree on it. Mutual agreement moreover, enables each of them to take the framework off the shelf whenever a request for service arrives.

Transport Units

The passive elements flowing through the network are piece goods. Important for the physical handling of the goods is their volume. To start, we define a hierarchy of goods, the order of which roughly corresponds to the volumes. We distinguish subsequently: ship/train loads, truckloads, pallet loads and parcels, see Figure 3.9. For convenience, and when no confusion is expected, the word 'load' may be skipped. The hierarchy of goods levels anticipates the type of resources needed; the levelling corresponds to the type of (handling) equipment needed in the process of forwarding. Our hierarchy is a variation on the one given by Damen [1994]. He defines a hierarchy of *Transport Units* used in mail services. This hierarchy goes from trains, trucks, roller containers, bags, bundles down to letters.

Ship/trainloads
truckloads
pallet loads
parcels

Figure 3.9: *Transport units.*

In practice it is not a particularly new problem to transport a full truckload from any point in Europe to any other point, or ships and trains from any station/harbour to any other. Moreover, it is no problem to find actors which can per-

form the operations. Airlines and railway companies enjoy (inter) national reputations, and the abundance of trucking companies is public knowledge. Moreover, even on the lower levels (pallets and parcels), Europe is in fact covered by a large number of service suppliers, who offer public collection and delivery services. Most of them, however, serve only local areas.

Using the hierarchy of transport units, the problem of the combined piece goods service can be further specified. The first part of the practical problem is to *connect* services on the same level, so as to achieve a wider (European) scale. In this sense it is a problem similar to the one we have discussed in the case of the mail services. The second part of the problem is to *integrate* services on different levels. The latter arises when collection (into larger transport units) or distribution of cargo is necessary, for instance, when collecting truckloads to load trains or ships. Or, on a lower level, when collecting or distributing pallets and parcels into (out of) truckloads. In other words, where dedicated transport of a single piece or transport unit is too expensive, combinations have to be sought to increase the volume and share the costs.

3.4.2 Available Operations

The problem of the pan-European parcel distribution is to connect networks on the same level of transport units and to integrate networks on different levels. Thus, we survey what actors and operations are available, and we discuss how to compose a structure of operations out of them.

Service Suppliers

The type of actors we are dealing with are professional service suppliers; they are *resource managers*. What they offer is performance of a certain operation. A service is nothing but an offer to execute an operation. The negotiable part of a service offer concerns the conditions on when/where to accept the goods and when/where they will be delivered. The service supplier – generally speaking – does not tolerate any interference with how he executes the service. Moreover, in general, the client is not really interested in how the service supplier does it, the client's interest is only that it will be done. We conclude that from a practical view, it is quite natural to preserve the autonomy of the actors in a structure of operations.

Collection and Distribution

Moving the passive elements from one place to another is the original subject of logistics; the elementary operation is that of transport. Single and incidental transport operations can be executed independently. In cases where the loads (and fees) are relatively small with respect to the capacity (and costs) of the transport means, collection and distribution structures emerge. Mail services deal with an analogous

problem, as we have seen in Chapter 2. The mail service domain also shows that when the network is expanding, maintenance of the links between operations becomes an important item. Rules and procedures for exchange between operations are required. Thus it makes sense to separate *collection* and *distribution* operations.

Figure 3.10: Collection and distribution.

Collection and distribution operations can be described in our format of operations, see Figure 3.10. The figure is basically the same as Figure 2-4 which shows collection and distribution in the mail services. The essence is that all goods arrive at a given place at a particular point in time. Collection operations *finish* at this point. Distribution is the dual of collection, it *starts* at a given place and point in time.

Sorting

Collection and distribution have to be connected; in modern terms the connecting unit is called a *hub*. A hub is always surrounded by other operations, its function is to divide passive elements over the subsequent operations, as Figure 3.11 indicates.

Figure 3.11: A hub.

With increasing volumes (quantity of pieces) within a service network, the connection between collection and distribution requires more attention. If one courier has to hand over one parcel to a second courier, there is not much of a problem, the exchange can be overviewed and the required handling is straightforward. However, when several couriers have to exchange many parcels amongst themselves, there is a *sorting problem* to be solved.

Sorting is typically an operation which creates a different order of the passive elements (goods). At a hub, piece goods arrive from several origins. At certain gates, usually called in-docks, goods are received. From certain other gates, out-docks, the goods are despatched. In practice, the in-docks and out-docks are often literally at other sides of the building. The start- and end-states of the sorting operation are identified with the situations at the in- and out-docks. The sorting is nothing but changing the start-state at the in-docks into the end-state at the out-docks, see Figure 3.12. The configuration of the end-state is determined by how the goods will flow further through the network. Sorting is dividing the goods over the subsequent operations in the network. Thus, sorting is a major constituent of a connected network, and has to be included in the structure of operations. The major sorting rules or prescriptions derive straight away from the framework of conditions.

Figure 3.12: Sorting at a hub.

Transshipment

In a service with a local range, collection and distribution operations *meet* at the hub. For the pan-European service, long-distance transport has to bridge the geographical distance between collection and distribution. In some cases it might be

straightforward. However, if the volume to be transported is too small, a second collection step might be required to obtain a higher level transport unit. Likewise, the structure at the other side depends on how long the volume can be kept together. When the whole distance cannot be covered by long-distance transport, stepwise distribution might be needed.

The higher levels of passive elements, pallet-loads, truckloads et cetera, require dedicated handling equipment to exchange the transport units among the transport means, for instance, to pick a container from a truck and to place it on a train wagon, or to free a particular wagon from within a full train. Here a *transshipment problem* has to be solved. When connecting networks on different levels of transport units, a hub also plays a role, by offering the handling equipment.

Buffering

As stated before, for long-distance transport, goods have to be collected so as to combine them into higher order (greater volume) transport units. One can achieve this by a careful planning of the collection activities. In a collection system, all transports have to end at a certain place before a certain time. The planning has to ensure that all actors deliver promptly on schedule, so that subsequently the goods can be forwarded immediately. Another option, which relaxes this constraint, is to store goods until the proper volume is achieved; in this case one can speak of a *buffer*.

A buffer function is often integrated with a hub. The buffer capacity is necessary to collect the incoming stream of smaller units to form the higher level units, and to split the higher units into smaller ones. A sorting activity is often necessary for both streams. As soon as goods arrive, they are sorted and placed in buffers corresponding to the outlets of the hub (this is often called *pre-sorting* in mail applications or cross docking in warehouse environments).

3.4.3 Rearranging the Chain

A networks consists of a connection of operations. However, connecting operations working on different levels of transport units, is not straightforward. When simply inputting the output of one network into another, double work is be done. For instance, a pallet network consists of operations on pallets and cannot forward individual parcels. In a parcel network on the other hand, only parcel can be processed and all pallets must be unloaded. Thus the levelling should be explicit.

In Figure 3.13 the principle of a levelled network is given, higher levels are nested within lower ones. The figure shows that four types of hubs are necessary. In the figure starting from the left side, the first pair of hubs combine smaller loads into bigger ones: parcels into pallets and pallets into truckloads. At these hubs pre-

sorting is done. Then, after transport has taken place, the second pair decomposes truckloads into pallets, and pallets into parcels, respectively.

Figure 3.13: A framework in which transport is nested within collection and distribution.

Figure 3.14 follows the route of an individual parcel through the network as given in Figure 3.13. The parcel undergoes the operations: parcel collection, pallet collection, transport, pallet distribution, and parcel distribution. The parcel climbs and descends through the levels of transport units.

Figure 3.14: The route of an individual parcel.

Which configuration of network connections and in particular which hubs to choose depends on practical data. Practical facts determine also whether the route shown in Figure 3.14 is the most adequate and effective. In practice, it turns out that *transport* itself can hardly be optimised. The preparatory, the intermediate, and the completing handling offer the opportunity to optimise the transport. (Often these activities make up the major costs in the network as well.) Thus, the hub and the operations executed there determine the further flow of the goods. The opportunities for optimising the collection and distribution lie not in transportation, but in the activities and planning that surround and direct the transport, that is, in particular in the hubs. Thus, though transport is the elementary operation, it is the organisation of the network that is the main objective of study. This implies that all operations have to be studied in the context of their mutual connections.

A hub derives its character from its place in the network. We have discussed three types of functions for a hub: *sorting*, *transshipment* and *buffering*. Sorting is reordering the goods, transshipment providing the resources for, and performing exchange between carriers, and buffering bridges time gaps between consecutive operations. Operations, though independently executable, will depend on one another in their start-states and end-states.

3.4.4 Redesign of Services

We are developing a structure of operations for the practical case of a European parcel service. In the previous subsections we have distinguished levels of transport units and indicated how (in principle) to connect the different levels. In this section we discuss which types of actors to select for our case. Where possible we will use existing services. A survey of available actors and services on each of the levels of transport units is required.

Logical Service Supplier

On each level of transport units we set up the role of chain manager. He will offer the goods to the network and receive them after processing. The chain manager is first an *integrator* who connects operations, provided by actors, into a service, and ensures that end-states match start-states. This role requires knowledge of the particular states and about the freight itself. He is called a *Logical Service Supplier* (LSS); he clearly is a service supplier. The term logical stresses that his role is that of connecting operations independent of whether or not he executes any operation himself.¹⁵ The typical role of an LSS is to combine cargo of different origins so as to build higher transport units.

¹⁵ It is interesting to note that our notion of LSS has much in common with the notion of Temporary Supply Chain Co-ordinator developed independently by [Hoogeweegen, 1997].

Long Distance Transport

In the parcel service we need long distance transport. There are specialised actors with air, water, rail and road connections available; each of these modalities has its pros and cons.

Air connections are expensive. By *road*, nearly all places can be connected, but over long distances, road transport (trucks) is slower than air and rail transport. Moreover road transport is less predictable and is certainly not cheap. The operations of *railways* are restricted by the trajectories of railways, the locations of stations for loading/unloading and the time schedule on the railways. With railway and road connections the whole of Europe is within reach. *Water transport* is the cheapest, however water ways are geographically less uniformly spread, moreover loading/unloading operations (harbours) need special care.

For our service, a combination of railway and road transport services seems to be the minimum required. Thus the backbone of our service will be a long-distance transport network, or better expressed, a transport schedule on the levels of ships/trains and trucks. We leave the execution of these operations to the specialised actors.

Pallet Service

As stated above, we have to combine services on each level. We have chosen actors with rail and road connections. The operations of the railway company are restricted, because rail wagons can only be transported (within a fixed schedule) among the railway stations. Several areas are serviced by road hauliers.

An LSS has to combine the operations into a service chain. The LSS is distinct (or distinguishable) from the railways or road hauliers. He hires railway as well as road transport, and organises the interchange between rail and road transport. The LSS works on the pallet level and combines shipments (pallets) into (nearly) full truck- or wagon loads. For convenience, we will call this LSS the *Pallet Service*.

Transshipment between rail and road transport takes place at *transshipment and sorting* hubs located at the most important railway stations. (This is a consequence of selecting railways.) These hubs are bi-directional: transshipment from trains to trucks and vice versa. Moreover the Pallet Service incorporates (service) *sale-point hubs*, where clients can deliver or collect shipments. The sale-point hubs thus serve as inlets and outlets of the network, and as internal nodes. These hubs can also be the starting points of house delivery tours. The LSS operates the transshipment and sorting hubs and some of the sale-point hubs. The peripheral sale point hubs might be operated by third parties (for instance road hauliers). Figure 3.15 gives an overview of the Pallet Service. As an LSS, the Pallet Service is involved with all operations. However, only the operations outlined in black are

really within its control. Going from left to right in Figure 3.15, The Pallet Service operates some of the sale-point hubs where parcels are stacked onto pallets, it operates all transport to the hubs at the railway station, where pallets are grouped into train loads. After train transport, the Pallet Service degroups the pallets and transports them to the sale-point hubs, some of which are operated by a third party.

Figure 3.15: The Pallet Service.

Parcel Service

The aim is to create an open service network for pan-European parcel distribution. On the parcel level, several service suppliers might be active. A typical example is an LSS who delivers spare parts. The service starts within public warehouses. In the warehouses spare parts are stored. The spare parts have to be delivered at "private" addresses (repair centres) all over Europe. Typically, the volume and the sizes of drops (a delivery at one stop) is limited to one or a few parcels only per addressee. Moreover, delivery is not regular but incidental. A single addressee seldom receives parcels (he seldom orders spare parts).

It is clear that the LSS on the parcel level, that is, the Parcel Service, has to find partners to set up its service of spare parts delivery. The level of the transport units and the volumes are too low to go immediately into business with the railways or road hauliers. The Pallet Service, described above, is an appropriate partner. As we discussed above, at its hubs the Pallet Service can take in single pallets and deliver them to other hubs. All intermediate operations can be left to the Pallet Service. What is missing is the collection of parcels to form a pallet load and the final delivery of parcels, which the Parcel Service has to organise itself. Thus, the Parcel Service is built on the framework of operations of the Pallet Service, see Figure 3.16.

Figure 3.16: The Parcel Service.

Variations in Services

The main long distance transport for the Parcel Service, is outsourced to the Pallet Service. The Pallet Service outsources it to the railways and the road hauliers. The Pallet Service acts as a sub-LSS, which means that from the point where the parcels are plugged into the Pallet network till the final hub, the Parcel Service cannot interfere and has to rely on the Pallet Service. The Pallet Service in turn relies on the railways and road hauliers.

Figure 3.14 shows how a parcel with spare parts flows through the levels of the network. Note that only the Parcel Service knows that spare parts are flowing through. The railways, for instance, are not really interested in the fact that they participate in a parcel distribution service, very likely that they neither know nor want to know this at all.

The Parcel Service has to select from the operations offered by its partners, for instance the Pallet Service, just those that suit his structure of operations. His service partners have their main business in other flows of goods. For them, the

Parcel Service is probably one among many clients. The Parcel Service just guides its parcel through the appropriate niches offered by the railway and road companies for instance.

However, the fact that the Parcel Service is one among many anonymous clients limits the possibilities of dedicated services. Each actor offers to execute certain operations, but does not allow many changes to them. The connection of services simply consists of aligning the points of exchange. The services which the distinct parties offer are delineated by describing input (location of the start, time of start, goods to be forwarded) and outputs (location and time of termination). These are the "variable" parameters, in general they are explicitly stated. Flexibility, for instance, for the Parcel Service is to be found in a "clever" planning with the explicit variables of the other actors.

3.4.5 A Hub in the Service

Within a service, the actual functions of a hub are: receipt of goods at some dock, storage for short periods (buffers), and delivery from some appointed dock. In the *actual process we can distinguish between information processing and physical handling*. Receipt and delivery are mainly administrative actions. Having an administrative overview of the goods, the physical re-ordering of the goods can be planned. Thus from here, instructions can go out to the shop floor as to which goods have to be moved (sorted).

Because of its central position, a hub has the potential to influence and direct the flows in the network. It can be a point where information is collected about the current flows in the network. On this basis one can decide to re-direct the flow of goods. The alternative routes can be found in the structure of operations.

When using services of third parties, control over a hub can be rather profitable. As an example, we consider again the Parcel Service. For different customers the Parcel Service operates several warehouses, say one in Amsterdam, another in Brussels. (A warehouse can incorporate hub functions.) It frequently happens that spare parts have to be delivered from these warehouses to (nearly) the same destination, say to Sicily. The ideal of the Parcel Service, and money can be saved by this, is to combine delivery of both products at the appropriate time in one delivery tour.

The Parcel Service has to anticipate the process flow in the network. The parcels are inserted into the network at such points in time that they will appear at the same time at the end of the channel. It is clear that the main prerequisite is that the structure of operations is transparent and that the performance of the partners is reliable. In fact, the Parcel Service makes money by only directing goods through the (for his purposes) appropriate niches.

3.5 To Conclude with the Supply Chain

In this chapter we have explained the formal aspects of our structure of operations. We have applied this concept to a practical case. Below we relate our approach to the subjects of *flexibility* and *planning*. Moreover we oppose our approach to what is called *functional decomposition*. At the end of this section we recall the main points and draw some conclusions.

3.5.1 Other approaches

Flexibility

We have analysed the Parcel Service. It uses only a few of the operations which each of the actors can offer. The structure of operations delineates what is expected of each actor, in order to make all performed operations fit into a useful chain (coherence). For performing the sub-operations, it is not necessary, that the agents know about the overall operation in which they participate

It is important to note that in defining the Parcel Service, we did not introduce new operations in the Pallet Service. All planning was done using standard operations of the other service providers. Flexibility can be achieved, not by inventing new and unknown operations, but in (differently) planning the pre- and post-conditions of the pre-defined operations. To underline this, we provide just one more example.

A consul in foreign service is to open an exhibition. As a surprise he wants to present a case of fresh apples from the home country. He calls the national airlines, and they assure him that they can deliver the case with the next flight. To be sure everything will work properly, the case is sealed all over and prompted with the text: "special treatment". The case is delivered together with an ordinary box with documents at the appropriate time at the airport. The employees put the case apart to ensure special treatment. The box with documents arrives with the first flight, but no case with apples. The consul makes a telephone call to the home airport. "Yes, yes, the case arrived, it is standing here besides me and it will be given special treatment." But the case arrived only two days later!

Planning

The structure of operations provides a basis for planning services. In Artificial Intelligence (AI), the type of planning which we have been discussing is called *state-based* planning. In the literature several automated planners are described. A classical example of a (state based) planner system is STRIPS. The essence of it is that it represents actions with slots. For references we refer to [Ramsay and Barrett, 1987] or [Penders, 1988]. The inference machine (as the core mechanism is

called) starts with the desired result, and looks for actions that have the result as output, given in the post-condition. Each schema has stated several pre-conditions. The machine continues by searching for schemas that will satisfy these pre-conditions until the start-state is reached.

In our treatment of a supply chain, we have distinguished the passive elements, the things that flow through, from the actors who perform operations using resources. A structure of operations consists of a framework of conditions (indicating coordinates mainly related to the passive elements) and a set of operations. The schemas of a state based planner are representations of operations. The variables (x,z) represent passive elements. The result of the planning is just a framework of conditions, or a co-ordination framework [Seghrouchni and Haddad, 1996]. Similar to our treatment, state-based planning presupposes a distinction between the co-ordination of actions and resource assignment. State based planning does not concern itself with resource assignment: resources are basically assumed to be available.

Functional Decomposition

In designing a structure of operations, we have abstracted from the particular resources that have to perform the operations. A similar approach is found in the so-called *functional decomposition* approach [Ross, 1977]. In a functional decomposition, one tries to determine what the main functions are, or the main activities applied to the passive elements. The goal of such a decomposition is to obtain a good requirements definition. In [Ross, 1977] and [Ross and Schoman, 1977] the method for functional specification is given. Typically the approach goes through the stages of:

Context analysis, the reasons for which the system is to be created;

Functional specification, what the system is to be;

Design constraints, conditions specifying how the system is to be constructed.

One might classify our operations into classes of functions. The functional decomposition shows all parameters (or parameter slots) by which functions need to be connected. Every function has a set of parameters (open contexts); in implementation, these parameters have to be filled in. A functional decomposition is neutral as to whether these parameters are fixed or negotiable.

A functional decomposition makes clear what is going on in a system, what is passed from one actor to another. Using functional decomposition, Damen [1995] analyses the functions needed in a supply chain. As such it provides a valuable checklist for a designer, but it cannot give directions as to how to design the system. Functions can be put in arbitrary sequences, but this does not lead to a design of a chain, it does not provide coherence.

The main shortcoming of functional decomposition is that functionally there is no argumentation as to why to choose a particular type of decomposition. For example in the days of Napoleon, an army was split into "functional" parts such as cavalry, infantry et cetera. Modern western armies are divided into brigades. Each of them has, on a small scale, all functions the army as a whole has. In a brigade there are still the functional parts like cavalry, infantry et cetera. Functional decomposition cannot discriminate between these two alternatives.

Damen [1994] extensively investigates mail systems using functional decomposition. Above we discussed three examples from the history mail services: a silken thread in the Hamburg–Amsterdam connection, the authority-based French La Poste and the mail service contract net in the Low Countries. These systems incorporated the same functions, that is, to pick up and deliver letters. However, as we saw, the systems were rather different but the functional decomposition was the same.

3.5.2 Summary on the Structure of Operations

A supply chain consists of a concatenation of operations. The *chain manager* keeps the *passive elements* flowing by ordering the operations. An operation is defined by a start and end-state. On the basis of the start and end-states the chain manager can combine operations and compose a *structure of operations*. The structure of operations is the *co-ordination* framework among the partners in a service.

To execute operations, resources and in particular agents are needed. The *resource manager* prepares the resources with which he participates in the supply chain. The management of the resources is kept as much as possible out of any prescript. The (sub-)operations are independently (and autonomously) executable. The operations can be performed by single autonomous actors. The main requirement for decomposing an operation into sub-operations, is that fully separable states are defined. By accurately planning the co-ordinates of actions, complicated operations can be achieved. The structure of operations is worked out for the real case of a pan-European parcel distribution service. The case mainly required a reordering of operations.

In particular commercial actors, distinguish among each other, not in the main operations, but in the particular start and end conditions of a service (or operation). These are just the terms in which the structure of operations is described. In the real case, the chain manager was called Logical Service Supplier (LSS). We encountered many service suppliers, they usually performed the role of resource manager. We have seen that parties can change roles, for certain shipments they are LSS, for other they are only service supplier. The structure of operations is also useful for the negotiations among the LSS, as they have to clarify where exactly the service starts and ends and which requirements will or will not be satisfied.

At execution time, an explicitly defined and mutually agreed structure of operations is rather convenient. It states clearly what partners (can) expect from each other. Moreover, in cases of deviations, it comforts to communicate the consequences of the deviation for the further operations.

3.5.3 Problem Statement 1

In this and the previous chapter, we explored our problem statement 1: *how can the activities of the actors in a (large) supply chain be arranged such that a (goal) directed flow of objects results?* In Chapter 1 it was said that a basic ordering principle forces the flow into the appropriate direction. We conclude that supply chains consists of a web of activities. This was illustrated in Chapter 2 by surveying examples.

In this Chapter we have evaluated organisational models: silken thread, authority-based organisation and the contract net and concluded that neither of them is particularly designed for a supply chain. A supply chain consists of a concatenation of operations, the concatenation is obtained via co-ordination. We introduced a structure of operations, in which the operations, necessary for moving the objects, are separated from their ordering. Task decomposition consists of distinguishing operations, we added the demand that operations are independently executable. Also, an ordering of the operations is required, for this we introduced the framework of conditions which provides the ordering and the co-ordination in a chain. We have worked this out for the case of a pan-European distribution service. The chain manager, also called Logical Service Supplier, has to compose, guard, and maintain the framework of conditions.

We discussed the supply chain from a global point of view, however, we have not yet touched the point of how operations are to be executed. This subject, captured in our problem statement 2, is dealt with in the next chapters.

4 Robots and Motor-Schemas

In Chapter 1 we formulated two problem statements on the art of moving objects. The first one has been dealt with in the previous chapters, and resulted in a convenient arrangement of the links in the supply chain. The second problem statement, which is *“how can a set of resources be made to perform as an actor within a supply chain”*, is the theme of this chapter. This problem concerns the local point of view i.e., how robot-agents perform in a supply chain. First, a design approach for mobile robots is discussed. Then we explore sensor data and the information provided by the data. It turns out that the aims and goals of a robot have to be provided by a higher level, in another frame of reference.

4.1 Limitations in Robot Design

Robots have rather limited and peculiar domains of application, which we have called *niches*. In Chapter 3, a supply chain was built by connecting niches, assuming that we know the niche of each actor. In robotics, the casualness of this assumption comes into question: the design of robots forces one to delineate tasks quite precisely.

In the nineteen seventies, robotics research had a flourishing period. The research aimed at designing technically sophisticated robots, which were successfully applied in rather pre-structured and (highly) controlled industrial environments, for instance in the automotive industries. When transferring the robots into a real life environment problems arose. The robot's internal machinery functioned properly, but it turned out to be difficult to make the robot sensitive and responsive to the environment.

The end of the nineteen eighties showed a divergence in the robotics community. Experimentalists acknowledged the limitations when robotics had to deal with a real environment and took that as the challenge. Researchers restricted themselves to the problem of sensing and focused on implementing low-level reactive sys-

tems. Robustness of low-level systems became the new aim. The divergence is particularly marked by Brooks' [1986] paper which introduced his *subsumption architecture*: an architecture to incorporate low-level systems into a single robot system.

Below we call the experimentalists' approach *sensory-based robotics* (design) after Lyons and Arbib [1989]. It is also known as *behaviour-oriented AI* [Steels, 1994] or *sensory-motor co-ordination* [Pfeifer and Scheier, 1997]. Sensory-based robotics works with so-called *motor-schemas* [Arkin, 1989], they are (routine) sequences for activating a robot's motors. A motor-schema is triggered by sensors, and when triggered, the schema results invariably in an action. The notion originates from the so-called body-schema (and ideomotor notion of [James]) approach in psychology. Early applications to robotics are found in [Culbertson, 1963, 1976]. Motor-schemas were explicitly introduced into robotics by Arkin [1989].

The novelty of sensory-based robotics is its explicit focus on low-level systems. As a design principle, it aims at robust and autonomous motor-schemas. This differs principally from traditional *functional decomposition*. "*Functional decomposition attempts to construct general-purpose functional modules*" [Gat *et al.*, 1994], whereas sensory-based robotics focuses on specific routines.

We discuss sensory-based robotics. The robots are applied on a goods terminal, a type of hub. In Section 4.2 we specify the application, using the structure of operations method described in the previous chapter. We turn to the design of the mobile robot in Section 4.3¹⁶. In Sections 4.4 and 4.5 we discuss sensors and information collection, and introduce the notion of *cognitive* basis. The cognitive basis of a robot might differ considerably from the designer's. This does have consequences for the design approach, as we discuss in Section 4.6.

4.2 Agents on a Cargo Terminal

Before going into the robot design, we specify the intended domain of application. The robot might be applied on a goods terminal (a type of hub) which is part of a supply chain. The role of the terminal is discussed above in Section 3.4. At the terminal piece goods arrive from several origins, they are first sorted and then forwarded.

¹⁶ This section is based on [Bomhof and Penders, 1992].

4.2.1 Sorting

Inside the hub the sorting operation has to be executed. At certain gates of the hub, called in-docks, goods come in. From certain other gates, the out-docks, the goods leave again. In practice, the in-docks and out-docks are often literally at opposite sides of the building. The sorting consists in the changing of the start-state at the in-docks into the end-state at the out-docks, the main action to be taken is that of moving the goods. The sorting task is decomposed by the hub-manager and translated into instructions about which items have to be moved, from where to where. The hub-manager sends the instructions to the individual agents, which are supposed to perform autonomously. After accomplishing a task, the agents notify the hub-manager. Taking all the feedback together, the hub-manager can conclude that the sorting operation is accomplished. We note that the agents need not know anything about the sorting and the supply chain: they do not need to share in the global view which the hub-manager has.

4.2.2 Robot-Agents

The instructions of the hub-manager to the shop-floor are in the format of start- and end-state conditions. The tasks of the agents are specified as: take piece x at $\langle \text{co-ordinates} \rangle$ on the in-dock and bring it to $\langle \text{co-ordinates} \rangle$ at the out-dock. Each of these operations can be assigned to a single agent. The main action needed is to bring the goods to the desired place, the function might be carried out by any agent that is able to find the goods and subsequently can bring them to the desired place. Moreover, each of these actions is independent, and can be executed by an autonomous agent. Autonomous means that the agent can carry out the operations allocated to it without external guidance, as we defined in Section 1.6.

On the shop-floor several robot-agents will be applied. Each of the robots has to be *reactive* to the environment. After receiving instructions, they are no longer guided, but carry out the work by themselves. The robots share the workspace; pieces of cargo which are temporarily put aside for later transportation, become obstacles. Moreover the robots meet each other occasionally. For this environment they need to be equipped adequately. In the present chapter we discuss the design of the individual robots. In Chapter 5, we deal with the interactions of the robots and the team aspects.

4.3 Sensory-Based Robotics

In this section we discuss the design and implementation of a mobile robot. We start by introducing the concept of a motor-schema. Next we build the architecture of the robot, which includes several motor-schemas. After that we discuss the im-

plementation of the robot, that is, its hardware and sensors. The latter discussion provides examples, useful for a further study of sensors in Sections 4.4 and 4.5.

4.3.1 Motor-Schemas

A robot consists of *actuators*, *sensors*, and the *wiring* between them. The actuators are the basis of a robot design, they determine (and constraint) what the robot might do. For convenience, we introduce the notion of *logical sensor*: any combination of hardware and software which measures the state of the system and the environment [Henderson, 1990; Budenske and Gini, 1997]. Similarly, we introduce the notion of a *logical actuator*, which is any combination of mechanical and electronic hardware, and software that can transform the state of the robot and its environment. The sensors and the actuators are connected through a wiring, consisting of logical connections between both ends.

In order to perform, a robot needs a pattern of consecutive activations of its actuators. For instance, manipulating a robot arm, requires a complex pattern of motor actions, in which the kinematics of the moving links has to be taken into account. A *motor-schema* comprises such patterns and denotes routines for interacting with objects. The term motor-schema is introduced by Arkin [1989]. A motor-schema is a pattern of knowing-how, and incorporates sensors, wiring and actuators. Different motor-schemas may use the same sensors and actuators. Our robot has only a steering mechanism and wheels as actuators and each implemented motor-schema uses these actuators.

The sensors have to observe whether the circumstances allow for a certain action. Once it is observed that a certain situation is the case, the robot can act. Thus, the requirements concerning perception follow directly from the motor-schemas and we can apply task-specific sensor algorithms. This is an important advantage of sensory-based robotics. The motor-schemas indicate on what to focus the design of the sensors.

In a motor-schema the wiring couples the sensors to the actuators. The approach eliminates intervening representations and reasoning. At an elementary level simple wiring might do. The more sophisticated the aims, the more complex the wiring might be. In the latter case we might speak of an "intelligent" wiring, not because the wiring itself is smart, but because the designer makes smart connections.

The motor-schemas are the basic building blocks of sensory-based robotics. In his [1990] paper, Arkin describes the design approach for sensory-based robotics in two steps. The first step is determining the motor-schemas required to perform the task, and requires a task decomposition similar to our approach in the previous chapter. In the second step, one has to determine how to collect the information from the environment by which the motor-schema will be triggered. Below, we

focus on the latter step and elaborate on the robot architecture on the sensor design.

4.3.2 Control Architecture

A robot is provided with several motor-schemas, which might use the same actuators. The subsumption architecture of Brooks [1986] connects the motor-schemas. In his introduction, Brooks [1986, p. 14] writes: *"A new architecture for controlling mobile robots is described. Layers of control are built to let the robot operate at increasing levels of competence. Layers are made up of asynchronous modules that communicate over low-bandwidth channels. Each module is an instance of a fairly simple computational machine. Higher level layers can subsume the roles of lower levels by suppressing their outputs."*

Brooks' layers coincide with motor-schemas. In the architecture of control they are ranked with respect to how they have access to the actuators. The architecture has to solve internal conflicts in resource assignment as discussed in Chapter 1. Higher level motor-schemas subsume the roles of lower ones. In our architecture it also happens that lower-level schemas ultimately dominate higher level ones. In this subsection we provide an overview our robot, its motor-schemas, and control architecture. We discuss examples of compound behaviours and explain the architecture.

Overview

Figure 4.1 contains two motor-schemas which are combined into the basic architecture of our mobile robot. Each layer a and b constitutes an obstacle avoidance motor-schema, on layer a *touch* and on layer b *avoid*. Layer c shows the basic architecture. The motor-schemas consist of several modules, ranging from a logical sensor system (TOUCH and SONAR) to logical actuators [Bomhof and Penders, 1992]. The WHEELS and STEER modules designate the driving and steering mechanisms respectively. FACTORISE is the interface to both modules, its input is a velocity vector.

The logical sensor of the touch-schema is implemented with close range sensors. As soon as an obstacle is sensed close to the robot (a few centimetres) the touch-schema causes the robot to step back and a little aside.

Figure 4.1: The basic architecture of the mobile robot.

The avoid motor-schema consists of the modules SONAR, SURVEYOR and AVOID. SONAR is the (logical) sensor, it scans the environment in a range from 0.5 m to 5 m. The output signals are collected by the SURVEYOR module which forms a map of the surroundings. The map is forwarded to the AVOID module, which subsequently determines a route through the obstacles. The route is calculated using a so-called artificial potential field procedure to be discussed in 4.3.3. For now it is important that the avoid-schema keeps the robot always more than 0.5 m away from obstacles.

The motor-schemas touch and avoid both use the same actuators. In order to establish an unequivocal signal to the actuators we use so-called *drown* and *deprive* nodes [Brooks, 1986]. At these nodes, wires originating at different modules come together. The signal coming in on the wire connected to the D-part of a drown node (called inhibitor in [Brooks, 1986] and [Connell, 1987]) is passed on, while the other is suppressed. In a deprive node the dominant wire, when active, snatches the signal from the incoming wire, thus depriving the regularly-connected module of an input signal.

In Figure 4.1 a signal of the TOUCH module coming in on the D-part of the node will overrule the AVOID module's signal. As a result the touch-schema overrules the avoid-schema: touch functions as an emergency measure, used when the robot approaches an obstacle too closely.

Compound Behaviours

In Figure 4.2, the basic architecture of Figure 4.1 is extended with the modules *GOAL-GENERATOR* and *APPROACH*. The *APPROACH* module introduces a new motor-schema called *approach*. The sensor system ranges between those of the touch- and avoid-schemas: from 0.30 to 1 meter. This area is of interest when the robot comes close to its goal. What the robot must do when so close to the goal is dependent on the particular application. We do not expand on the implementation of the approach-schema, but discuss only the interaction with avoid- and touch-schemas.

Figure 4.2: The extended architecture of the mobile robot.

In the real robot the module *GOAL-GENERATOR* is implemented as a (manually-operated) remote control. The module determines the goal at which the robot aims. It forwards the goal's position to *SURVEYOR*, the avoid motor-schema in turn, will direct the robot to the goal, while preventing it from getting too close to any obstacle. The *GOAL-GENERATOR* can also directly steer the robot. However, usually its the signals are drowned by the signal coming from the *AVOID* module. It is also possible to overrule the *AVOID* module. At the node between *SURVEYOR* and *AVOID*, *AVOID* can be deprived of its input signal; this node is accessible via the remote control.

Below, we discuss some compound behaviours that can be composed using the extended architecture. We discuss how the robot passes through a corridor, docks at a pile of boxes, and follows a leader. These behaviours apply at a goods terminal.

Corridor

When piloting the robot through a corridor, for instance a doorway, the entrance of the corridor must be made the goal for the robot. This is done by the goal generator. The avoid-schema guides the robot to a safe position in front of the corridor. If nothing interferes the robot will stop there, actively surveying its environment. To make the robot go through the corridor, control has to be taken over by the remote control. The AVOID module must be deprived of an input signal, and also the APPROACH module. The TOUCH module retains control. Thus, whenever the robot threatens to collide with the sides while being steered through the corridor, the touch-schema will generate small corrections. Thus the robot indeed passes the corridor carefully.

Docking

At a goods terminal, the robot might aim at some pile of boxes. While the robot is moving towards the pile, the avoid-schema guards it for not colliding with any obstacle, such as another robot. When it comes close to the pile, the approach-schema becomes active, by depriving the AVOID module of input signals. The approach-schema manoeuvres the robot into the right position. However, if the pile is approached too closely, the touch-schema overrules and slightly adjusts the position in order to take care not to knock over the pile. Thus any miscalculations of the approach-schema are corrected to prevent the enterprise failing at the decisive moment.

Following a Leader

In the previous examples the goal is set by the remote control. We can vary this slightly. The goal might be installed as a beacon on which the GOAL-GENERATOR orients itself. The GOAL-GENERATOR finds the goal and presents it in the map of SURVEYOR. This goal-finding mechanism can be used for quite different purposes.

The beacon itself might move, and the robot follows a moving goal. A human worker might carry the beacon and the robot follows him. Another variant is obtained when each robot is provided with a beacon. Thus, in a robot team, robots will follow each other and travel in a line formation. This behaviour is of interest at the goods terminal. Many items of cargo are transported from one side to another, and items are temporarily put aside: it is rather difficult to maintain proper order on a busy terminal. If the robots follow a leader, on the main trajectories, we obtain some order in the robot traffic.

Apart from the goods terminal, interesting applications are also found in road traffic. Cars provided with the combined goal-following and avoid-behaviour might be steered by automated pilots on the main trajectories. The cars will follow one another, that is the goal-finding; but will keep fixed distances between each other which is what the avoid-schema does.

4.3.3 Implementation of the Robot

The mechanical construction of the robot is straightforward, photograph 4-1 shows the robot. Our purpose is to have a mobile testbed for the implementation of the sensors and the wiring. Relatively little attention has gone to its mechanical design.

Photograph 4-1. The mobile robot

The frame of the robot is a metal sheet of 260×360 mm. Underneath the following items are mounted: a front axis with the driving wheels and two steering wheels at the back (as for a forklift), a differential gear, a steering construction, batteries, and the remote control receiver. Most of these components are standard equipment for model making. The robot's maximum speed is about 3 m/s, but in the experiments it is usually between 0.5 and 1 m/s.

The rack, on top of the sheet, is used to store the electronic hardware. Each motor-schema has a dedicated print board with a microprocessor, an EPROM, some additional memory, timers and so on. Each print board is provided with a slot to plug in the print of the next level module or the remote control receiver. These slots allow the realisation of the drown and deprive nodes.

At the front of the metal frame the sensors of TOUCH are mounted. It consists of infrared (IR) LEDs and photoreceptors. The LEDs light up in turn, while the photoreceptors are all active. The sensory system observes an area of about 8–30 cm in front of the robot. In the experiments the robot will normally only move forward, so sensors are only mounted in the front of the robot.

On top of the cage-like construction the turning sensor of SONAR is mounted. It is an Hydrotech HE320 ultrasonic switch, shortly US-sensor. The sensor has been constructed as a device turning left-right over 180°. Only one US-sensor is mounted, resulting in a forward view. A side-effect of this construction is that the maximum sample interval is longer for angles at the left and right hand side of the robot.

The Touch Motor-Schema

The IR sensors of TOUCH are divided into a left and a right hand set. TOUCH compares in which of these sets most sensors spot an object. Dependent on this, it generates a step backwards, either a little to the left or to the right. If the sensor information is indecisive, the robot chooses a direction at random. The algorithm initiating the step backwards is called *flee-opposite* and is further discussed in Chapter 5.

We have to ensure that the correcting moves of the touch-schema can be performed safely. The sensors are mounted such that also an area at the sides of the robot is watched, the robot will not go beyond this area in the correcting moves. Each correcting step is no larger than the maximum range of the sensor system. The correction process will be repeated as long as any of the sensors are activated. As might be expected the resulting behaviour is rather slow because of the step by step movements, but that is what is intended. The signalled object is quite close and a sharp panic reaction is not required.

A problem with reflection based sensing is that the sensors are more sensitive to shiny surfaces than to, say, dark cloth. A shiny surface might be detected at 1 or 2 meters, while the cloth is only detected at 0.1 meter. Because of this, no attempt is made to determine the exact distance to the robot, we only set a threshold. The LEDs are mounted such that the emitted light beam is directed towards the floor, Figure 4.3 shows that an area of maximally 30 cm is watched. Outside this area nothing is seen. To avoid noise generated by the particular properties of the floor the photoreceptors are directed such that they are sensitive only in the area 5 cm above the floor. To further separate noise from actual sensing we use a threshold

on the received signal. We can adjust this threshold to make sure that the less reflecting surfaces are indeed sensed.

Figure 4.3: Top and side view of the sensors of TOUCH.

As Figure 4.3 shows, the highest point observed is about 7 cm above the floor, but the robot itself is about 30 cm. This means that the touch-schema does not provide full protection to legged obstacles. In our experimental environment, as well as at the cargo terminal, such objects are relatively rare. However, the decisive argument is that these objects will be spotted by the SONAR module.

The Avoid Motor-Schema

The SONAR module delivers pairs of data containing an angle and a distance. SURVEYOR collects the into a map, and passes it to the AVOID module. The AVOID module determines a new course for the robot. It interprets the map as a so-called *artificial potential field*. Obstacles detected are considered as a source of resistance, while the goal is taken as a source of attraction. The magnitude of these forces decrease as a function of the distance. (To visualise it one often represents the potential field as a landscape in which the repulsive sources form mountains and the sources of attraction valleys.) The AVOID module determines a route of decaying resistance, (it steers the robot through valleys in the landscape). This and similar procedures work rather well in implementing obstacle avoidance [cf. Khatib, 1986; Hwang and Ahuja, 1992]. Examples of resulting behaviour are

shown in Figure 4.4. In this section we focus on the sensor; in Chapter 5, we discuss the potential field procedure extensively.

Figure 4.4: Trajectories of two robots using the potential field procedure.

The SONAR is implemented as a turning device. The sample rate is related to the moving speed of the robot and the obstacles. We may assume that at the goods terminal, obstacles move about as fast as the robot itself. With this assumption we can derive the sample rate. We have to prevent an obstacle approaching the robot without being noticed. This might happen when at the first scan the obstacle is just outside the maximum range of sonar, and in the next scan has passed beyond the minimum range. The time between successive scans must at least be smaller than or equal to

$$T_{\text{sample}} = (R_{\text{max}} - R_{\text{min}}) / 2V_{\text{max}}$$

where R_{max} and R_{min} are the maximum and minimum range of the sonar, and V_{max} is the maximum of the obstacle's speed and the robot's own speed (both of which are assumed to be more or less the same). The minimum range of the sensor is about 0.5 m, the maximum range depends on the reflectivity of the surfaces; about 5 meters may be assumed. Combining this with a 1 m/s velocity of the robot, the time between successive scans may be as large as two seconds. In practice, our

sonar rotates faster. At the left and right hand sides the sample interval is two seconds, straight in front it is one second. The robot reaches maximum speed only when moving forward, and this is where the robot approaches obstacles fastest.

4.3.4 Sensing and the Robot Design

So much for the construction of the robot. We have discussed the implementation of the touch and avoid motor-schemas and the logical sensors TOUCH and SONAR. An IR-sensor (infrared) and an US-sensor (ultrasonic) are applied. The sensing techniques have far reaching consequences for the robot design; we discuss their pros and cons in this subSection. The discussion also sets the stage for the next sections.

Measurement of distances with US-sensors is relatively accurate (deviations are far less than a cm per meter distance). A known problem with US-sensors is their relatively wide opening angle, see for instance [Badreddin and Holenstein, 1991], for our sensor it is 15° (a US-sensor works with sound beams, which cannot be directed more precisely). Moreover, our SONAR module signals only the closest point of an obstacle. Figure 4.5 gives an impression of how the environment is presented: it is partitioned into conic-shaped areas. If an obstacle reflects the sound beam, we need to consider the whole area from the obstacle onwards as occupied.

Because of this, finding a passage between two obstacles can be problematic. Figure 4.5 shows two situations, one with two objects and one with one object. Assume there is a passage of 0.40 m wide between them, as depicted in the upper situation. With a lucky scan, which has to go right through the middle of the passage, the passage can be found, but only when the robot is as close as 1.51 m. Otherwise, the sensor reading is the same for both situations. In the worst case, the corridor will only be found from a distance of 0.76 meter.

Figure 4.5: Representation of environment by the US-sensor of AVOID.

To improve the resolution of the sensor system, the US-sensor might be combined with IR-sensors, similar to the ones used in the TOUCH Module. An IR-sensor can emit fairly precisely directed light beams. Reading is based on the reflection of the light. Measuring distances with IR is more complicated and requires different equipment, and we do not consider it. However, with their smaller opening angles of about 5° , IR-sensors improve the resolution of angles. Fusion of the information is on the face of it straightforward: once the IR-sensor has detected an object, the angle is known with 5° accuracy and the US-sensor provides the distance. With a 5° opening angle a corridor of 40 cm wide will be found with certainty at a distance of 2.30 meters (see Figure 4.6), this is quite an improvement compared to the 0.76 m (worst case) of the US-sensor. This fusion of sensor techniques not only improves resolution, but also reliability. In case one type of sensor fails because of “difficult” objects, the other one still provides information about the obstacle, though with a decreased precision. However useful it seems to combine US with IR, it triggers difficult problems. These problems are the subject of the next section.

Figure 4.6: Representation of environment after combining US and IR.

4.4 Sensor Basis

We have dealt with the limitations of a robot design and how the sensors are incorporated into the design. In the remaining sections we focus on the sensors as sources of information. For this purpose, we introduce the notion of sensor basis, which we use in Section 4.5 to clarify how a robot observes its environment. The sensor basis helps to understand what are the prospects as well as the inherent problems of sensor data fusion.

In current chapter we are dealing with a robot in an environment with objects. Intuitively objects are such things as boxes, bags or robots¹⁷. In this (naive) context

¹⁷ Things with which the art of moving objects is concerned.

we define the formal notion of sensor basis. This notion is developed consciously. At first sight this might seem a little out of proportion when treating sensor fusion. However, the reason for doing so is that the notion of sensor basis is also the starting point for further philosophical analysis in Chapter 6. There we will discuss some fundamental points of perception theory, and we also have to reconsider the notion of object. For the present chapter the naive notion suffices.

4.4.1 Definition of the Basis

Our introduction of the *sensor basis* is based on the definition of *knowledge basis* in [Rauszer, 1992]. For convenience, while introducing the definition we parallel Rauszer and use the term “knowledge”. However at the end of this section, we redefine this term. In a robot design we have distinguished *sensors* and *wiring*; consequently we have to introduce two bases: a *sensor basis* and – for the wiring – a *knowledge basis*, more about this at the end of the subsection.

Assume an agent r has at its disposal a set of logical sensors which measure certain attributes $\alpha \in A_r$. The sensors assign to an object o in the universe of discourse¹⁸ U , an attribute α with a value $v \in V_\alpha$.

Definition 4.1. *Sensor value.* $\alpha : U \rightarrow V_\alpha$; $\alpha \in A_r$;
 $\alpha(o) = v$; $v \in V_\alpha$; $o \in U$.

A sensor might not discriminate between two different objects. This is expressed by $\alpha(x) = \alpha(y)$. Rauszer [1992] defines an *indiscernibility* relation for an agent r as:

Definition 4.2. *Indiscernibility.*

$$\text{ind}(A_r) = \{(x, y) \in U \times U \mid \forall \alpha \in A_r (\alpha(x) = \alpha(y))\}.$$

$\text{Ind}(A_r)$ is an equivalence relation: obviously it is symmetrical, transitive, and reflexive. Hence, the set of equivalence classes is a partitioning of the universe U . Thus, if two objects fall within the same equivalence class $[o]_{A_r}$, it means that the agent cannot distinguish one from the other. The set of all equivalence classes reads:

Definition 4.3. *Basis.* $E = \{u \subseteq U \mid u = [o]_{A_r}, \text{ for some } o \in U\}.$

¹⁸ U is the naive universe of objects mentioned above. Throughout the sections below, objects $o \in U$, will be objects in the naive sense.

E is called the *knowledge basis of agent r* over the universe of discourse U , determined by its set of attributes A_r . The knowledge basis indicates which objects can be distinguished by the agent. When agent r spots an object, its sensors assign attribute values to the object. Hence, any knowledge the agent *can* have about an object, can be represented in the knowledge basis.

Some further notions illustrate the usability of the basis. Assume an arbitrary set of objects $X \subseteq U$. Then, the *positive knowledge of r* about the set of objects X , is a subbasis of E and consists of the set of equivalence classes generated by the elements of X :

Definition 4.4. *Positive Knowledge.*

$$I^r(X) = \{u \subseteq U \mid u = [x]_{A_r}, \text{ for some } x \in X\}.$$

X is *positively known* by r , if each $x \in X$ is in a separate classe of $I^r(X)$. This means that r can (positively) distinguish among the objects in X , i.e., it can tell each object apart from the others in X . Note that in the example in Figure 4.5 the set of objects is not positively known by the US sensor.

Besides the objects belonging to X , the sensors might be sensitive to other objects as well. Moreover, an object not belonging to X might have the same sensor readings as a certain object belonging to X . The best knowledge which the agent can collect about X is that X is positively known and that the sensors do discriminate objects not belonging to X as such.

Definition 4.5. X is *completely known* by r iff X is positively known by r and any $o \notin X$ is recognised as such.

Analogous to definition 4.4 we can define the positive knowledge $I^r(X^c)$ of $X^c \subseteq U$ (X^c is the complement of X in U). We now have:

Proposition 4.1. X is completely known iff

$$X \text{ is positively known and } I^r(X) \cap I^r(X^c) = \emptyset.$$

Proof. We have to prove that agent r recognises objects from X^c as not belonging to X . $I^r(X) \cup I^r(X^c) = E$. Moreover, E is a partitioning of U therefore, any $o \in X^c$ is in some equivalence class in E , and this class is either in $I^r(X)$ or in $I^r(X^c)$. $\square$

The notion of positive knowledge is the most relevant for our purposes. The designers of the robot can describe the environment as consisting of a set X of objects. The positive knowledge of the robot indicates which of these objects it can (possibly) observe. Here a warning must be given: care has to be taken in concluding that the environment is *positively known* by the robot: sensors do not al-

ways discriminate objects as we would like them to do; Figure 4.5 shows an example. Ensuring that the set of objects is completely known is a further problem, in particular because it is not known what objects the universe contains. Some objects might not be signalled by the sensors or they might be seen as objects from X . This topic is discussed in the next section.

In a robot design we distinguish sensors and wiring. Therefore, we have to adapt our notions. S denotes the set of attributes for which the logical sensors produce values. By applying Definition 4.2 and Definition 4.3, we obtain $\text{ind}(S)$, we call this the *sensor basis* E_{sens} of the agent (an example is given below). We reserve the title *knowledge basis* when applying the definitions to the wiring. Letting C denote the set of logical connections or “information processing functions” implemented in the wiring. Then we call $\text{ind}(C)$ the *knowledge basis* E_{know} . Any robot comprises both bases.

Example

In the robot described in Section 4.3 provides an example of a sensor basis. Its sensor observes cones, which are either occupied from a certain distance onwards or empty, thus the sensor readings consist of the angle under which the object is spotted, and the distance towards the object. Figure 4.5 illustrates the sensor basis E_{sens} . With this basis the robot observes the environment. Information picked up in the environment (as far as this agent is concerned) always consists of such pieces of information.

4.4.2 Cognitive Basis

The design of a robot specifies an internal cognitive structure, which will incorporate E_{sens} and E_{know} . The process of information collection starts with E_{sens} and proceeds in the wiring, with E_{know} . The sensors filter out a certain view of the world built upon the sensor basis E_{sens} . Analogously the wiring also processes data that provide a certain view of the world. The point is to fit these bases into a sound design and match them into a cognitive basis, Figure 4.7 gives an overview of these bases.

Figure 4.7: Overview of the representation bases.

Sensor and Knowledge Basis

In a robot, the information collected by the sensors is forwarded to the wiring. The information flows from the sensor basis E_{sens} to the knowledge basis E_{know} . As in information theory one can determine how much information the sensors forward to the wiring. The basis is the number of states known to be possible at the source. In information theory this is called *entropy* [Shannon, 1950]. The amount of information forwarded is bounded by how many of these possible states can be excluded at the end of the channel. The source is the sensor basis E_{sens} and the end of the channel is E_{know} .

The entropy gives a measure of how much information is forwarded. In order to achieve a consistent design the sensor basis and the knowledge basis need to be matched to one another. The possible views generated with the sensor basis need to correspond somehow to the possible views that can be generated in the knowledge basis, and we need to go beyond an entropy requirement. E_{sens} and E_{know} need to deal with the “same” classes and relationships; for instance, if the sensor basis is as in Figure 4.5, adjacent cones of E_{sens} need also to be treated as adjacent in E_{know} .

Formally this is expressed using a mapping from E_{sens} to E_{know} . The mapping needs to preserve the mentioned relationships. In the ideal case E_{sens} and E_{know} are isomorphic and equivalent, and process the same information. However, in a robot design the required isomorphism is not obtained straightforwardly. Several cases

are possible. We use the algebraic notion of an isomorphism to lay down the requirements concerning relationships.

Assume that on E_{sens} a certain n -ary relationship S holds, and an analogous n -ary relationship K holds on E_{know} . An isomorphism from the sensor basis to the knowledge basis is defined as follows.

Definition 4.6. *Isomorphism (restricted definition)*¹⁹.

A mapping $f: E_{sens} \rightarrow E_{know}$ is an isomorphism from $\langle E_{sens}, S \rangle$ to $\langle E_{know}, K \rangle$ iff

- (i) f is a bijection from E_{sens} to E_{know} ;
- (ii) for all $s \in E_{sens}$ $S(s_1, \dots, s_n) \leftrightarrow K(f(s_1), \dots, f(s_n))$.

Clause (i) in Definition 4.6 requires that f is a bijection. It means that every class of E_{sens} maps on a single class of E_{know} . Moreover, every class of E_{know} is the image of a class of E_{sens} . Clause (ii) requires that relationships are preserved under the mapping, and that any relationships in E_{know} , can be traced back in E_{sens} . When an isomorphism as required in Definition 4.6 exists, we can define an internal cognitive structure for a robot, we call it the cognitive basis and denote it as Φ .

Definition 4.7. *Cognitive basis.* If f is an isomorphism from E_{sens} to E_{know} , the cognitive basis Φ consists of the equivalence classes generated by f .

The cognitive basis Φ consists of equivalence classes of objects that have the same values in the logical connections. Thus, Φ consists of those (classes of) objects $o \in U$ that can be distinguished by the sensory system as well as by the wiring. The sensors and the wiring (sensor and knowledge bases) must have positive knowledge about them. Thus, the definition of Φ stipulates the design principle that there is no point in making the sensors "see" more than the wiring might "know" and vice versa.

The practical problem of the robot designer is to establish the isomorphic mapping as required in Definition 4.7. Sometimes, sensors and wiring straightforwardly map onto each other. But, more often certain potentials of the sensor do not match to the wiring and vice versa. The isomorphism has to be obtained in an iterative process. When establishing the isomorphism, a certain loss of information or a restricted utilisation of computing power has to be taken for granted. As far as our robot is concerned, the reflection pattern of the sound beams emitted by the US sensor is rather rich. Even the presence of a second object can be detected from it [Koomen, 1996]. The applied obstacle-avoidance procedure processes the addi-

¹⁹ For a full definition of an isomorphism we refer to [van Dalen *et al.*, 1987]. The version presented here suffices to show our points.

tional information. However, detecting the second object requires highly sophisticated signal processing. The effort to detect such details has to be weighed against the design and application aims and we decided not to implement it for our robot. The robot performs quite well without such additional information (as we will see in Chapter 5).

4.5 Constructing Sensor Views

In the previous section we defined the cognitive basis starting from known objects (the Definitions 4.5, 4.6, and 4.7 start from objects). In this section we deal with an unknown environment. The objects themselves are unknown, not only is their position not known, but also their sensor characteristics are not known. We discuss how to draw conclusions from the sensor data in such situations, compare various approaches and combine different types of sensors.

Our robot uses the sensor information implicitly in its internal processing. In conformance with sensory-based robotics, sensor signals are directly forwarded to motor-schemas. For instance, in the avoid-schema, the sensor signals are directly forwarded from the SONAR to the AVOID module. No explicit representation of the environment is constructed. However, in this section we make the sensor view explicit for several reasons. First, because in a discussion of cognitive capabilities (see later in Chapter 5) we need to compare views of the world. Second, it brings to the surface some profound problems in artificial perception. Third, and most important to the theme of this chapter, the difficulties which we encounter in constructing a sensor view of the environment supports the choice for a sensory-based design approach. Sensory-based robotics purposely avoids constructing explicit sensor views.

4.5.1 Surveying

When a robot first encounters a certain situation, the sensors provide only approximate information, because they have limited resolution. The robot can refine the information, for instance by moving around and surveying the environment. In Figure 4.5 above, a single scan of the environment is presented, in surveying, similar scans are collected and might be combined.

A method which is often applied is to collect survey data into a so-called *certainty grid* (see for instance [Borenstein and Koren, 1991]). The environment is covered by a (usually square) grid pattern. The successive scans of the sensor are transformed to the grid pattern. It is determined whether each square is occupied or not in the scan that is evaluated. The value obtained, combined with previously obtained values is fed into a formula which calculates a *certainty value* for the

square. The formula incorporates a certain conservatism, so that the square does not change at once (in one scan) from occupied to free. Moreover, the formula can compensate for uncertainty inherent to the sensor. The certainty grid method is useful for path planning in the "classical mover's problem", that is, planning the moving of a rigid robot among stationary obstacles of fixed shapes [Hwang and Ahuja, 1992].

Note that in this method, the robot is internally processing the squares of the grid pattern. The cognitive basis consists of sets of grid squares. Sets of grid squares do not represent objects straightforward. Some information gets lost. Two objects are the same if they deliver the same sensor readings in terms of grid squares. A set of grid squares might represent one or more objects. For global path planning the loss of information is not really relevant. However, on a detailed level, we cannot give the robot an instruction to pick up an object, which is a set of grid squares.

A further drawback of certainty grids is that the grid has to cover the whole environment. A remedy for this might be found in the so-called *quattree* of Zelinsky [1992]. In this algorithm, the grid pattern is developed depending on the results obtained in subsequent scans.

In any case, in the process of driving around and collecting information the view of the world becomes more detailed. Besides the grid methods, there exist logical systems to deal formally with this process of growth and refinement of information. We mention in particular the *update-semantics* of Veltman [1996]. This system basically works by making finer distinctions among possible worlds. The sets of possible worlds (they are called *information states* in [Veltman 1996]) are further subdivided as the amount of information increases. The descriptions of the possible worlds become more specific. Similarly, by collecting more sensor data the robot further enriches its information about the environment and cancels out other possibilities.

By surveying a robot collects more information from its environment and refines its view. However, the general drawback of the information collecting process is that it presumes that the environment is not changing during the collection process. Collection of information is a time consuming process. In a dynamic environment such as a goods terminal, obstacles (other robots) might move as fast as the robot itself. The survey might be obsolete even before it is fully constructed.

4.5.2 Sensors divided over Motor-Schemas

In the dynamic environment of a goods terminal a robot is not given much time to reason; when something happens the robot is supposed to react instantaneously. To refine the information (at run time) we combine different types of sensors. We investigate how to deal with data originating from different sensors. We first re-

flect on the role of sensors in sensory-based robotics, next we discuss sensor fusion.

In the architecture of our robot (see Subsection 4.3.2) we apply different sensors with different ranges of operation. The range of the TOUCH sensor is from 0.05 m to 0.3 m; the APPROACH sensor ranges from 0.3 to 1.0 m; and the SONAR ranges from 0.5 to 5.0 m. The choice for sensor ranges is based on the action radius of the motor-schemas, all of which differ. The sensor bases are likewise distinct, and not really combined. "Integration" of the sensor signals results from the fact that the robot can execute only one action at a time. This is rather easily achieved as all motor-schemas use the same actuators and the signal to the actuators is made univocal. The idea is typical for sensory-based robotics. Each motor-schema has its specific sensors. With the motor-schemas complex behaviours are constructed. A situation is seen as composed of a variety of specific conditions to each of which some motor-schema reacts. The set of the most frequently occurring conditions is determined and motor-schemas responding to it are designed.

In combining motor-schemas we implicitly assume that the conditions observed by the sensors are mutually independent and that the actions are compatible. Here we encounter a weakness, there is no guarantee that the reaction to one condition leaves other conditions unaltered. Actions of one motor-schema might prevent the sensors of another schema from being triggered. In Subsection 4.3.2 we discussed the behaviour of passing through a corridor. As discussed there, the avoid-schema keeps the robot away from obstacles. Without some (high level) interference there will be no progress in passing through a corridor. It is a variant on the frame problem – the problem that actions might have unforeseen side effects, cf [Toth, 1995]. We note that this problem is implicit in all approaches which we have subsumed under the title sensory based robotics.

4.5.3 Sensor Data Fusion

Above we discussed motor-schemas with one type of sensor. However, a motor-schema might combine data originating from different types of sensors. In Figure 4.6 above, we added an IR sensor to the US sensor. On the face of it the design increments the angle measurements, as suggested in the discussion. However, simply adding IR data to the US data might cause inconsistency since the behaviour of a light beam is quite different from a sound beam. Cloth reflecting a sound beam rather well can completely absorb light beams. Hence, some objects in Figure 4.8 might not be detected at all by the IR sensor.

Figure 4.8: Inconsistency by combining an US and an IR sensor.

Redundancy and Complementarity

When combining data from different sensors it is relevant to distinguish between *redundancy* and *complementarity*. Redundancy and complementarity are each others complement. Redundancy between two sets of sensors data means that the one data-set contains information that is also present in the other, possibly in a different form. *Fully redundant* sensors deliver the same information. For instance, a photograph and its negative (though they are visually different) are fully redundant. Complementarity of sensors means that they deliver different information about the same source.

Using our notion of sensor basis, we provide a further clarification. Redundant sensors deliver the “same” sensor bases. We formalise this in the following statement. Assume E_1 and E_2 are the sensor bases generated respectively by sensor1 and sensor2.

Proposition 4.1 *If sensor1 and sensor2 are fully redundant, then E_1 and E_2 are equivalent.*

Proof. Equivalence of E_1 and E_2 means that there is an equivalence relation between the classes of E_1 and E_2 , such that for every pair of objects x and y : they are in the same class of basis E_1 if and only if they are in the same class of basis E_2 . For the redundant sensors we have to prove that E_1 and E_2 are equivalent. Assume, x and y are in the same class: $x \in [x]_{\text{sensor1}}$ and $y \in [x]_{\text{sensor1}}$. This means that whether sensor1 observes x or y , it delivers the same results: $\text{sensor1}(x) = \text{sensor1}(y)$. Sensor1 and sensor2 are redundant so $\text{sensor2}(x) = \text{sensor2}(y)$. Thus x and y are in the same class of sensor2: $x \in [x]_{\text{sensor2}}$ and $y \in [x]_{\text{sensor2}}$: E_1 and E_2 are equivalent. The converse, proving that if E_1 and E_2 are equivalent then sensor1 and sensor2 are redundant, goes along the same steps. $\square$

Complementarity means that sensors deliver data originating from different characteristics of an object. They generate sensor bases that are not (regularly) connected. We have seen an example of complementary of sensors in Figure 4.8 in the upper situation.

On a formal level we can comfortably deal with complementarity between sensors, representing objects in different dimensions. The idea of a perception basis was already given in [George, 1961, p. 319–326]. Contrary to the approach we adopt, this basis is generated by straightforwardly adding the data of different sensors in extra dimensions. However, the practical problem is that there is no key to combine the information. There is no answer as to how to interpret a combination of the sensor data. What does a combination of certain sensor data values mean? How can we determine which combination of data (that is, attribute and values) belong to the same object. This is essentially the problem of what is called *sensor data fusion*.

Objects

Dealing with unknown objects is a problem that arises when applying logic (or other formal approaches) to the field of sensors and artificial perception. Most logical systems – such as [Veltman, 1996; Rauszer, 1992; Hudges & Cresswell, 1984] – make the basic assumption that one deals with *individuals* or *individual objects*. However, to extract the existence of individuals from the sensor information, goes far beyond a simple interpretation of the sensor data. What the sensors deliver are only various series of values of attributes, but it is not possible to trace from the sensor signal which objects acted as arguments.

In the definitions given above, the sensor basis and cognitive basis are constructed starting from external objects. The equivalence classes that make up the cognitive basis are sets of attribute values observed when sensing the object. Thus internally – within the cognitive basis – the objects are sets of observed properties. When constructing a cognitive creature, this means that prior to any cognitive activity, the creature has to be “taught” about the objects. The designer has to determine which combinations of attribute values count as an object.

When combining different sensors, the cognitive basis changes. The aim is a refinement of the cognitive basis Φ . A straightforward solution is first to apply sensor1 and at a second moment sensor2. Thus we obtain Φ_1 as a refinement of Φ_0 and a morphism²⁰ from Φ_0 to Φ_1 is straightforward. But the intended refinement is not always straightforward: one object might split into two – this happens in the process of finding a corridor – or, as in the IR and US example of Figure 4.8, objects found by one sensor might not be found at all by the other. One cannot simply extend one basis say Φ_0 to Φ_1 , by just adding new attributes. In the case of the IR and US a reasonable mapping from Φ_0 to Φ_1 has to be constructed. In order to maintain consistency, this process needs to be rigorously constrained, as illustrated in the argumentation accompanying the system of Veltman [1996].

Our approach to sensors and sensor bases brings the sensor data fusion problem out into the open. For instance, the certainty value approach discussed above can easily combine data of different sensors. However, this method is also not really appropriate to deal with consistency problems. The formula with which one calculates the certainty values, will have to deal with the inconsistency within the computation. But what should be done if an object is not spotted, and how can this case be discriminated from the case that there is indeed no object?

Redundancy and complementarity of sensors might provide some foothold. Redundancy can be used to determine correspondences between images and to anchor one basis into the other. Houzelle *et al.*, [1994] define a redundancy rate, using the formalism of information theory. Introducing their definition into our concept of sensor basis is straightforward: consider the generated sensor bases E_1 and E_2 , as start and end points of a communication channel. The sensors are fully redundant if E_1 and E_2 are equivalent, see Proposition 4.1. If, on the other hand, no common ground can be found, we are dealing with complementarity. However, how redundancy should be applied to detect unknown objects from the sensor data is not explained.

In any case, one has to accept and be aware that the sensors and the generated sensor cognitive bases do not match to each other. Moreover, we note that, although the concept of logical sensor sounds generic, one cannot easily substitute one (hardware) sensor for another, say US by IR. Substitution in general changes one logical sensor into another and consequently the sensor basis is changed.

²⁰ A morphism f from $\langle \Phi_0, S \rangle$ to $\langle \Phi_1, K \rangle$ retains relationships: (ii) for all $s \in S(s_1, \dots, s_n) \rightarrow K(f(s_1), \dots, f(s_n))$, compare clause (ii) of Definition 4.6:

4.6 The Robot Designer

In the previous section we discussed artificial perception and the information collection problem from an engineering point of view. Artificial perception often turns out to have unforeseen complications: it is non-intuitive to human designers. The irony is that human beings seem to have inherited perception mechanisms that disguise irregularities and emphasise similarities and patterns. We meet here the so-called *frame-of-reference* problem [Pfeifer, 1995]. When evaluating cognition, usually the frame of reference is the common human one. However, when discussing other types of cognitive (artificial or natural) creatures this leads to fuzziness.

4.6.1 The Frame of Reference

In robotics we have to deal with the robot and its designer. Below, we make the role of the robot designer explicit. The robot designer has a certain task in mind for his robot, which has to be executed as the designer designates. Consider, for instance the robot when lifting its arm. Whenever the robot lifts its arm, within a series of actions which result in gripping the work piece, the designer is confident. However, he worries whenever the same action occurs suddenly while the robot is supposed to be at rest.

Being more precise about objects and using the notion of cognitive basis we can unravel the designer's point of view and that of his robot. The cognitive basis the designer wants to implement in his robot is a simplification of his own, let Γ denote this intended cognitive basis. Γ is based on the naive view of what objects are (refer to the Subsection 4.4). The design process consists in clarifying what Γ consists of in terms of implementation. The result is the cognitive basis Φ , obtained via E_{sens} and E_{know} as discussed see Figure 4.7.

E_{sens} and E_{know} are equivalence classes of attributes. They are not realms of objects as in the naive Γ is supposed. This difference gives rise to a philosophical dispute which is further discussed in Chapter 6. Having observed this difference, we discuss here what the consequences are when designing a robot. The distinction between Γ and Φ is rather important to apprehend the problem of interpreting sensor information.

As we have seen the sensors of our robot do not provide the information to conclude the existence of objects (which the designer and other observers see). In a successful robot design, it is neither necessary nor possible that the robot's cognitive basis meets that of the designer. What is required is that the cognitive basis of the robot is sufficient to be acted upon. This means that the creature can distinguish all possible states of the environment (as seen from the designer's set Γ) in which its actions can be applied from those in which it cannot apply its actions. The implementers' question is: how to select those data that reflect the $\gamma \in \Gamma$

which the designer has in mind by some $\varphi \in \Phi$, which includes establishing a mapping from E_{sens} onto E_{know} .

4.6.2 Design Directions

Given the gap between the cognitive basis of the designer and his robot, the designer can proceed in two directions. The first is to try to elaborate a fine representation of the external objects (and the external world as a whole). In other words, he can try to make the cognitive basis of the robot as finely detailed as the world itself (or as the designer sees the world). We have discussed some drawbacks with which this approach is confronted from the start. And we are sceptical that such an approach will ever succeed. However, if it is not possible to make the cognitive basis match the world, we might make the world match the cognitive basis, this is what the first direction takes resort to. The second direction is that we content ourselves with the fact that the sensors and thus the cognitive basis can represent only some particular aspects of the external world. We do not change the world but concentrate on making the most out of what we indeed do have.

One can say that the development of so-called *industrial robots* takes the first route. The environment in which such robots are applied is purposely restricted to well-specified lab-like conditions. The knowledge basis E_{know} of these robots is (generally speaking) much more specific (or larger) than the sensor basis E_{sens} . In a lab one can control and manipulate the environment. By externally controlling the feeding of the sensors, one bridges the discrepancy between E_{sens} and E_{know} . What cannot be implemented in the robot has to be excluded in the environment.

It is the second approach which distinguishes sensory-based robotics. It takes it as given that a robot's sensors are limited, and looks for what can be obtained. As a consequence, the robots can never achieve the same high precision as the industrial robots. However, the sensory-based robots can be applied in under-specified environments. And here the status of the cognitive basis comes into question.

The designer's view of how his robot will be embedded or prepared for its future environment is important. This seems obvious, however, the focus is often on the wrong thing. When discussing the Control Architecture in Subsection 4.3.2 we cited [Brooks, 1986]. In the citation it is said that modules "(...) communicate over low-bandwidth channels (...)" The designer has to make explicit what he aims for (that is, what set Γ he is working with) and how he implements this in his design (what he thinks the Φ is that resembles Γ). Only after that does it make sense to check the bandwidth of the channels. When determining Γ and implementing Φ one answers either implicitly or explicitly the question posed by Arkin [1990] "(...) how do we select from the wealth of sensor data available that which is relevant?" and that which matches the designer's set Γ .

4.7 Conclusions from Robot Design

The central topic in the discussion of this chapter is the concept of performing autonomously. A mobile robot is often called an autonomous agent, where in fact it only performs autonomously. We discussed motor-schemas in sensory based robotics. In a motor-schema sensor information directly triggers actions, a motor-schema is typically designed to solve a particular task and implies know how. Thus, a motor-schema exemplifies autonomous performance.

The motor-schemas have no “insight” into why they are performing. Goals and aims of a motor-schema must be set on a level provided with a more global point of view. We defined a sensor basis and cognitive basis. We analysed the problem of sensor fusion and distinguished frames of reference. We elaborated on two examples: the touch and avoid motor-schemas of our robot. The discussion showed the importance of the designer and the way his role permeates the robot design. We also discussed compound behaviours, using passing through a corridor as an example. Performing this behaviour required that the robot is explicitly told that its goal is beyond the corridor.

Motor-schemas can be considered to be agents (that know how to execute a certain behaviour). This view is made explicit by Arkin [1992], who describes a single robot as composed of multiple agents, connected via the control architecture. The low-level agents of sensory-based robotics can have applications quite different from the robotics field. The avoid-schema might be applied as a anti-collision device for cars or aircraft. In such applications many “avoid”-agents will meet. In Chapter 5 below, we study low-level agents applied in a multi-agent team. The discussion in the present chapter was restricted to robotics. In Chapter 6, we resume discussing these points in the context of a natural (human) being, and consider epistemology.

4.7.1 Problem Statement 2

In this chapter we have dealt with the second problem statement: *“how can a set of resources be made to perform as an actor within a supply chain”*, and the local point of view. In Chapter 1 it was stated that a supply chain requires a division of labour, therefore tasks are specified. Locally, agents need know how to execute a task, robots are typical instances of agents able to perform an action. The robots are applied on a cargo terminal to perform sorting. We conclude from this example that the agents need not know anything about the sorting and the supply chain: they do not need to share in the global view.

5 Spatial Interaction of Robots

We are dealing with our second problem statement – *how can a set of resources be made to perform as an actor within a supply chain?* In Chapter 4 we discussed how to construct a robot agent; in this chapter we study robots operating with fellow robots. We focus on the spatial features of the interaction in a multi-robot team. Interaction appears to provide a starting point for arranging a chain of actions.

Robots are instances of spatially performing agents. Our robots exhibit only two types of behaviour: goal finding and obstacle avoidance. These behaviours are applied within a team. In Section 5.1 we delineate the spatial characteristics. Section 5.2 deals with intentions as the starting point to characterise and analyse the collective behaviour of a team. In Section 5.3, we develop a formalism to describe interaction and distinguish types of interaction patterns. Using mathematical analysis we assess in Section 5.4 the interaction in specific teams in detail.

5.1 Spatial Features

Robots mainly perform spatial movements. We saw in Section 4.7 that a robot might consist of several software agents. Currently, software agents are widely discussed, refer for instance to [Bradshaw, 1997] and [Wooldridge and Jennings, 1995]. Characteristics of software agents under consideration are, for instance, pro-activeness, reactiveness, communication and autonomy; moreover agents use an ontology. These issues also come up in a robotics discussion, however, in a slightly different interpretation. The robots are pro-active when performing goal finding, and reactive when avoiding obstacles. We discussed them in Chapter 4. In this section we consider the *ontology* of a team in which the robots do *not communicate*, but are *autonomous* and do interact, subsequently we informally define the notion of *having an intention*.

5.1.1 Ontology

Guarino [1998] circumscribes ontology as “*the categories accounting for a certain vision of the world.*” Our robots have a limited cognitive basis, thus the ontology of our robot domain is limited. The robots neither have a notion of a goal nor of an obstacle, they only process within terms of spatial positions and areas. Goal finding is aiming towards a certain geometrical position. Areas may be either free or occupied; *objects* coincide with occupied areas. Thus, avoiding obstacles means not to enter occupied areas. For convenience, we continue to speak about goals, objects, obstacles and other robots. This will not lead to confusion as long as it is borne in mind that the terms denote geometrical points or regions, without any further connotation.

Generally, a robot is assumed to move about in a stationary environment. The position of the robot changes amongst the presumed fixed objects. However, we might also consider the position of the robot as fixed and consider the objects as moving. Thus, the notion of a *moving object* (i.e., goal, obstacle, or robot) is obtained quite straightforwardly. Any of the objects may perform spatial actions, and in general we speak about *agents* including any kind of object. With this limited ontology we also by-pass the problems caused in perceiving objects, as discussed in Section 4.5.

5.1.2 No Communication

The robots in our teams do not communicate directly, and consequently have no means of negotiation. In this respect our teams differ considerably from most software agents: in the research on software agents, communication is one of the central topics [cf. Wooldridge and Jennings, 1995]. Our arguments for retaining this constraint are rather practical. When quite a large number of agents are involved, it is a difficult problem to establish a suitable communication protocol. The agents have to broadcast messages, as they do not know the identities of other agents. They might start all together to broadcast messages and thus disrupt each others' messages. It is a problem to order the broadcasting efficiently. On the whole, the robots may spend a great deal of time listening to communication that is not at all useful to themselves. A further problem is the *resilience* required for the protocol. Resilience means that the protocol will work properly, despite the fact that one or more of the agents breaks off the communication. In a non-resilient protocol, one agent is kept waiting for a reply, while the other agent has broken off. Resilience is essential, since any autonomous agent might suddenly alters its objectives and not complete the initially started communication cycle. Take for example a goods terminal: one robot starts to “converse” with a second robot. During the communication, the second robot has to escape a third, leaving the first waiting for a response.

Many real-life situations do not require any communication, often they even do not allow for communication. In shopping centres, people pushing their shopping trolleys, generally pass each other without communicating [Fujimara, 1991]. In air traffic there is obviously no time to negotiate. Zeghal and Ferber [1993] note "*In fact negotiation does not enable sufficient co-ordination speed (...), and the co-ordination is not resilient enough (...)*" In marine applications this also holds true.

5.1.3 Autonomous Agents

Due to the exclusion of communication and the spatial orientation, our studies are also applicable to biological organisms. Quite a number of animals use spatial navigation strategies which can be captured in geometrical procedures. For instance, Kanzaki [1996] describes the odour-based navigation of silkworm moths in the search for a mating partner. Crickets find a mate by orientation to the chirping song [Webb, 1995]. Many more examples are described in literature, refer further to [Mataric, 1995] and [Trullier and Meyer, 1997]. Also animals serve as a paradigm for robotics; by building so-called *animats*, i.e., simulated animals or real robots researchers aim to model animal navigation into computer programs and robots [Trullier and Meyer, 1997].

The strict spatial orientation is characteristic for sensory-based robotics. The definition of autonomously performing agents, given in Chapter 1, can be further refined for spatial behaviour. An autonomous agent is an agent that, starting in some position and being given a goal, is able to reach that goal without being externally guided. Moreover, by restricting the domain the construction of the robot is simplified. This opens the possibility to use several of them forming a team; the design of such teams is the object of study in this chapter.

Natural multi-agent systems, such as an ant colony, consist of many autonomous individuals. The interaction amongst the individuals result in complex team behaviour. These interaction mechanisms of social organisation are not yet understood [Mataric, 1995]. Below, we try to find some clues in the restricted robot domain.

5.1.4 Intentions of Agents

We differentiate agents by their intentions. An intention captures the type of (re)action an agent performs because of the presence of objects (or other agents). Whenever we can distinguish a certain pattern (of spatial characteristics) in the (re)actions of the agent, we say that the agent *has an intention* behind its (re)action. Examples of such intentions are shown in Figure 4.4, where two robots aim for a goal; we can say that they have an intention for the goal. When an agent reacts to another agent it can also show a certain intention: it might avoid the other agent or on the contrary chase it. A formal definition of an intention in geometrical

terms will be given in Section 5.3, after we have specified interaction patterns. An intention, by our definition, captures a pattern of actions. Intentions are different if the patterns of actions with which they are associated are different. However, an agent can also have an intention even though it does not perform the typical pattern of actions. For the domain of artificial agents we stipulate that (software) agents provided with the same program potentially have the same intentions. In this respect we diverge from approaches in software agent development where intentions are often understood as a set of (planned) tasks to be executed [Georgeff and Ingrand, 1989].

To summarise, the domain of our study has a limited ontology, within which we define when an autonomously performing robot has an intentions. Goal finding and obstacle avoidance are the main drives behind the intentions. Moreover, intentions are the basis of the interaction within a team, and they have a considerable impact on the overall performance of the team.

5.2 Team Behaviour

In Chapter 3 we discussed a cargo terminal, where a team of low-level robots performed a sorting task, which consisted of spatial actions only. Teams of low-level robots can establish seemingly complex tasks. Bay [1995] discusses also an example of a team of co-operative robots. In this case the team carries a (freight) pallet. Agents performing low-level spatial actions might also be implemented within the aircraft domain. The team behaviour supports air traffic management as described in [Zeghal and Ferber, 1993].

The aim of this section is to show some overall characteristics of team behaviour. We first briefly look at the distribution of control in the team. We distinguish different team members by the intentions which they have and then we evaluate the behaviour of teams, in which all members have the same intentions. After that we add to the team a newcomer which has a different intention.

5.2.1 Distributed Control and Autonomy

To control and guide the behaviour of a multi-agent team, two approaches are possible: a *centralised* and a *distributed* one. In a centralised approach, a central controller (usually called central planner) plans all actions for all agents. The central planner deliberates about a plan for the whole team and checks its validity and completeness [Svesta and Overmars, 1998]. In the distributed approach the individual agents act on their own, and have to find their way quite pragmatically.

Central control is associated with path planning. In multi-robot applications path planning is called dynamic motion planning [Latombe, 1992]. In a large team, the

planning process is very complex, as it has to lay out the course of each robot-agent in detail. The approach is suited to environments in which all factors are fully under control. However, if agents slightly adapt to the behaviour of others and deviate from their prescribed courses, the team behaviour becomes intractable. In our teams of adaptive and autonomously performing robots we do not have much choice but to choose the distributed approach.

In our teams there is no direct central control. All necessary guidance of the team has to be pre-programmed in the individual agents. Team behaviour must be embedded locally in the robots' intentions. In order to design such teams insight is required into the effect each robot has on the team's behaviour. The difficult point is the interaction emerging among the agents.

5.2.2 Different Intentions

In Subsection 5.1.4 we argued that intentions capture a pattern of actions. We can distinguish agents by these patterns, and thus by the intentions they have. The basic distinction is between agents which react to obstacles and those which do not. Hence in Table 5.1 we have two main groups, *reactive* agents which react and *non-reactive* agents which do not react.

Table 5.1: Classification of agents by their intentions.

Non-reactive agents	stationary agents	
	blindly moving agents	
Reactive agents	avoiding agents	
	antagonistic agents	
	non-selective	aggressive
		non-aggressive
	selective	aggressive
		non-aggressive

Non-reactive agents are subdivided into stationary and moving agents. *Stationary agents* – it is natural to call them *obstacles* – do not move²¹. They can be immovable objects but they might also be agents that for the time being do not move. Non-reactive but moving agents are called *blindly moving*. They are agents that move without (noticeably) reacting to other objects, they are insensitive.

Reactive agents form the second main group. We subdivide the group according to whether the agents move toward the object or on the contrary move away. *Avoiding agents* are defined as those which move away from objects. Avoiding agents

²¹ In fact stationary agents cannot have intentions in the sense of our definition, but their non-actions can be considered as intentions.

typically set out for a certain destination and while doing so cautiously avoid contact with others. Pollack *et al.*, [1987] refer to them as *cautious* agents. Quite the opposite behaviour is performed by *antagonistic agents*. Antagonistic agents are avoiding agents but they also aim at a *moving* target, therefore they need another agent as their target. In other words, antagonists are chasing other agents. Antagonistic agents might differ in that they can be *selective* or *non-selective* with respect to their target. *Non-selective* antagonists chase any other agent. *Selective* antagonists chase a particular agent. We remark that an agent can be selective only if it is able to recognise and select its target. Of course, all kinds of intermediate cases of more or less selective agents are possible. For instance, they may chase a certain type of agent. We consider here the extremes only: selective antagonists chase one particular target, non-selective antagonist chase any other agent.

Antagonistic agents may be further distinguished as *aggressive* and *non-aggressive*. Aggressive agents go straight to their target. Antagonistic aggressive agents typically occur in military applications; extreme examples of aggressive antagonistic agents are modern cruise missiles. For our purposes the aggressive agents are not very important: after a successful hit there is no more team. Of interest to our studies are the non-aggressive antagonistic agents. These agents chase their target but approach it only until some threshold distance.

5.2.3 Congenial Teams

In a team, different agents with different intentions might be applied in combination. To characterise multi-agent teams, one usually distinguishes between *homogeneous* and *heterogeneous* teams, see for instance [Fujimara, 1991] and [Noreils, 1992]. Fujimara [1991] gives the following descriptions of homogeneous and heterogeneous teams. In a homogenous team "(...) all robots are identical in the sense that they are operated under an identical set of simple rules." Heterogeneous teams consist of "a set of agents with different capabilities with respect to their sensors, planning algorithms, etc." The terms homogeneous and heterogeneous are relative, and depend on which characteristics are taken into account and how much deviation is accepted. In our further investigation we should be more precise.

We assume that all our agents operate with roughly the same cognitive basis (as defined in Chapter 4) and the same ontology (see 5.1.1). This means that all agents can signal the same objects. In this respect they are homogeneous. However, the agents might have different intentions. We define a team of agents as *congenial* if they have the same intentions and in other respects they are (nearly) homogeneous. The latter means that they might look quite different, but the actions which they perform are comparable, that is, the sensors and actuators, and thus the spatial movements are very similar. Congenial agents have the same intentions, in terms of design they are robots provided with the same programs. Congenial agents ap-

ply in the same situations the same actions, thus in a team they can be substituted for each other.

Below, we examine congenial teams and give the major tendencies in their behaviour. The teams might be composed of any of the types of agents described above. The teams develop behaviour that is typically generated by the implemented intentions. An overview is given in Table 5.2.

Table 5.2: Teams of congenial agents.

Congenial team	Team Behaviour
Stationary agents	non-dynamic
Blindly moving	Chaos Or path planning
Avoiding	Scatter
Antagonist non-selective	cluster together
Antagonist selective	cluster around one or ring of agents

A congenial team of *stationary agents*, that is, a team of obstacles, is not very exciting for studying interaction, nothing happens and it has no dynamics. *Blindly moving* agents, whether or not acting in a team, just start and go. A team of blindly moving agents typically requires centralised control and dynamic path planning, otherwise the team will end in complete chaos. Moreover, in a large team, extensive path planning must have preceded any action; small deviations from planned courses will accumulate and lead to disorder, as we have already discussed in 5.2.1.

Avoiding agents try to avoid and get away from others. Thus, a congenial team of avoiding agents tends to scatter. Generally the agents also aim for a (fixed) goal; the positioning of the goals might force them to approach each other. Thus, goal finding and avoidance behaviours interfere, and the agents can arrive in conflicting situations, we discuss this later. We start concentrating on the avoiding behaviour. The notion of an autonomous robot is commonly associated with that of an avoiding agent.

As stated before, *antagonists* avoid obstacles and chase for a target agent. *Non-selective antagonists* chase for any target agent. When they get close to the target they maintain a certain distance. It is remembered that we consider only non-aggressive antagonists. Thus in a congenial team non-selective antagonists cluster together. Since they are chasing each other, the team shows cohesion. Team cohesion has interesting applications as we will see in examples later on.

Selective (non-aggressive) *antagonists* chase for a specific target. The target may differ from agent to agent, the composition of teams of selective antagonists form may differ accordingly. We discuss two extremes. The first extreme is a team in which all antagonists focus on the same target agent, in the other extreme each antagonist is itself the target of another agent. In the first team all antagonists chase the same agent and the team will cluster around this target agent. In the second team each antagonist chases a different target. The team will sort out such that each is close to its target while avoiding others. The antagonists will form a ring.

Table 5.2 indicates which situation a team of congenial agents tends to establish. Homogeneity in a team provides a starting point for predicting team behaviour. Stationary agents remain at their places, the start and end situation are the same. Blindly moving agents behave according to the planning, otherwise chaos results. The reactive agents in our teams perform goal finding and obstacle avoidance. As a result of this, avoiding agents scatter and antagonist cluster together.

5.2.4 Newcomer in a Team

The congenial teams studied above, are homogeneous and which makes it easier to study them. For a heterogeneous team, it is questionable whether different intentions are consistent with each other. Moreover, it is doubtful whether we can get any hold on the general behaviour of heterogeneous teams. Below, we investigate slightly heterogeneous teams, by adding one single newcomer agent to a congenial team. In Table 5.3 the rows give the congenial teams of Table 5.2. The columns of the table show which *newcomer* is inserted in the congenial team, and indicate its behaviour. The behaviours of the congenial teams are on the diagonal.

Table 5.3: Teams of congenial agents with a newcomer.

Congenial team	Newcomer agent				
	Stationary	Blindly moving	Avoiding	Antagonistic non-selective	Antagonistic selective
Stationary agents	non-dynamic	Pre-planned path	smoothly going through	smoothly going through (stops at one)	smoothly going through (stops at one)
Blindly moving agents	new planning	Chaos or path planning	reactive avoidance	reactive avoidance	reactive avoidance
Avoiding agents	adapts	Moves straight through	scatter	looses target	stays close to one agent
Antagonist non-selective	adapts	Caught at once	flee	cluster together	joins cluster
Antagonist selective	adapts	leader	leader	joins cluster	cluster around one, ring of agents

The first row gives an environment of *stationary agents* which by itself of course has no dynamics. Adding one more stationary agents or obstacles does not change much. However, newcomers of the other types do move around. A blindly-moving

newcomer must apply path planning. The stationary environment is a typical setting for that: there is hardly any interaction. The avoiding and antagonistic agents will equally smoothly pass through this environment: they apply obstacle avoidance. There is little interaction in the team and the different intentions of the newcomers are hardly observable, thus the row is nearly uniform. A slight difference is observed when a newcomer antagonist aims for one of the obstacles as its target.

Row two shows a *blindly moving* team. In this team all agents proceed along the pre-planned paths. When adding a stationary object in column one, the planning needs to be reconsidered. When a moving agent (columns 2-5) is added to the team, there is a choice between two possibilities: adapting the team or simply ignoring the newcomer. The first, to adapt the whole team to the newcomer, means that fully new paths need to be planned. Obviously, this is the best choice when the newcomer is blindly moving as well (column 2). However, if the newcomer is a reactive agent (columns 3-5) it is nearly impossible to set up a plan. Hence, only the second possibility remains: the (reactive) newcomer is simply ignored: the blindly moving agents continue as planned. The newcomer has to avoid the blindly-moving agents and finds itself in a rather discomforting situation. The newcomer avoiding agent (column 3) will give way to the traversing objects. The newcomer antagonists (columns 4 and 5) will keep away from the objects while chasing one or several of the moving objects.

A congenial team of *avoiding agents* in row three, will scatter. Since the agents avoid obstacles, they will adapt their courses when a stationary obstacle is added. When a blindly moving object is added (column 2) the avoiding agents will give way to it. The team behaviour becomes particularly clear if we assume that the team is at rest; the blind agent crosses the shop floor straight, and drives the avoiding agents out of his way.²² A non-selective antagonist inserted in an avoiding team (column 4) will constantly change his target, and subsequently loses each. To be more successful, he should focus on one particular agent as his target. This case represents a real life problem in the natural world. By sticking together, animals in a herd protect themselves from predators. The predator's problem is to single out his prey animal, in order to focus its actions. A selective antagonist (column 5) indeed singles out his target agent, he will follow it and stay close, thus after a while the avoiding team proceeds to scatter where one agent is followed by the selective antagonist.

²² The case is like the chickens on a farm who are driven away when the farmer passes. Obviously, the blind agent is the one which is most successful in reaching its destination. This is the typical case where the agent with the least knowledge is the most successful as noted in [Fujimara, 1991] and [Penders, 1991].

A team of *non-selective antagonists* given in row 4, will cluster together (column 4). The team's behaviour will not change much with the arrival of a newcomer. The team easily adapts to a new obstacle on the shop floor (column 1). A moving newcomer (columns 2–5) might for a while be chased. However, the non-selective antagonists soon change their interest to others. The blind newcomer will be given way (column 2). The avoiding newcomer will try to flee, and might succeed in escaping the team or become locked in the cluster (column 3). A selective antagonist (in column 5) gets locked or ultimately finds its target and stays close to it (in the cluster).

Selective antagonists in row 5, mutually avoid each other and cluster around one agent or form a ring (column 5). As for all reactive agents, also this team easily adapts to a new stationary obstacle. In columns 2 through 5, the newcomer is moving and the selective antagonists might focus on it. If the newcomer is the common target, the antagonists will follow him, he becomes leading. This is true when the newcomer is a blind agent (column 2) as well as an avoiding one (column 3). (The newcomer agent will be followed, as the rats followed the Pied Piper of Hamelin). A non-selective newcomer (column 4) will chase the selective antagonists and cluster with the team.

So far for the analysis of adding a newcomer to a congenial team. For non-reactive teams we conclude that when a non-reactive agent (an obstacle or a blindly-moving agent) is added the planning has to be redone. On the other hand, when a reactive agent (avoiding or antagonistic) is added to a non-reactive team re-planning is of no use. Reactive teams show much interaction dynamics, and thus they are equipped to deal also with newcomers and no particular provisions have to be made. However, a newcomer might influence the team behaviour considerably.

5.2.5 Other Team Designs

The analysis above started from particular intentions to describe congenial teams (5.2.2) and a particular ontology (5.1.1). In this subsection, we compare our teams and our findings with some other approaches described in the literature.

Other intentions

The classification of intentions as given in Table 5.1 is a slightly adopted version of Penders [1994a]. The present classification covers that of Zeghal and Ferber [1993], which originated in airborne systems. Zeghal and Ferber [1993] distinguish *co-operative*, *non-co-operative*, *passive*, and *antagonistic agents*. They do not explicitly distinguish between reactive and non-reactive agents. The term *antagonistic agent* is theirs', however they are not explicit whether antagonists are reactive. The term *passive* agent is used by Zeghal and Ferber [1993] for blindly moving agents. We expect the term *passive* to give rise to confusion with stationary agents. Zeghal and Ferber [1993] describe avoiding teams as well, however they

call avoiding agents *co-operative*. This term is misleading, since the agents are not co-operative, they do not even have a notion of another agent. They are only *self-interested*, i.e., try to satisfy their own constraints. That avoiding agents can co-operatively perform a task as on the cargo terminal, is incidental. One could say that they co-operate despite themselves. In fact, as we have seen in Chapter 3, their actions are co-ordinated. Zeghal and Ferber [1993] use the term *non-co-operative* for agents which differ from the other team members. All the newcomers in Table 5.3, for instance, would in their terms be called non-co-operative.

Kube and Zhang [1994] also study a multi-agent team in which the agents perform spatial behaviours that count as intentions in our sense. The behaviours are: *avoidance*, *goal finding*, *slow*, *follow*, and *find*. *Avoidance* and *goal finding* are the same as those in our teams. The *slow* of Kube and Zhang [1994] results in our teams as a combination of avoidance and goal finding. The *follow* is made by substitution of a fixed goal by a target agent. *Find* is just going around without a fixed goal (we will encounter an example in Section 5.5). We conclude that these behaviours are indeed covered by our intentions.

Mataric [1995] considers the behaviours *avoidance*, *following*, *aggregation*, *dispersion*, *homing*, and *wandering* as a basic set. *Avoidance* and *following* are discussed. *Wandering* is going around without a fixed goal (as the *find* of Kube and Zhang). *Avoidance*, *following*, and *wandering* are intentions in our sense. However, Mataric [1995] also distinguishes *aggregation* and *dispersion*. Aggregation is the same as our clustering and dispersion is similar to scattering. The latter two behaviours are team behaviours, and as we have seen in the discussion following Table 5.3, these behaviours result from avoidance and goal finding, depending on the intentions of the team members. Contrary to [Mataric, 1995] we conclude that clustering and scattering (and thus aggregation and dispersion) are not basic behaviours nor intentions of individual agents; they are resulting team behaviours.

Human agents

A human agent is the paradigm of an intelligent agent. Human beings flexibly switch from one intention to an other. However, in severely constrained circumstances spatial conflicts occur in human interaction too. We make a slight investigation into human interaction and discuss some examples. This also indicates how our study of artificial teams relates to real life situations.

In Figure 5.1 two situations are sketched. On the left a bridge on which it is impossible for two agents to pass at the same time. The right part of Figure 5.1 shows a road crossing where several agents arrive at the same time. In these examples normal traffic rules are not decisive. The situations give rise to conflicts, which are difficult to resolve rationally, even for humans. In such cases humans

apply other skills, such as boldness.²³ It is noted that communication does not always solve these conflicts either, which confirms our presumptions concerning communication (5.1.2)!

Figure 5.1: Left: bridge; right: road crossing.

Other Ontologies

In our study, as well as in the ones of Zeghal and Ferber [1993], Kube and Zhang [1994], and Mataric [1995] the same ontology is used. The agents are in *direct competition* for a resource. This resource is space and the points or regions in space may or may not be occupied. Moreover, the agent themselves are the objects that occupy the space on the shared shop floor. The interaction emerges because only one agent at a time can occupy a spatial position; the interaction generates within a spatial ontology.

Cesta *et al.* [1996] discuss multi-agent systems operating in a slightly different ontology. Their agents are in permanent need of a certain amount of *supplies* – by Cesta *et al.* [1996] called energy. The supply level of an agent changes (decreases) as a side effect of its actions. At certain places agents can find new supplies. Cesta *et al.* [1996] make stepwise simulations of the team. The agents, one after the other, decide what action to take: to move, to confiscate the supplies or to give them away. Contrary to our teams, in the teams of Cesta *et al.*, no interaction generates, the agents does not directly deal with others. To make the simulations comparable to ours, the model of Cesta *et al.* [1996] would have to allow cases where two (or more) agents concurrently aim for the same portion of supplies. A comparison of the agent types further reveals the importance of the initial ontological assumptions. Cesta *et al.* distinguish four types of agents: *solitary*, *parasite*, *selfish* and *social* agents. Solitary agents actually ignore others, in our teams this is not

²³ Near our laboratories there is a traffic road with a narrow bridge, and it is rumoured that two drivers turned to physical violence in determining who was allowed to pass first.

possible: the space on the shop floor is shared. Selfish agents take the supplies which they need (our agents are of this type). The characters social and parasite agents are defined on the supply dimension; social agents share supplies with other, parasite agents take supplies from other. It is clear that these agents cannot appear in our teams (and ontology).

5.2.6 Reviewing Team Behaviour

Above we defined intentions of an agent in terms of programmed spatial actions. The intentions characterise the regularity in an agent's reactions to other agents. In a team, the intentions of the individual agents evolve into team behaviour. Table 5.2 and Table 5.3 give rough indications of the eventually resulting team behaviour. Teams of non-reactive agents behave according to the planning made in advance. However, teams of reactive agents generate interaction. Thus, before the avoiding agents scatter and the antagonists clutter, much interaction has taken place. The interaction causes dynamics in the teams. Interaction and dynamics of teams are the subject of this chapter. So far we can conclude that since antagonists aim at a moving goal, these teams show the most dynamics. A team of avoiding agents with a stationary goal acts more calm. However, the dynamics of the behaviour of an avoiding team is illuminating for all reactive teams. Antagonists are also avoiding agents, in this respect antagonistic teams have much in common with avoiding teams. However, the chase for moving goals introduces an extra dynamic aspect, thus making it complicated to capture the dynamic behaviour

5.3 Interaction: Collisions and Conflicts

In this section we develop formal tools to describe interaction for teams of congenial avoiding agents, moving in a two dimensional plane. The robots perform goal-finding and obstacle avoidance. But, as we have seen in the previous section, the robots interact with each other and there is no guarantee that they eventually will arrive at their goal. As an illustration, and before turning to the technical work, we first give an example of the sort of problems we will deal with.

Figure 5.2, which is taken from [Wang, 1989] shows the trails of three congenial avoiding autonomous robots. The figure indicates the start positions and the targets of the robots. Robot 1 succeeds to reach its goal, but robots 2 and 3 fail, they run into the insoluble conflict of avoiding one another. Latombe [1992] calls such conflicts simply oscillatory phenomena. It is our opinion that studying these conflicts reveals some of the essentials of team behaviour, as they lead to interaction.

To analyse these situations we will in this section distinguish various geometrical interaction patterns. Moreover, we formulate and proof a Theorem 5.1 on how to

apply the interaction patterns when studying teams. After that we go, in Sections 5.3 and 5.4, into detail with congenial avoiding teams, and show that in several cases the behaviour of an avoiding team can be captured in mathematical models.

Figure 5.2: Trajectories of three congenial robots.

5.3.1 Spatial and Temporal Preliminaries

For clarity we reiterate the major constraints on teams of agents. We assume minimal communication in the team. The robots cannot communicate, let alone negotiate with each other; there is neither a communication link present, nor a (permanent) link to consult a central controller. Once having received its instructions, a robot is assumed to perform on its own. The robots can perceive objects, i.e., obstacles and other agents. Each robot has the same kind of information about its environment, i.e., they have the same cognitive bases.

Comparing this chapter with Chapter 4 robots are treated differently. In Chapter 4 real robots were of concern, here we make abstractions. We treat robots as *free flying objects* [Latombe, 1991]. This means that the robots can move freely in the two-dimensional plane, and are not bounded to a grid pattern or the like. We do not take kinematics constraints into account. The robots are considered to be *point robots* [Fujimara, 1991], which means that they have no spatial extension and thus the robots compete in the plane about points only.

We observe that the team behaviour and interaction patterns can only be described from the standpoint of an external observer. Our approach is to go stepwise from

situation to situation, and analyse how each agent reacts to the situation. Such analysis is usual in computer science, in particular in studying *interprocess communication* (see, e.g., [Dolev *et al.*, 1987]), but had to be adapted for spatial applications. A difficulty is that in interprocess communication, actions of agents occur in a temporal sequence, whereas in spatial interaction all agents can react at the same point in time. Nevertheless, the stepwise approach simplifies the mathematics required and reveals the overall characteristics of our teams.

To deal with the temporal aspects of interaction, we introduce *points in time* (or moments) which are common to all robots; which is equivalent to saying that their clocks run synchronised. At the (common) points in time robots might change their action. An action chosen at a point in time is sustained during the subsequent period. When the robots move they use maximum speed all the time, unless they are explicitly instructed otherwise.

Interaction patterns are extended in time and are represented in *time frames*, which are sets of subsequent time points. The course of events has only one (future direction, hence a *linear time* model is appropriate. To have sufficient points of time available to represent the crucial events, we need at least a dense model of time (dense means isomorphic to the rational numbers).

Notation

D is the overall domain of objects and includes all agents and robots. Each agent i is able to observe any of the other agents $j \in D$, (i.e., they have equal cognitive bases). Thus we can speak of a subset $d \subseteq D$ of objects to which each agent is in principle sensitive. This, however, does not mean that the domain is fully known in advance. The number of objects present in a situation might vary in time. We have to deal with *appearing*, *disappearing* and *temporary* objects [Fujimara, 1991]. Appearing objects are objects which initially are too far away and therefore not seen yet, but which may enter the situation later on. Others move away from a situation, they disappear, as robot 1 in Figure 5.2. Thus, for each point t in time the set d of objects relevant at that time point t is fixed, however the set might vary with the time points.

We use the following symbols:

D	overall domain of objects and robots
d	$d \subseteq D$ set of relevant objects
i, j, o	variables to denote objects or agents
$(\mathbb{R} \times \mathbb{R})$	the two-dimensional plane
V_{ixy}	position of robot or vehicle i
G_i	the goal of robot i ; the positions of the goals are usually fixed
p, q	variables reserved for points in the plane

$\text{dist}(p,q)$	(Euclidean) distance between point p and q in the plane
$[p,q]$	line segment between p and q
$Ob(i,o)$	robot i observes object o
S_i	actual speed of robot i ;

The usual logical connectives $\wedge$, $\vee$ and $\rightarrow$ as well as the quantifiers $\exists$ (exists) and $\forall$ (for all) are used in the object language, to describe the situation at a certain point of time. In the meta language, when speaking about temporal relationships we use the words **and**, **if**, and **iff** (if and only if). For example we can further specify the predicate $Ob(i,j)$:

Definition 5.1.

$$t \models Ob(i,o) \text{ iff } t \models V_o \in H_i \text{ for some given horizon } H_i \subseteq (\mathbb{R} \times \mathbb{R}).$$

At time t , agent i observes an object, if and only if the object is within the horizon of i . The horizon is that part of the plane which i can observe. The Figures 4-5 and 4-6 show examples of specific horizons. In any case, horizons might differ between the agents, however for convenience of the discussion in the present chapter, we usually assume that all agents have the same horizon. An exception is the flee-opposite procedure discussed in sections 5.4 and 5.5. Moreover we assume that the horizons coincide with the whole plane, $H_i = (\mathbb{R} \times \mathbb{R})$ for every agent i . Consequently the predicate $Ob(i,j)$ is symmetric: $Ob(i,j) \leftrightarrow Ob(j,i)$.

5.3.2 Interaction Patterns

In this subsection we distinguish and define certain interaction patterns. We start by singling out certain special situations which may lead to a *collision*. A team of congenial avoiding robots is quite dynamic. During the interaction certain *incidents* may occur because of which the agents change the course of events such that no collision occurs. We introduce the dynamic notion of a *conflict* to denote any course of interaction. A conflict is a series of situations which may or may not lead to a collision. Given this ambiguity we have to elaborate on conflicts and define when they end. At the end of the section we have sufficient tools available to state our formal definition of *having an intention*, as was promised in Subsection 5.1.4.

Events

Definition 5.2. A *situation* is the positioning of agents and objects at one point in time.

For convenience (and as long as no confusion is expected) we identify situations by the points of time at which they occur.

Certain situations contain a threat towards a collision. We distinguish two types of collisions called *real-collisions* and *semi-collisions*. A situation contains a *real-collision* if starting from this situation, the robots have to cross one particular point (p) in space, at the same (future) point of time. However, this prediction works only when the robots maintain their current speeds S_i and S_j . Figure 5.3 shows a real-collision of two robots A and B.

Definition 5.3. *Real-collision* ($RC(d)$).

$t \models RC(d)$ iff

$t \models \exists p \forall i \in d (p \in [V_i, G_i])$ and $t \models \forall i, j \in d (\text{dist}(V_i, p) / \text{dist}(V_j, p) = S_i / S_j)$.

Figure 5.3: A real-collision situation of two robots.

From the real-collision situation onwards, the speeds of the robots need not be constant, each robot might change its speed. If the speeds change relative to each other, the collision will not occur. Thus we may therefore drop the clause about speeds, and obtain the weaker notion of a semi-collision. In a semi-collision the courses of the robots intersect, however the robots can be at arbitrarily distances from the intersection point (no relation to the speeds is required). For instance in Figure 5.4, robot A nearly passes p while robot B is due to arrive much later.

Definition 5.4. *Semi-collision* ($SC(d)$).

$t \models SC(d)$ iff $t \models \exists p \forall i \in d (p \in [V_i, G_i])$.

Figure 5.4: A semi-collision situation.

In definitions 5.3 and 5.4 we mainly focus on the goal finding behaviour of a robot. However, the robots also avoid each other. When combining avoidance and goal finding, the intersection point p is not fixed throughout a series of situations. Therefore, in real life applications semi-collisions are very often identified as dangerous. Automated collision-avoidance systems applied in airborne traffic, treat real-collisions and semi-collisions similarly (see [Zeghal and Ferber, 1993]). In marine applications the point p is called a *virtual collision point*. Zhao *et al.* [1992] review maritime applications, and use the *distance at closest point of approach* (DCPA) and the *time to closest point of approach* (TCPA). Note that both notions only make sense in semi-collisions.

We have identified real-collisions and semi-collisions, but reactive agents start to avoid and thus interact, even when there is no indication of a semi-collision. To capture this feature, we introduce the notion of incident. In a certain situation, an agent has an *incident* with the objects in set d if it is sensitive to the objects.

Definition 5.5. *Incident* ($I(i, d)$). $t \models I(i, d)$ iff $t \models \bigvee_{j \in d} (Ob(i, j))$.

The notion of incident is relative to an individual agent. Thus, in a team of agents every agent has its particular incidents. Mutual or reciprocal incidents are incidents caused by and observed in the same situation by agents who are mutually involved. In a mutual incident all agents observe each other, and we denote a mutual incident as $I(d)$.

Definition 5.6. *Mutual incidents* ($I(d)$).

$$t \models I(d) \text{ iff } t \models \bigwedge_{i \in d} I(i, d') \quad \text{where } d' = d \setminus i: d' \text{ is } d \text{ but } i \text{ is excluded.}$$

Above in Definition 5.1 we introduced the assumption that our agents all share the same horizon which spans the whole plane, $H_i = (\mathbb{R} \times \mathbb{R})$ for every agent i . Thus we also have $Ob(i, j) \leftrightarrow Ob(j, i)$. The consequence is that in the discussion below inci-

dents always coincide with mutual incidents. However, agents might have different horizons, for instance a blind agent has an empty horizon. For teams in which agents do have different horizons the difference between incidents and mutual incidents is crucial, refer to Section 5.5 and [Penders, 1991] for more examples.

Interaction Patterns

Now that we have identified several particular events, we proceed to model interaction patterns. Basically we need so-called time frames, that is series of subsequent points in time. Technically a frame $\mathbf{T} = (T, <)$ is a set of time points T with an ordering $<$. Of course, to model robots there should be a particular relation between the time point in our frames. The objects in our set d move from one situation to another, however, they must follow a continuous path; they cannot jump from one position to another.

Definition 5.7. [Latombe, 1991] A *path* of an agent from position (x_1, y_1) in situation1 to (x_2, y_2) in situation2, is a continuous mapping τ :

$$\tau : [t_1, t_2] \rightarrow (\mathbb{R} \times \mathbb{R}), \text{ with } \tau(t_1) = (x_1, y_1) \text{ and } \tau(t_2) = (x_2, y_2).$$

Definition 5.8. *Related situations.* Two situations are (time-)related if there exists a path for each agent $i \in d$, from the first situation to the second situation.

As a consequence of Definition 5.8, related situations have so-called constant domains: in two related situations t and t' the same agents occur; hence if a new object turns up the situations are not related.

We now have the tools available to define the notion of a conflict, which is central to our treatment of interaction. An *interaction pattern* obviously must be a time frame consisting of incidents (Definition 5.4) in related situations (Definition 5.8). *Conflicts* are interaction patterns where the agents mutually interfere with one another's aim to reach the (respective) destination point(s).

Definition 5.9. A *conflict* is a time frame $\mathbf{T} = (T, <)$ of related situations with a fixed domain d , such that at every point of time t in the frame there is a mutual incident.

$$\mathbf{T} \models \text{Conflict}(d) \text{ iff } \mathbf{T} \models I(d) [t] \quad \text{at all time points } t.$$

Conflicts end when anyone of the mutual and related incident conditions is violated. Examples at hand are cases where objects appear and disappear from the horizon (Definition 5.1). In such cases the incidents are no longer related (Definition 5.8). A conflict ends *naturally* if at a certain point of time $t \in T$ at least one of the robots is no longer obstructed by any of the others, which means that – under further normal conditions – it would reach its destination. For convenience and

when no confusion is expected, we leave out the qualification 'natural'. A certain point of time t is the *Ending of Conflict*(d):

Definition 5.10. $t \models \text{Ending of Conflict}(d)$ iff $t \models \exists i \in d (\text{dist}(V_i, G_i) = 0)$.

Strictly speaking Definition 5.10 demands that an agent has to come to a standstill exactly at the position of the goal. In general we have to admit some tolerance. We could replace $(\text{dist}(V_i, G_i) = 0)$ by some criteria which says that agent i has to enter a certain neighbourhood of G_i . Another way around this is to allow the goal to have (spatial) extension. We will not introduce these options here, as it does not affect the general idea we want to sketch. However, it is important to point out that if the agent overshoots its destination point (or the defined neighbourhood) we stipulate that after the overshoot we are dealing with another conflict.

In general a robot involved in a conflict does not approach its destination point straightaway. As we have seen in Figure 4.4, due to the interaction with other robots, a robot might make all kinds of enveloping movements. Also, at a certain point in time it might be rather close to its goal while some moments later it can again be further away. The robots in Figure 5.2 seem to do so. Whatever the case, a conflict has an end if (in the long term) a series of time points exists in which one robot really approaches its goal. This robot must have a path, which is continuous (Definition 5.7), and that gets arbitrary close to its goal. We define *targeting series*, as a sequence of those time points at which the robot really approaches its the goal.

Definition 5.11. Let τ be a path defined on time frame $T = (T, <)$, such that at time t_0 robot i is at distance c^* from G_i and at time t_e it is at G_i , then *Targeting series* $TS(i, G_i)$ consist of time points between t_0 and t_e such that for every $c, c^* > c > 0$

$$T \models \text{dist}(V_i, G_i) < c [t] \quad \text{for some } t \in TS(i, G_i)$$

$$\text{and for } t, t^* \in TS(i, G_i) \text{ if } t < t^* \text{ then } \text{dist}(V_i, G_i)[t] < \text{dist}(V_i, G_i)[t^*]$$

Targeting series are only useful if they indeed characterise ending conflicts. We prove this in the following theorem.

Theorem 5.1. A *Conflict*(d) has a natural ending iff it contains a targeting series $TS(i, G_i)$ for some robot $i \in d$.

Proof.

$\Rightarrow$ when a conflict has a natural end, for some i , $\text{dist}(V_i, G_i) = 0$; a targeting series with a limit of $\text{dist}(V_i, G_i) = 0$ is easily constructed.

$\Leftarrow$ if a conflict contains a targeting series for robot i , robot i will approach its goal arbitrarily close, that is: the limit of $\text{dist}(V_i, G_i) = 0$. $\square$

Many conflicts end in the long run, though at first sight it is not clear. A targeting series, according to Theorem 5.1, provides a characteristic to recognise ending conflicts. However, there are conflicts that do not have an end, as already was illustrated by Figure 5.2. We call them balanced conflicts.

Definition 5.12. A *balanced conflict* is a conflict that does not have a natural ending.

A conflict is balanced if from some point in time t^* onwards, none of the robots approaches its goal any closer, that is the conflict does not contain any targeting series.

Lemma 5.2. A time frame $T = (T, <)$ contains a balanced *Conflict*(d) iff

$$T \models \forall i \in d(\text{dist}(V_i, G_i) \geq c > 0) \quad [t^*] \text{ for a certain time point } t^*$$

$$\text{and } T \models \forall i \in d(\text{dist}(V_i, G_i) \geq c) \quad [t] \text{ for all } t, t^* < t.$$

Proof. According to Theorem 5.1, a balanced conflict has no end iff it contains no targeting series, and that is what Lemma 5.2 says. $\square$

5.3.3 Reviewing Collisions and Conflicts

In Chapter 4 we referred to a cargo terminal on which robots were to be applied. However, a balanced conflict may withhold a team from making any progress. If the robots set out for a task, but half way get entangled in a balanced conflict, they obviously never accomplish their task. The designer of a team, of course, seeks to avoid balanced conflicts.

Theorem 5.1 and Lemma 5.2, are the main tools for deriving mathematical characteristics of the performance of a team. The use of Theorem 5.1 is that subsequent proofs can be restricted by taking into account only series of relevant situations. Hence, it is not required that a proof completely outlines the courses followed by the robots involved. For instance, in Figure 5.2 balance might be proven by showing that Robot 2 and 3 do not enter a certain region around their targets. Yet for showing that a zigzagging robot reaches its goal, it is sufficient to show that the zigzagging course gets arbitrarily close to the target.

5.3.4 Intentions Defined

Above we defined a conflict as a pattern of mutual interactions between two or more robots. Using patterns of interaction, we can formally define the notion *having an intention*, as promised in 5.2.2. Intentions concern the behaviour of a single agent, no matter how others react. An intention is applied in a series of related situations (Definition 5.7). Moreover, an intention designates peculiarities in

the behaviour of an agent. Hence, we use the notion of a targeting series, in particular because it does not require that the agent heads straight for its target but instead it allows some freedom.

Definition 5.13a. Agent i has an intention for target o (at position (x,y)) if there exists a targeting series $TS(i,o)$ for i with o as its centre point.

Definition 5.13a applies for goal finding, but we have also treated avoidance as an intention. When avoiding, the agent moves away from the target object o . A definition of avoidance is obtained by substituting a targeting series with a diverging series.

Definition 5.13b.

Let τ be a path defined on time frame $\mathbf{T} = (T, <)$, such that at time t_0 robot i is at distance c^* from o , then agent i avoids object o on time frame $\mathbf{T}$, if for any $c, c > c^*$:

$$\mathbf{T} \models \text{dist}(V_i, o) \leq c^* [t] \quad \text{for some } t > t_0$$

$$\text{and for } t, t^* \in \mathbf{T} \text{ if } t > t^* \text{ then } \text{dist}(V_i, G_i)[t] > \text{dist}(V_i, G_i)[t^*].$$

Targeting and diverging series also can be applied when the target object O moves. Thus Definition 5.13 also applies in the case of moving objects.

The definition lays down an intention as a pattern of performed actions. There is no claim for completeness of the definition. For instance, in a philosophical context an agent might be said to have an intention even if – because of certain circumstances – the agent is unable to effectuate it. Also, an agent might reach a target without (in the philosophical sense) intending to do so. We make no effort to deal with the cases mentioned. What we have obtained with Definition 5.13 is an indication of what intentions are in a geometrical context. To have an intention for target O , is clearly different from avoiding O . The programming of a robot generates the performed actions and thus generates intentions. Different programs applied on the same robot usually imply different intentions. However note that intentions are not identical with (parts of) programs. Moreover, in the geometric context we can judge only performed actions. Nevertheless, all we need here and in the next section, is that programs and intentions are related. We restrict the discussion to agents with the intentions of goal finding and obstacle avoidance. Later, in Chapter 6 we further specify intentions in a philosophical context.

5.4 Obstacle Avoidance

We study congenial teams and the interaction and team behaviour generated by the intentions: goal finding and obstacle avoidance. Detailed characterisations of interaction, of course, vary with the procedures applied. In this section we discuss three obstacle avoidance procedures: flee-opposite avoidance, potential-field avoidance and manoeuvring-board avoidance. We show simulations of these procedures, they illustrate some of the behaviours which we encountered in Section 5.2. In the next sections, we proceed investigating interaction and team behaviour, the notions of collisions and conflicts will be applied there.

The *flee-opposite* avoidance is discussed first, it is a straightforward procedure basically performing obstacle avoidance without consequent goal finding. This procedure is applied in the TOUCH module of our robot (cf. Chapter 4). Secondly we discuss *potential-field* avoidance. The avoidance based on a so-called *artificial potential-field*, see [Khatib, 1986], and essentially combines goal finding with obstacle avoidance. We study a simplified procedure which is applied in the AVOID module of our robot. Finally, we discuss *manoeuvring board* avoidance, this procedure is particularly designed for avoiding moving objects, and is described in [Tychonievich *et al.*, 1989].

5.4.1 Flee-Opposite Avoidance

The basic move generated by the flee-opposite procedure is straightforward: as soon as the robot observes an object, it steps backwards until the object is out of sight, that is, out of the horizon Definition 5.1. Out of sight means that the object is out of the range of the sensors, see Figure 4-3. Depending on relative position of the object the robot will steer either a little to the left or to the right while moving backwards. After the step backwards is completed, the robot has to resume a course, at this point several choices for heuristic rules are possible. Chattergy [1985] discusses a family of such rules. We elaborate two of them, the results confirm those of Chattergy [1985].

Figure 5.5 gives simulation results, obtained by a robot that resumes a course straight ahead, as if its target is always in front of it. In the simulation, the robot encounters a row of objects on its left and steps back slightly to the right. After a series of repetitive steps it succeeds in passing the objects.

Figure 5.5: Flee-opposite with fixed direction.

The robot in Figure 5.5 is in fact exploring its environment, as it is not fixed on a particular target. We consider it an instance of the *find* and *wandering* behaviours discussed in 5.2.5.

In the simulation of Figure 5.6, the position of the goal is fixed. The robot starts at the same position as in Figure 5.5 and encounters the row of objects. After going back slightly to the right, the robot resumes a course. In this case it drives towards the fixed position of the goal and soon becomes trapped. We may call this a balanced conflict, though a degenerated one since there is only one agent acting. In any case, the simulation shows how goal finding and obstacle avoidance interfere.

The same diagnosis applies in the situations shown in Figure 5.7. In the left situation, contrary to Figure 5.6, the obstacle-avoidance procedure does not guide the robot away from the baseline of [Obstacle,Goal]. When the robot finds an object exactly on the direct course to its target it is also locked in a balanced conflict. In the right situation of Figure 5.7 the robot encounters a corner. It makes a couple of moves, however the net effect is minimal, and the robot is trapped. This occurs typically in sharp corners, when the robot is on the diagonal.

Figure 5.6: Flee-opposite with fixed goal.

Figure 5.7: Left: a stationary obstacle and one robot; right: a trapped robot.

5.4.2 Potential-Field Avoidance

The second avoidance procedure is based on a so-called *artificial potential field* first described by Khatib [1986], an overview and variants are given by Latombe [1991]. The input for the procedure is an overview of the surroundings of the robot. That is to say, the points where the goal and other objects are located. In its calculations, the procedure treats the goal as a source of attraction and other objects as sources of repulsion. At each point of time the new velocity of the robot is determined as the vector sum of these forces, see Figure 5.8.

Figure 5.8: Vector sum in the potential-field procedure.

The magnitude of the repulsive forces decreases as a function of the distances to them. One often visualises this as a landscape in which the repulsive forces form mountains and the attraction forces valleys. The procedure calculates a path of decreasing resistance, i.e., it steers the robot through the valleys in the landscape.

Figure 5.9 shows a simulation of two robots crossing the shop floor, in the middle of which there is an object. The potential-field procedure generates smooth paths. We observe that the procedure applied by Wang [1991] shown in Figure 5.2 is based on a similar procedure.

Figure 5.9 is instructive for our interaction model. Robot 1 fully overviews the shop floor. From the start, it avoids the obstacle and robot 2. A consequence of the overview is the overshoot near the goal 1. Robot 2 is provided with the US-sensor described in Chapter 4. The sensor has a limited range. As late as step 5, it becomes sensitive to the obstacle. Note that the last steps are straightforward towards the goal 2. In the figure several of the distinctions made in our interaction model apply. There is no real or semi-collision in the simulation, only incidents (definition 5.6). For robot 2, the obstacle appears at a sudden, moreover, since robot 2 does not react to robot 1 there is no mutual incident (definition 5.8). Any

case the simulation shows how important it is for interaction, that a team is congenial, in Figure 5.9 there is no interaction between the robots.

Figure 5.9: Two robots applying the potential-field procedure. Robot1 has a complete overview of the shop floor, robot2 has the US-sensor of Figure 4-5.

5.4.3 Manoeuvring-Board Avoidance

The flee-opposite and the potential-field procedures are sensitive and reactive only to the current positions of the objects. However, the procedures might go a step further and anticipate the behaviour of (moving) objects. Procedure doing this are particularly in use in marine applications (see for instance [Zhao *et al.*, 1992]). Though they are not well-suited for robotics application, it is interesting to have a brief look at an example. They illustrate the practical use of the notion of semi-collision (Definition 5.4). Essential to the calculations are the *virtual collision points* which we have encountered above (in Subsection 5.3.2).

The example which we discuss is the *manoeuvring-board* avoidance procedure described in [Tychonievich *et al.*, 1989]. Its basic algorithm is straightforward. Around any observed object a safety circle is constructed to indicate an area into which the robot should not move. The tangent lines from the robot to the circle around the object mark off a danger zone for a stationary object. Travelling into

this zone the robot would, sooner or later, intrude into the safety circle (the *Predicted Area of Danger*, as it is often called in marine applications). However, as the object moves the danger zone should move along with it. Therefore the zone is translated along the velocity vector of the object. In Figure 5.10 the translated zone is specially marked. The new velocity vector of the robot is chosen such that it does not point into the translated danger zone.

Figure 5.10: Manoeuvring-board procedure.

The manoeuvring-board procedure processes velocities of objects to determine the zones of danger. The robot needs to (store and) compare the current situation with a previous one to determine velocities. It has to reconstruct the path of the objects, from the one situation to the other. When doing so, it might run into the so-called *association* or *correspondence* problem: when comparing two situations, which of two pairs of positions from the successive situations (maps) belong to the same object? Figure 5.11 show a puzzling situation. A_0 and B_0 are the positions of two objects at time t_0 . The problem to solve at time point t_1 is which object moved to P and which one to Q , did A move to P or to Q ?

Figure 5.11: The correspondence problem.

In an application of a team of robots the correspondence problem is a serious one. It is caused by the fact that the robots might get relatively close to each other. The objects A and B in Figure 5.11 are so close that by performing a standard move,

they can switch position within the period between two scans. Since this cannot be avoided in robot applications, the manoeuvring-board procedure is not suitable. The problem is avoided if the obstacles never get closer than twice their maximum speeds. This applies in marine applications. Ships are situated in a – relative to their speed – sparsely occupied plane, and are prevented from approaching each other too close.

Another problematic issue for a manoeuvring-board approach is the information collection. The collected data have to be sufficiently accurate to determine a realistic estimate for the velocities of the objects. For the robot discussed in Chapter 4, the low resolution will pose a problem. Because of the low resolution, the possible deviation in the velocity estimates can be more than twice the real velocity.

5.4.4 Reviewing Obstacle Avoidance

We have discussed three obstacle-avoidance procedures: flee-opposite, potential-field and manoeuvring-board. The flee-opposite procedure generates series of repetitive movements, as Figure 5.5 and Figure 5.6 show. The attractiveness of the flee-opposite procedure is its simplicity in implementation. On our robot (see Chapter 4) it requires hardly any processing. In Section 5.5 we briefly discuss a team application of the flee-opposite procedure. The potential-field procedure generates a more smooth path as is shown in Figure 5.9. In Section 5.6 we go at length into the application of the potential-field procedure in a team and we give a mathematical analysis of the interaction that generates in the team. We do not evaluate the manoeuvring board procedure. It is not suitable for a multi-robot application because of the correspondence problem and the requirements of information collection. Moreover, since the manoeuvring-board approach anticipates the speed of objects, it presupposes that the agents will continue their movements, as if they were blindly moving. And, as we argue in Section 5.2, not much interaction generates in a blindly moving team.

In the next sections we explore the interaction in a congenial team of avoiding robots, first we briefly study the flee-opposite procedure and next, in Section 5.6, the potential field procedure. In advance we should stress that we take a rather particular point of view by considering interaction in a team and evaluating the avoidance procedures when applied concurrently. The designers of the procedures probably never have had such a context in mind; our treatment will not do justice to the procedures themselves. However, despite this, the avoidance procedures put us on the way to study interaction.

5.5 Flee-Opposite Interaction

A team of two congenial robots applying the flee-opposite procedure is shown in Figure 5.12. The series (a) and (b) show some successive points in time and the courses of two robots. The series start at time point t_0 , with the two robots at positions A and B . At the time point t_1 , the robots are sufficiently close to perceive each other, then they both flee simultaneously in (strictly) opposite directions (as shown in the figure). At the next point in time (t_2) the robots are too far from each other to perceive one another, so both robots resume moving in the goal direction. The result is a zigzag course of the robots around borders of their horizons, that is the threshold distance at which they are mutually visible. This threshold is the maximum range of their sensors.

Figure 5.12: The flee-opposite procedure, (left): balanced conflict; (right): non-balanced conflict.

In Figure 5.12 (left), the situations at points t_1 , t_4 and t_6 are quite similar. The robots are in (a repetition of) real-collisions: their courses eventually intersect as is clear from elongating their velocity vectors. The situations only differ in that the robots have approached each other at slightly different distances. Both robots repeatedly flee, without making any progress in passing one another. Figure 5.12 (right) is quite different, the robots start at point t_0 in a semi-collision. At point t_6 the robots are out of each others range, from that point of time on both can proceed without conflicting.

The explanation for the different behaviours in Figure 5.12 (left) and (right) is simple. At the successive points we can draw lines connecting the positions of the robots A and B at that point of time, in Figure 5.12 they are drawn as broken lines.

The orientation of the successive lines at the time points t_1 , t_4 and t_6 is important. It is common knowledge among sailors that if, as in Figure 5.12 (left), the lines are parallel to each other the robots will eventually collide. That is the situation of a real-collision. A proof – that in Figure 5.12 (left) the robots will cross the same point at the same point in time, and that there will not be such a point in Figure 5.12 (right) – is easily constructed by similarity in a triangle.

Balanced conflicts also occur in a team of *three* robots. To apply the flee-opposite procedure in a team, we have to define the flee-opposite move when more than one object is detected. We recall the robot design described in Chapter 4: the robot only detects whether there is an object on the right or on the left hand side. Depending on this information, the step back is adjusted left or right. If an object is spotted on the left as well as on the right, it goes straight back. In this indecisive situation the robot takes the middle course. Figure 5.13 shows a balanced conflict for three of such robots.

Figure 5.13: Three robots with flee-opposite.

We will not engage in a thorough mathematical analysis of the flee-opposite procedure. A two robot team is studied in [Penders, 1991]. The typical problems of a multi-robot team application are clearly explained with the procedure discussed in the next section.

5.6 Potential-Field Interaction

In this section we analyse the interaction generated by the simple artificial-potential field procedure. The procedure is applied in teams of two and three ro-

bots. Though our procedure is a simplified version, it indicates from the perspective of our analysis the rather typical behaviour to be expected from the sophisticated procedures of [Khatib, 1986] and [Rimon, 1992] or the procedure of [Wang, 1991] as depicted in Figure 5.2.

The analysis of the interaction in a team of two robots is given in appendix A1. It provides a complete description for a team of two robots. The analysis of a three-robot team is more complicated, and the methods used for a two-robot team apply only partially. Below, we analyse several cases which together indicate the main characteristics of a three-robot team. We discuss a balanced conflict (Lemma 5.3) and a slightly different situation in which the robot team turns (Lemma 5.4). Thereafter we discuss a case where two robots force the third robot forward or backward (Lemma 5.5).

5.6.1 Basics

Prior to beginning the analysis, we provide some conventions and notations. We deal with a fixed set of robots, they form our domain $\mathbf{D}$. Variables r and o are used to range over robots $r, o \in \mathbf{D}$. Sometimes we use the term object for a robot to achieve clearness, we remind that there is no ontological difference. Individual robots are indicated by letters A , B and C ; their respective goals are indicated by G_A , G_B and G_C . Line segments will be denoted by square brackets $[p, q]$.

An object, that is, another robot on the shop floor is represented as a repulsive force $\bar{R}$. The goal of the robot is represented as a force of attraction $\bar{F}$. The forces of attraction $\bar{F}$, and the repulsive forces $\bar{R}$, are indexed:

$\bar{F}_A$: attraction exerted on A ;

$\bar{R}_{AB}$: repulsion on A caused by B .

We determine the *force sum* of the forces. The new velocity of the robot, is a vector which, in our analysis, is identified with this force sum. We note that we do not consider acceleration. The robots move from one point to another with maximal speed, as determined by the force sum. For each robot the forces result from the presence of objects. The team of robots is congenial, thus the parameters will be the same for all robots.

Definition 5.14. *Force sum.* $\bar{F}^* = \bar{F}_r + \sum_{o \in \mathbf{d}} \bar{R}_{ro}$

in which: $\bar{F}^*$ is the basis for the new velocity vector of the robot r ;

$\bar{F}_r$: the attraction of the goal of the robot r ;

$\bar{R}_{ro}$: the repulsive force exerted by object o on the robot r .

The *repulsive* force decays in general with the distance between robot and object and points directly opposite to the other robot's position (see Figure 5.8). The magnitude of the repulsive forces between robots A and B is defined by:

Definition 5.15. $\|\bar{R}_{AB}\| = 1/\text{dist}(A, B)^2$.

Since the robots A , B and C are congenial we have:

Observation 5.16. $\bar{R}_{AB} + \bar{R}_{BA} = \bar{R}_{AC} + \bar{R}_{CA} = \bar{R}_{CB} + \bar{R}_{BC} = \bar{0}$.

Proof. Obvious from definition 5.15 and definition 5.14. $\square$

For the *attraction* of the goal one commonly makes a choice between a force decaying with the distance or a force with a constant magnitude. We deal with constant magnitudes for the attractive forces; it supports better goal finding especially when the robot is far from the goal. The attractive force is exerted on the respective robot and directed to its goal.

Definition 5.17. $\|\bar{F}_A\| = \|\bar{F}_B\| = \|\bar{F}_C\|$, at every point of the plane.

Below we use several terms from elementary geometry. A triangle has a point (centroid) where the medians intersect. For readability we use the common term *Centre of Gravity* to denote this point. In general three robots form a triangle, which has a Centre of Gravity (CG). Similarly the triangle formed by the goals G_A , G_B , and G_C has a Centre of Gravity, we denote it as ECG .

We also use the Steiner Point of a triangle. It is the point that satisfies the following description of Fermat: the point p in a triangle ABC , whose distances from A , B and C have the smallest sum. Any triangle, all angles of which are smaller than 120° , has its Steiner Point inside the triangle. The Steiner Point of the triangle formed by the robots is denoted as SP , and the Steiner Point of the triangle G_A , G_B , and G_C , by ESP . Below, when we discuss the Steiner Point we assume that all angles of the triangle are smaller than 120° .

To summarise, we use the abbreviations:

CG : centre of gravity in the triangle formed by the *robots*;

ECG : centre of gravity in the triangle formed by the *goals*;

SP : Steiner Point of the triangle formed by the *robots*, that is, if all angles of this triangle are smaller than 120° ;

ESP : Steiner Point of the triangle formed by the *goals*, that is, if all angles of this triangle are smaller than 120° .

5.6.2 A Robot Team

A team of congenial robots, applying the potential-field procedure has some similarity with the mechanics of a system of particles. Thus, our study of robot teams resembles studies on the many-body problem of physics. The many-body problem is: “*For a system of N particles with known interactions, determine the evolution of the system from any given initial state*” [Hestenes, 1986]. The many-robot problem is in some aspects more simple. The robot team operates within a fixed two-dimensional plane. Also, we do not consider acceleration of the robots. Moreover, we do not have to determine the complete evolution of the robot team. Theorem 4.1 showed that we can restrict attention to targeting series only, which means that it suffices to study only relatively few points in time (or situations).

A complication is that – as we will see below – the interaction in the robot-team is caused by two system of forces: the repulsive forces which the robots exert on one another and the attractive forces of the goals. The repulsive forces are the internal forces of the team. An obvious difference with the many-body problem is that the repulsive forces in a robot team work in a direction opposite to the attractive forces in mechanics. The attractive forces of the goals on the respective robots are external forces. The difference between the two force systems is clear from the fact that the repulsive forces do not affect the Centre of Gravity, whereas the attractive forces of the goals do cause the Centre of Gravity to move.

A Robot Team: Internal Forces

We briefly consider a robot team in which the robots perform only avoidance and no goal finding. The system of repulsive forces generates particular behaviours in the robot-team, typically the *CG* is an invariant point.

Observation 5.18. The repulsive forces do not affect the *Centre of Gravity*.

Proof. (analogous to Newton’s third law). $\square$

Observation 5.19. If three robots are not on the same line, then the repulsive forces make three robots move towards a position in an equilateral triangle, while *CG* does not move.

Proof. Assume that the robots form an equilateral triangle. Focus on one agent, say robot *A*. The sum vector of the repulsive forces on robot *A* is on the baseline of the bisector of $\angle BAC$, which coincides with the median through *A* and *CG*. Moreover, for every robot the magnitudes of the forces and the sum of the forces are equal. Thus the robots remain in an equilateral triangle.

Now, assume that the robots form a general triangle. The repulsive forces decay with the distances, so the robots that are closest to each other receive the larger

forces. The equilibrium is achieved when all distances amongst the robots are the same. $\square$

The medians divide the opposite side of the triangle into equal parts and go through CG . Observation 5.18 says that CG will not move. So, while the robots are moving away from each other, their positioning preserves certain relationships and dependencies, otherwise CG would move.

A Robot Team: External forces

To explain the system of external forces, we first consider the team of three robots as a point and look at the force field generated by the attractive forces. The attractive forces of the goals have constant magnitude, definition 5.15. An important consequence is:

Observation 5.20. At the *Steiner Point* (ESP) inside the triangle G_A , G_B and G_C , the sum of the attractive forces equals the zero vector.

Proof. At the Steiner Point the attractive forces are under angles of 120° , summing any pair of them (say $\vec{F}_A$ and $\vec{F}_B$) delivers a sumvector of length $2(\cos 60^\circ * (\vec{F}_A + \vec{F}_B)) = 2 \times \frac{1}{2} \vec{F}_C$, directed straight opposite of $\vec{F}_C$. $\square$

Observation 5.21. The Steiner Point is the *unique invariant* point of the system of attractive forces; that is, whenever none of the goals G_A , G_B and G_C coincide or are on the same line, the goals are in a general position and form a triangle.

Proof. At any point each goal exerts on its robot a force of constant magnitude $\|F\|$, directed towards itself. Assume, at a certain point, the sum of the forces equals the zero vector: $\vec{F}_A + \vec{F}_B + \vec{F}_C = \vec{0}$. It is obvious that this point must lie inside the triangle, thus at this point: $\vec{F}_A + \vec{F}_B = -\vec{F}_C$. We resolve the forces on the base line carrying $\vec{F}_C$. The angle ($< \pi$) which $\vec{F}_A$ makes with this line is called α ; and the angle which $\vec{F}_B$ makes with it β ; $\vec{F}_A + \vec{F}_B$ is directed opposite to $\vec{F}_C$. Thus we have $\cos \alpha \|F\| + \cos \beta \|F\| = \|F\|$, that is $\cos \alpha + \cos \beta = 1$. Moreover, $\sin \beta - \sin \alpha = 0$ thus and $\beta = \alpha = 1/3 \pi = 60^\circ$.

The attractive forces go through the three angle points of triangle $G_A G_B G_C$ and their base lines make angles of 60° . The Steiner Point (as defined by Steiner) is the unique point where three lines through the angle points of a triangle meet in such a way. Thus, our point must be the Steiner Point. $\square$

ESP is unique for the force system $\vec{F}_A, \vec{F}_B, \vec{F}_C$. Within the robot team, the attractive forces make the Centre of Gravity CG move. After the robots have reached their respective goals, CG will be at the Centre of Gravity of the triangle G_A, G_B and G_C (we denoted this point as ECG: Endpoint for the Centre of Gravity). We note that

this point is not always the same as the Steiner Point ESP , as observation 5.22 below states.

Observation 5.22. If SP coincides with CG in a triangle, the triangle must be equilateral.

Proof. The centre of gravity is the point of a triangle where the medians intersect. A median divides the opposite side into equal parts. If CG coincides with SP , the medians make angles of 60° , thus the distances from CG to the angle points are all equal, and the triangle must be equilateral. $\square$

Consequently we need to discriminate between the SP and ESP , and the CG and ECG of the team of robots. The CG and the ECG are important for the internal forces, the SP and ESP for the external forces. We have explained the roles of the Steiner Point SP and ESP , and the centre of gravity CG and ECG . In appendix A we give an analysis of a team of two congenial robots. In this case there is no Steiner Point and we show that balanced conflicts occur only if the positioning of the robots and the goals in a real-collision is symmetrical. The proof is elaborated on the fact that point CG moves towards ECG . The analysis results in a complete description of the behaviour of a team. For a three-robot team the case is rather complex, as the different roles of CG and SP indicate.

5.6.3 Balanced Conflicts

We now combine the internal and external force systems. This implies that we cannot any longer treat the robot team as a point system. We start with considering particular cases where the points CG , SP and ESP coincide. These cases are easy to treat as the robots form an equilateral triangle (observation 5.22).

Lemma 5.3. If in a real or semi-collision, CG coincides with SP and ESP , and for all robots r , ESP is on the base line through $[G_r, r]$ then the robots are in a balance conflict.

A typical example of a balanced conflict and Lemma 5.3 is depicted below in Figure 5.14. The robots A, B and C are all three in a real-collision: each needs to pass the same point at the same point in time, moreover this point coincides with SP and ESP .

Figure 5.14: A balanced conflict: the robots form an equilateral triangle, the intersection point of the robots' courses coincides with SP and ESP .

Observation 5.23. Each robot is situated on the base line in the order r , ESP , G_r and the robots form an equilateral triangle.

Proof. The robots are in a real-collision or a semi-collision, which means that they have to cross each others' course to the goal (definitions 5.4 and 5.5). The common intersection point is ESP . The courses are under angles of 60° , since $ESP = SP$. $\square$

Proof of Lemma 5.3. The robots form an equilateral triangle and thus the base lines of the repulsive forces make angles of 60° . Moreover, the baselines of the attractive forces make angles of 60° , and coincide with the diagonals in the triangle of robots. Thus, the attractive force is just opposite to the sum of the repulsive forces. At a certain point on each robot's course, the sum of the repulsive forces and the attractive force equal 0. The robots get locked at these points, each robot at the same distance from CG . They remain in an equilateral triangle. We recall that a balanced conflict is a series of situations in which none of the robots reaches its goal (definition 5.12), thus the considered case is a balanced conflict, which proves Lemma 5.3. $\square$

zones I, III, V (or in II, IV, VI). Let us consider robot C in zone I (for the other robots the same considerations apply).

The repulsive force $\vec{R}_{CA} + \vec{R}_{CB}$ will be on the base line through CG and the current position of robot C . The force points opposite of CG . $\vec{F}_C$ points towards GC but is not on a line through CG . Thus, $\vec{F}_C + \vec{R}_{CA} + \vec{R}_{CB} \neq \vec{0}$, and points away from the base line through G_c and CG , towards zone II. The same is true for the other robots, so the robot team will turn anti-clockwise and, as soon as C has passed the line $[A, G_A]$ it will approach its goal. From this point of time onwards there is clearly a targeting series $TS(C, G_C)$.

If the robots are in zones II, IV, VI, the team will turn clockwise. This proves Lemma 5.4. $\square$

5.6.5 Far-Away Goals

In the lemmas 5.3 and 5.4 the robots are inside the triangle formed by the goals. In Lemma 5.5 we look at cases where the robots are relatively close but where the goals are far away on one side of the robot team.

Lemma 5.5. Assume that the goals of the robots are infinitely far away, though all situated in the same direction. Moreover, assume that two robots (A and C) are on a line orthogonal to the direction of the goals, and the third robot (B) is between the course of the robots (A and C) to their goals, (see Figure 5.16 and Figure 5.17) then three cases are possible:

Figure 5.16: Three robots on one line, goals infinitely far away.

Case a: if all three robots are on the same line l , they proceed towards the goals while remaining on a line parallel to the initial one. In Figure 5.16 the forces are sketched and line l is parallel to l' .

Proof. Easily understandable from Figure 5.16. All forces work in directions parallel to the lines l and l' . The distance between l and l' equals the magnitudes of the attractive forces, all of which are equal.

Remark, robot B will tend to a position just in the middle between A and C , so that $\vec{R}_{BA}$ equals $\vec{R}_{BC}$.

Case b: assume that robot B is a little behind the line l which connects A and C , then B will always remain behind the line connecting A and C .

Proof. The repulsive forces $\vec{R}_{BA}$ and $\vec{R}_{BC}$ have a non-zero component, pointing into the direction opposite of $\vec{F}_B$, whereas $\vec{R}_{AB}$ and $\vec{R}_{CB}$ have a component summing positively to $\vec{F}_A, \vec{F}_C$.

Figure 5.17: One robot just behind, goals infinitely far away.

Case c: assume that robot B is ahead of the line l which connects A and C , then B will always remain ahead of the line connecting A and C .

Proof. The repulsive forces $\vec{R}_{BA}$ and $\vec{R}_{BC}$ have a non-zero component summing positively to $\vec{F}_B$, whereas $\vec{R}_{AB}$ and $\vec{R}_{CB}$ have a component pointing into the direction opposite to $\vec{F}_A, \vec{F}_C$.

This ends the proof of Lemma 5.5. $\square$

Remark, assume a situation as in Lemma 5.5, but all robots have a common goal which is finitely close. The robots will show behaviour similar to either case b or case c, in the sense that robot B in the middle overtakes A and C , or is overtaken. The decisive point from where B overtakes or is overtaken is not on the line connecting A and C but behind it. Moreover, if B is exactly in the middle we are eventually dealing with a balanced conflict between two robots A and C ; such cases are dealt with in appendix A1.

5.6.6 Reviewing the Potential-Field Procedure

The main characteristics of the behaviour of a team of three robots are captured in lemmas 5.3, 5.4 and 5.5. A further and thorough mathematical analysis should prove that these cases indeed cover the main characteristics. So far we have the following. Observations 5.19 through 5.22 indicate the difference between the internal system of repulsive forces and the external system of attractive forces. Lemma 5.3 shows how the combination of both systems leads to a balanced conflict. Such conflicts are relatively rare. Lemma 5.4 shows a class of situations, the situations are close to a balanced conflict, but while the conflict evolves the robot team turns around a common point. Lemma 5.5 deals with cases where a conflict of three robots develops such that a conflict between only two robots remains. Also Figure 5.2 shows a case where a three-robot conflict evolves into a two-robot conflict.

A two-robot system applying the artificial-potential field procedure is analysed in appendix A1. It is shown that balanced conflicts coincide with symmetric real-collisions. Moreover, non-symmetric real-collisions and semi-collisions never evolve into symmetric real-collisions. We conclude that balanced conflicts form a closed category. If the team starts in a symmetric real-collision all subsequent situations will be of this kind. However, if the team starts in a non-symmetric situation it will never encounter a balanced conflict. For a three robot team, this last part, to characterise completely when (under which conditions) balance occurs, is still being investigated.

In potential-field-based avoidance goal finding and obstacle avoidance are integrated in one calculation, so that the approach is rather transparent and can be handled analytically. Although more practical constraints might force one to choose other approaches, the present study presents important starting points to analyse the emerging interaction in congenial avoiding multi-agent environments.

In Section 5.2 we also discussed antagonistic teams. Antagonists chase for a moving target while avoiding other agents. We close the discussion with an example in Figure 5.18. The robots are selective non-aggressive antagonists and apply a potential field based procedure. We have seen in Section 5.2 that these robots cluster together. The moving target (indicated by a small square) in this example is a blind agent. At the start, the three antagonists surround the target, they are in a balanced conflict. The team behaviour becomes intriguing when the goal starts to move.

Figure 5.18: Robots moving in formation.

The team follows in formation, that is without colliding with each other or the goal. Thus, by letting the goal move, we obtain a robot team with quite a different character. Such a team is applicable on a cargo terminal, for instance, to move a set of cargo items. In any case, intuitively it seems that a neat description of this *moving in formation* can be achieved along the lines used to define interaction and conflicts.

5.7 Conclusions on Interaction

In Chapter 4 and 5, we have dealt with problem statement 2: how can a set of resources be made to perform as an actor in a chain. We studied robots as resources in the context of a chain of actions. In Chapter 4 we discussed the design and construction of robot-agents, while in this chapter we analysed spatial interaction and reviewed team behaviour.

The intentions of individual actors determine and characterise team behaviour. We defined intentions as regular patterns in an agent's reactions to other agents.

Congenial teams are teams in which the member apply the same intentions. We discussed several teams in Table 5.2 and concluded that avoiding teams scatter while antagonists (who chase others) cluster together.

We further analysed the dynamics of two types congenial avoiding teams: a team which applies flee-opposite avoidance and a team which based avoidance on an artificial potential-field. Central in the analysis is the notion of conflict, some conflicts do not resolve, we called them balanced conflicts. In Table 5.3 we predicted the behaviour of a newcomer in a congenial team. The newcomer can advantageously influence the team behaviour considerably, as Figure 5.16 showed. The thus obtained team seems well applicable on a cargo terminal as part of a supply chain.

5.7.1 Problem Statement 2

To answer problem statement 2, the approach which we developed to assess interaction is mathematically transparent and provides the basic insights for team development. In Chapter 1 we referred to the metaphor of an ant colony and wondered what binds these agents to perform team behaviour. The study of teams of robots has revealed that the basics of team behaviour can be traced back to the intentions of the team members. Team behaviour evolves, without there being an overall plan. One might conclude that in supply chains similar mechanisms apply, as the silken thread seems to suggest. However, such a conclusion should be supported by further research.

6 Motor-Schemas and Cognition

Studying the art of moving physical objects, we analysed supply chains in Chapters 2 and 3 as an ordering of actions and activities. Chapter 4 dealt with robots and how a motor-schema collects sensor input and reacts to objects. However, we investigated sensors without a concrete definition of objects; in fact objects were reduced to sets of attributes with values. Teams of robots were the subject of Chapter 5; we studied team behaviour generated by the intentions of the individual robots. We defined intentions as patterns of spatial behaviour, thus also reducing the ends at which they are aimed to spatial positions. In the context of studying moving physical objects, intentions have to do with objects as their ends. In this chapter we analyse the intentions as a cognitive notion. Starting from robots we work towards natural beings, and arrive also at a cognitive notion of object.

Below, we aim to capture the action-driven aspect in the notion of having an intention. We set out by defining when a motor-schema implements an intention. A motor-process is the analogue of a motor-schema in a natural being. A motor-process is associated with a motor-image. The latter denotes a state in which the creature simulates sense inputs and actions without performing. A motor-image can concern an object, and hence it can be what an intention is aimed at. Thus, supplementing the discussion of Chapter 5, we arrive at a qualitative notion of intention and we can make a shift from mechanical and motor behaviour to intentional behaviour. We also reconsider the notion of cognitive basis introduced in Chapter 4, and define a motor-image basis which includes actions as a component of a cognitive notion.

6.1 Intentions and Action

History provides a variety of notions of intention, many of which do not directly relate intentions to actions. However, in our view this relationship is essential, as we argue below. Our view is in accordance William James, and for a first approach we consider some of his examples.

The first example is about iron filings, which fly through the air to stick to the surface of a magnet. James [1890] notes "(...) and if a card covers the pole of the magnet the filings will press forever against its surface." The next example is about Romeo who wants Juliet: "(...) and if no obstacles intervene he moves toward her by as straight a line as they [the filings]. But Romeo and Juliet, if a wall be built between them, do not remain idiotically pressing their faces against its opposite sides, (...) Romeo soon finds a circuitous way of touching Juliet's lips directly." James notes that with the filings the path is fixed, whether it reaches the end or not. With the lover it is the end which is fixed, the path may be modified indefinitely. The third example concerns a living frog at the bottom of a jar of water. The need to breathe will soon make it swim straight upwards, taking the shortest path to this end. "*But if a jar full of water be inverted over him, (...) he will restlessly explore the neighbourhood until by re-descending again he has discovered a path around its brim to the goal of his desires.*"

The actions in the examples are performed to pursue certain ends. The ends are comparable in that they consist in arriving at a certain position. Romeo and the frog distinguish in their behaviour clearly from the filings, since they explore alternatives to achieve their ends. James concludes to the following test: "*The pursuance of future ends and the choice of means for their attainment are thus the mark and criterion of the presence of mentality in a phenomenon. We all use this test to discriminate between an intelligent and a mechanical performance*" [James, 1890]. Referring to our Chapters 4 and 5, robots are interesting adaptively performing artefacts that seem to have intentions: their behaviours show that they pursue certain ends without the path being fixed. James' test inspires our attempt to be more precise. We use the criterion of exploring alternatives to distinguish different levels of intentional actions, ranging from the filings through to human beings; robots find a place between these extremes.

In this chapter we advocate a view of perception and the mental activities which emphasises the action-driven aspect, which we also found in the examples of William James. Eventually, this view leads to the conclusion that cognition is founded on autonomous perception-action mechanisms: there is not a basis for cognition without action and in return action gives rise to cognition. The field of robotics is sometime called artificial life research, in particular where the aims are to extend the results obtained to build life like creatures as animats [Trullier and Meyer, 1997]. Artificial life research tends to adopt a similar point of view. A robot is designed to perform actions, and all its constructions are directed to that. In robotics it is rather obvious to question: What is the sense of collecting information, if it is not used in action? In natural creatures, actions follow as a result of perception. Moreover, sometimes action and perception occur in parallel, as for instance in eye-movements [Berthoz, 1996]. Thus cognition in natural creatures cannot be studied without action, and requires an action-driven approach.

In the action-driven view the concepts of cognition and mind cannot be isolated from action and acting, and thus it presupposes that mind and body are intimately related. This presupposition challenges many theories of cognition and mind, in which the mind is opposed to the body. Advocates of such theories have long puzzled about the relationship between mental categories and the outside world. This puzzle also pervades the field of artificial-intelligence research. The action-driven point of view of sensory-robotics conflicts with *symbolic AI*. Symbolic AI is concerned with high-level cognition which is identified as (formal) symbol-structure manipulation [Pfeifer, 1995]. The implicit view of the concept of mind is that of an abstract reasoning centre. Though abstract reasoning is generally counted among the higher (mental) levels, this view ignores any connection of the mind with the body. Thus, the abstract mental concepts become opposed to the external objects which they are supposed to denote.

Similar problems arise when cognition is conceived as a strictly mental process, and also in the so-called Belief-Desire-Intention approach in AI [Roa and Georgff, 1996]. In the latter approach actions result from cognitive activities, but the interdependence between actions and cognition is not considered. As above, it has to be answered how cognition, conceived as a strictly mental process, can cause actions to be performed.

6.1.1 Intentions of Actions

Actions are performed to bring something about: the performing creature has certain aims. In robotics the designer determines the aims and implements them in his creature. Thus, the appropriateness of an action is judged according to how well the robot's action brings about what the designer had planned. However, living creatures cannot be judged in such a way, they generate their own aims. Natural organisms carry the knowledge of the evolution of their species in them. Successful intentions are part of their being. But, what aims can be ascribed to a creature? The usual goals, such as preservation of life, nutrition, and reproduction of the species, are achieved only via a rather long chain of causal relations.

For instance, nutrition is supposed to preserve life. But do we human beings eat because we are hungry or because we want to preserve our life? The absence of food makes us eager to eat, the result of eating is that we feel satisfied, but what about preserving our life?

In Chapter 5, we referred to the silk moth. The silk moth finds a potential partner by zigzagging into an odour plume. Also here we question what exactly the silk moth's intentions are, is it just working on the odour plume, or does it have something in mind concerning mating partners, not to mention reproduction? The silk moth, walking into the odour plume resembles to a large extent the robot with the potential field procedure discussed in Chapter 4. An intention defined as finding the source of an odour (the place where the odour is most intense) has some

final state of saturation and it fits into our study of the art of moving physical objects. The same is true for the frog example of James. We restrict the discussion to intentions that can be traced back in the cognitive basis of a creature. Thus the arousal of the intention and the point of its realisation are kept close together. For an intention like reproduction or preserving life, it is difficult to delineate a final state and we leave such intentions outside the discussion.

And thus, returning to the example of Romeo and Juliet, we consider only the aspect that they want to be close to each other. We observe the agent and ascribe intentions to it, but the aims might have a rather wide scope. We restrict the discussion to intentions that have restricted aims and ends. The aims as to why an agent acts as it does will not be our concern. Agents just act and react. We will not bother about the reasons an agent might have for acting; in a supply chain, an agent acts because it is prescribed to act. Also the frog, it seems, cannot but act. Nevertheless the actions result into something.

6.1.2 Preview of Sections

In Section 6.2, we discuss the epistemological consequences of distinguishing frames-of-reference. Moreover, we review distinct notions concerning the mind: abstract reasoning and consciousness. We show with examples that consciousness and abstract reasoning do not contribute to understanding the action-perception cycle.

In a motor-schema of sensory-based robotics, sensors directly cause actions. In Section 6.3 we stipulate when a motor-schema is an implementation of an intention. In Section 6.4 we elaborate the action-driven view of cognition of natural creatures and we refer to Aristotle. He states that perception produces imagination. In psychology, the notion of motor-process refers to a pattern of knowing-how (to execute an action), and includes a (mental) *motor-image* as well. We explore the role of imagery in thought and thinking.

In Section 6.5 we return to our notion of cognitive basis and define what we call a *motor-image basis*. The novelty is that motor-images, which are based on (potential) actions and perception, are at the basis of cognition. We also reconsider the role of consciousness. Consciousness enables humans to transform abstract notions into actions. We look briefly at the training methods in Zen Buddhism and the eastern disciplines of martial arts where manipulation of the body is a central item.

To conclude, we show in Section 6.6 that the idea of a motor-image basis provides a basis for mutual understanding between creatures and thus is useful to found interaction and communication on.

6.2 Philosophical Notions

In the sections below, we discuss intentions of robots and natural creatures, they all collect signals from the world and act in it. Though the world is the same, we stress in Subsection 6.2.1 that not all creatures collect the same information about it. Moreover, whereas a robot is not but an acting device, for living species there are more nuances at place. Human beings do have a mind, and when discussing intentions we touch on the mind-body relationship. In Subsection 6.2.2 we briefly introduce the notions of rational self and consciousness. We argue in 6.2.3 that these notions do not contribute to the understanding of the perception-action cycle but on the contrary pose problems.

6.2.1 Objects

Perception is directed at the world around a creature; we make a slight excursion into metaphysics to state explicitly some philosophical points concerning this world. The first point is the ontological presupposition that there is a real world of objects. The second point is the epistemological presupposition that this world of objects is (in principle) *accessible* in the sense that there exist means to collect information related to the objects.

Though these two presuppositions state that the world is transparent, some problems still remain. One problem is theoretical, and can be traced back to Kant [1787]. He questions whether we really have access to the objects (*an sich*). We need not be concerned with this problem here. The other problem is practical. We discussed in Chapter 4 the difficulties in comparing different types of sensors. The same problem arises when comparing creatures with different senses. As a practical fact this has the epistemological consequence that, though the world is in principle transparent and cognitively accessible, it does not mean that all creatures share the same inputs and information about it. Stewart [1995] emphasises a similar point, when arguing that all objects are equally accessible to cognitive creatures. Unfortunately, he calls this presupposition *objectivism*. A term which to our opinion might raise confusion, in particular because it is an epistemological point and not as Stewart assumes an ontological (or metaphysical) one.

Humans have a strong inclination to take their world as being the real one (and why not, it is convenient; except that, and this is the point here, it is not the complete picture). We refer to an example already mentioned by Aristotle in his book *De Anima* at (421a7)²⁴. Humans smell poorly compared to many animals. The point is that the world is rather more differentiated than we humans smell it.

²⁴ References to Aristotle are given in the standard way, the page and line numbers trace back to the complete works published in 1848-'74; we use the translation of Hamlyn [1968].

The world is (in principle) accessible to all cognitive creatures. The accessibility is not the issue; the point is that creatures have different senses and thus the inputs which they receive from the world differ. As a consequence, and this is important for the discussion below, the human view of the world is not in all respects the ultimate one.

6.2.2 Consciousness

A discussion of cognitive faculties and actions, inevitably touches on the duality of mind and body. The concept of mind includes several distinct notions. In this section we briefly discuss consciousness.

In the history of philosophy the notion of mind evolved from certain notions of soul. The concept of soul has a generic connotation. In Aristotle the soul is the *principle of life*, and in this sense even plants have soul (De Anima, 413b1). Soul is the cause and first principle (415b8) and the essence of a living body²⁵. We do not involve this notion in our discussion. However, its definition is of interest, because we can define mind analogously, and say that having mental faculties (such as mind and also perception) means that some processes take place which are mental. In this view, the mental process of perception is nothing but the fact that one perceives. In fact, we endorse this view later in our definition of having an intention (Definition 6.2) and in our definition of a motor-image basis (Definition 6.5).

Another notion of soul is encountered in statements such as: human beings have a soul. Here, the soul refers to *the individual*, or *the self*²⁶. We do not discuss this notion further, although it is related to a connotation in which the soul is the *rational self*. For instance, Plato (in his middle dialogues) uses this connotation. The soul is that element whose good condition constitutes our true well-being, and good condition is achieved by intellectual training. The notion of rational self is further developed in Western thinking and one usually also distinguishes (self)

²⁵ For the interested reader: a related generic connotation is that of *soul-stuff* [McDowell, 1973], each individual has/shares a quantity of "soul-stuff". This connotation appears often in arguments on the immortality of the soul. As for our discussions here, it is sufficient to say that these connotations do not affect the notion of mind (in perception) which we are exploring.

²⁶ Soul is sometimes considered to be that which identifies a person, *the true self*, as opposed to the body. In such a view, at the event of death, body and soul are separated. A problem here is who is said to die, the body or the soul. If, as is generally supposed, it is not the soul that dies, who then is buried, the person or his body?

consciousness and *mind* [Flanagan and Güzelde, 1997]. The rational self and consciousness are what in fact is opposed to the body in the mind-body problem.

A prominent view of (self) consciousness is elaborated by Descartes in his meditations [1642]. Descartes' argument starts by supposing that: "(...) *some evil genius not less powerful than deceitful, has employed his whole energies in deceiving me; (...)*" (meditation I). This evil genius has even deceived him in that he has no body and no senses. What remains for Descartes is: "*Then without doubt I exist also if he deceives me, and let him deceive me as much as he will, he can never cause me to be nothing so long as I think that I am something.*" (meditation II). This is often summarised as: *I think, therefore I am.*

This rational self or consciousness remains after sense perception has been dropped. It is clear that it is hard to incorporate such a notion of (self) consciousness with the perception based notion of mind which we will be aiming at below. But this view of consciousness may also lead to an inherent problem.

The rational self and consciousness are in the realm of the mental and concern cognition. Perception starts as a bodily process, but perception also has a cognitive content which represents the outside world. To access this content, one needs some entity that is "reading" these representations. Otherwise, without an inner recipient, perception exists as a bodily process and the cognitive content cannot be used. As a consequence, many theories of cognition have consciousness play the role of the inner recipient of sense perception. This is basically what Flanagan and Güzelde [1997] call *the doctrine of the ghost in the machine*. (The metaphor of the ghost in the machine is from [Ryle, 1949].)

6.2.3 Abstract Reasoning

The Cartesians further developed the view of Descartes, by creating a nearly absolute distinction between mind and body. Present day's symbolic AI holds a similar view, which permeates discussions on artificial life and consciousness. In this subsection we criticise this view, and argue that the consequence of this view is that the mind might create a world of its own.

We note that we are not criticising the fact that we reason with abstractions, because as a matter of fact we do. Humans have impressive verbal skills. To the core of them belong symbolic and abstract reasoning from which the science of *logic* developed. From computer science we know that it is possible to use physical representations for abstract symbols and to reason with them. In fact, we have to assume that humans do perform symbolic and abstract reasoning and operate successfully with it. Under criticism, however, is the (often implicit) point of view that cognition is exclusively located within the realm of abstract concepts.

Self Contained

Typical for abstract and symbolic reasoning is that the system is (supposed to be) self contained or (in technical terms) complete. Everything there is in the world can be described within the symbolic system. The ultimate example of it is a logic (system), which can be proven to be complete with respect to certain models; this means that within the logic all propositions that are valid in the model can be deduced from the axioms. The statement: I think, therefore I am, has a flavour of this, below we explore what this means for a living being.

Illustrative for how a symbol system might operate is the *Chinese room example* [Searle, 1980]. A person (unfamiliar with Chinese writing) is locked in a room with a batch of Chinese characters. Moreover, he has been given instruction so as how to combine the characters. In the process he is given more Chinese characters with instructions how to combine them with the batch he already has. According to the instructions he returns other characters. The people who provide these characters call the first batch a “story” and the last one “questions”, moreover they call the characters returned “answers to the questions”. The question of course is: does the locked person understand what he is answering? We are inclined to give a negative answer, something is wrong. For the person in the room there is no relationship between the characters and the world which they denote.

Since symbol systems are self contained, they are not impeded by difficulties like what information relates to which object. The objects are the basic entities – referred to by variables – of the symbol system. A symbol system does not encounter the problem which we discussed in the context of artificial perception, of how to extract which sensor data concern the same object. In a manner of speaking, the symbol system with its axioms defines its own world; that is to say: the world is as the system supposes it to be. Symbol systems neither contain nor provide cross-references between the symbols and the external world (to which the symbols refer) and thus do not attribute to our discussion of how intentions as a mental phenomenon relate to the external world.

Physical Phenomena

When symbolic reasoning tries to deal with external physical phenomena problems arise. Very typical ones are the so-called Gettier examples. One of them goes as follows [Harman, 1973]: “*A person may see that there is a candle in front of him. Another person may be equally justified in believing that there is a candle in front of him and he may be right too, although a mirror intervenes between the second person and his candle so that what he sees is the reflection of a different candle off to one side.*” A mind as a spectator of the senses cannot solve these problems.

The notion of consciousness of the Cartesians, paraphrased by: I think, therefore I am, gives rise to a situation similar to the Chinese room of Searle. Consciousness

can invent its own world. Descartes investigated consciousness in his *meditations*, referring to an evil genius that is deceiving him. In our opinion, the evil genius is not needed: consciousness itself can fulfil this role. Consciousness might live a self-defined life, when separated from any bodily experience. This leads to extreme solipsism, a consciousness living in its own self-constructed world, a world not accessible to any other cognitive creature.

We consider a last example which opposes (the perception of) physical appearance to verbal communication. Bateson [1973] argues that in human communication the use of (symbolic) language cannot replace bodily communication or *kinesics* as he calls it. He argues that one disbelieves a person because his kinesics do not conform with his language. *"When boy says to girl, 'I love you', he is using words to convey that which is more convincingly conveyed by his tone of voice and his movements; and the girl, if she has any sense, will pay more attention to those accompanying signs than to the words."* [Bateson, 1973, p. 388].

We are not denying the role of consciousness, nor of verbal and abstract reasoning. The point is that the role of mind and consciousness cannot be fully captured within systems of abstract (verbal) and symbolic reasoning. The last example shows that the mind and consciousness also respond to bodily conveyed signs. We have criticised the Cartesian notion of consciousness in that it must have to do something with the physical world and here perception and action come in. Of course, this is neither an explanation nor a definition of consciousness, but in this view, perception and action fall naturally into place. Besides, when we consider animals, any evidence that is available only can consist of observed actions. Consequently it is for instance problematic (if not impossible) to design tests to determine whether animals (apes) have a faculty of abstract reasoning (for an attempt see [Heyes, 1998]).

In this section we elaborated on the rational self and consciousness. Both play a role in perception, but below in Section 6.4 we show that this does not force us to say that cognitive content (that is the mental) is something separate from physical (bodily) experience. However, before going into that, we first resume our discussion on robotics.

6.3 Robots and Cognition

In Chapter 4 we introduced the notion of cognitive basis, it determines how a robot can observe its environment. In this section we analyse what the (re-)action pattern should consist of in order to say that the pattern is the implementation of an intention. In Chapter 5, spatial characteristics are used to circumscribe when an agent has an intention for a target. In this section we fill in some qualitative con-

straints to delineate when a pattern of actions is qualified as implementing an intention.

6.3.1 Sensor Views and Motor-Schemas

The central point in the discussion of Chapter 4 is that sensors deliver attributes with a certain value. Thus, a sensor filters out certain aspects or characteristics of the world. It delivers a particular view of the world, which we associate with a possible world. We borrow the latter notion from logics: "*Possible worlds are nothing but different ways things might have been*" [Stalnaker, 1984]. Different sensor readings associate with different possible worlds. The cognitive basis shows for which input a creature is sensitive, and provides the elements with which the possible world and a framework of possible worlds can be built.

The cognitive basis describes how a cognitive creature classifies its environment. The framework of possible world-views denotes what states the creature might be in. In a logical framework the point is that amongst the possible worlds in the framework, the real world is to be found. For a logician *propositions* are the tools to designate the real world, "*Propositions map possible worlds into truth values*" [Stalnaker, 1984]. Sensors change the internal state of a robot and perform the function of picking out one of the possible worlds, and thus determine the current cognitive state of the robot.

However, we must emphasise that the above does not mean that robots or other creatures use a possible world framework. In a motor-schema, for instance, no explicit image or representation of the external world is aimed at. The possible worlds are only our way of describing the agents' operations in our anthropomorphic frame-of-reference. It is what Dennett [1987] calls "*the intentional stance*."

Creatures with different sensors (or senses) divide the (same real) world differently, consequently their cognitive bases are different. In Chapter 4 we discussed how a cognitive basis might change if data from different sensors are combined. This is also essential when comparing different creatures. As we have seen, the easy case is that the one cognitive basis maps into the other, that is, the bases differ only in that they are not equally fine-grained. This happens when the sensors differ in resolution. However, when different sensors are applied it is very likely that the cognitive bases cannot be mapped into each other. In such cases, different creatures have quite different views on how the same real world looks. The challenge in studying (artificial) perception is to find out how to compare these views and to determine what different creatures do have in common when perceiving the same world.

6.3.2 Cognition and Internal States

As we have seen in Chapter 4, the robot designer bypasses the frame-of-reference problem by selecting himself which attributes the robot has to react to. He makes the cognitive basis sufficient for the robot to act upon. And, the basis is sufficient if the robot can distinguish all possible states of the environment in which a some motor-schema can be executed from those in which it should not apply any action. This is the point of view of the robot designer. Our discussion here is on what can be said of the robot itself.

In developing his theory of *situated automata*, Rosenschein [1985] gives the following definition:

- (i) agent x is said to know a proposition (ϕ) in a situation where its internal state is v if (ϕ) holds in all possible situations in which x is in state v .

We will not discuss whether in (i) the term 'know' is the most ideal one. Our interest is in whether (i) captures a useful cognitive notion. We first note that the definition covers unintended cases. Consider for instance a tree trunk: "*if a tree trunk has 78 rings then it indicates that the tree is 78 years old*" [Stalnaker, 1984]. The tree trunk satisfies (i), but we are reluctant to ascribe a cognitive state to it!

Definition (i) has a more serious defect: it is not clear to which frame-of-reference ϕ belongs. Rosenschein [1985] intended (i) to define when one (as an observer) legitimately can ascribe a cognitive state to a creature. In this case ϕ is part of the observer's cognitive basis. The consequence is that, if agent x is in state v but ϕ does not hold, the conclusion can only be that the observer made a mistake, he interprets v wrongly. In fact, the relationship between agent x 's state v and proposition ϕ thus is not (much) more than that between the tree trunk and the proposition that "*the tree is 78 years old*". Taking this route, we can ascribe as much cognition to a creature as we like, but it does not clarify the cognitive state of the creature. (We are in a certain sense illustrating our own state of mind). When ascribing a cognitive state, somehow, a relationship (or constraint for a relationship) between v and ϕ needs to be demonstrated.

6.3.3 Actions and Intentions

We focus on perception and action, and our aim is to determine when a pattern of actions can be said to implement an intention. Generally, actions are intentional when they are aimed at bringing about certain results. But when does a robot aim at something, and how can we assure that the aim is an aim of the robot? It is clear that any definition has to go beyond definition (i) above.

To illustrate robot action we return to the robot described in Chapter 4, which is provided with an obstacle-avoidance procedure and a goal-finding procedure in the module AVOID. Sensor data are collected by the SURVEYOR module and for-

warded to the AVOID module. The AVOID module interprets the data received from SURVEYOR as an artificial potential-field. Obstacles are considered as a source of resistance, while the goal is taken as a source of attraction. The potential field is visualised as a landscape in which the repulsive sources form mountains and the sources of attraction valleys. Thus the AVOID module determines a route of low resistance, (it steers the robot through valleys in the landscape). This and similar procedures work rather well in implementing obstacle avoidance, as we have discussed in Chapter 5. The resulting behaviour is shown in Figure 5-9. (The observation of the environment is discussed in Chapter 4, see Figure 4-6).

In the example, the sensor data are taken as establishing a potential field in which the robot is forced toward a region of low resistance. Obstacle avoidance is no more than the consequence of a mechanism assembled such that whenever a certain state occurs in the environment, a corresponding internal state activates the actuators. Compare a thermostat put on 21° . Its construction (with bi-metals etc.) is such that when the temperature drops, the heating is connected to the power supply. Later on the room gets warmer. Stalnaker [1984] calls this a *tendency-to-bring-about*. In the robot the signals picked up by the sensors are transduced by the wiring into activations of the motors, as in a control loop. It is clear that this goes beyond the tree trunk indicating the age of the tree.

From the example of the obstacle avoidance procedure, we might infer that the robot has the intention *to avoid obstacles*. But the avoidance of obstacles is not directly implemented. It is not the case that at a certain point of time, the internal processing concludes *this is an obstacle* and subsequently starts avoiding it. Within the cognitive basis we cannot draw this conclusion; as we have seen in Chapter 4, the sensor readings do not allow to conclude to the existence of an object. The potential-field procedure generates actions which fall in with obstacle avoidance, but strictly speaking it does not avoid any obstacle.²⁷

When a motor-schema is executed, the robot goes through a causal chain of (partial) internal states: states of the sensors, the wiring, and the actuators. This might be a long chain of minor changes. The cognitive basis shows to which inputs from the environment the robot is sensitive, thus we can distinguish (series of) internal states which are relevant in relation to the sensor inputs. For convenience of our discussion, we identify sensor inputs with the internal states $v \in N$ which they ultimately bring about. We will call N the set of these internal states. Motor-schemas can be treated similarly. A motor-schema implements a tendency-to-bring-about a certain internal state. We identify the motor-schema with this internal state.

²⁷ We note moreover, that the idea of an artificial-potential field procedure is derived from physics, where its implications are rather different.

The robot using the obstacle avoidance motor-schema with the potential-field procedure, aims at an internal state from which no (new) impulses go out to the actuators, a point in the landscape with an equilibrium amongst the repulsive and attraction forces. We call this state v . The performance of the robot is captured in the following preliminary definition:

- (ii) *performing motor-schema v* : in a situation where the sensor inputs cause the mechanism to be in an internal state different from v , the robot performs until its inputs cause v to occur.

Definition (ii) describes a causal and mechanistic process. Moreover, it refers only to internal features of the creature. The state v is internal, the potential field (landscape), if anywhere, exists only internally in the robot. Thus, as in the example of the robot, the main driving forces, the finding the goal and avoiding obstacles, can be pre-programmed. Performing a motor-schema however does change the environment, something that is not explicitly mentioned in (ii).

We want to lay down what the robot to bring about in the environment, so that we can ascribe an intention to the robot. However, as we argued when discussing definition (i), the relationship between internal state v and the intention γ needs not only to be explained but also to be confirmed by some output that results because the robot is in internal state v .

Before going into defining when an action has an intention, we resume the context discussed so far, Figure 6.1 gives an overview. In the middle the real world is represented. The robot creature starts acting when the real world is in a situation s_0 . In this situation its internal state is different from v , we abbreviate this as $\neg v$. The connection between s_0 and $\neg v$ is established by the sensors, which is consistent with the starting clause of (ii). Because of $\neg v$ the motor-schema is triggered and starts performing in order to bring about internal state v . While doing so the actions change the environment and cause s_1 to become the case in the real world, this is shown in the upper part of Figure 6.1.

Figure 6.1: Overview of intentional context.

Having intention γ is aiming for a situation that is not yet the case. At the moment that the intention arises the external world is in situation s_0 . Thus, in s_0 it should be the case that $\neg\gamma$ holds. Moreover, γ must be saturated at the moment that internal state v is reached and thus that s_1 is the case, see the lower part of Figure 6.1. In the following preliminary definition, we make an attempt to ascribe an intention to a motor-schema.

- (iii) A motor-schema v is said to be an implementation of *intention* γ if in a situation where γ is not (yet) the case, motor-schema v is triggered and (under normal conditions) performing motor-schema v brings about a situation in which γ holds.

Definition (iii) does not directly relate the agent's actions to γ . The drive behind the actions is to reach the internal state v . γ concerns the environment, the fact that the environment changes to γ might be a rather indirect result of the actions performed.

The robot with the potential-field procedure has a control loop based on a single attribute. So it avoids obstacles rather successfully, but no obstacle is defined in its cognitive basis Φ . For a robot the frame-of-reference is essential in its design. We (external observers) ascribe the intentions of goal finding and obstacle avoidance to the robot. When we say that it intends γ , but γ belongs to the cognitive basis Γ we have in mind. However, the intention of motor-schema v can coincide only with the environment being in state s_1 .

As yet the relationship between v and s_1 is not tight. Assume, we have a robot similar to ours driving around. However, instead of adjusting its course, the robot just releases a balloon whenever the sensor detects something. This robot has no feedback on its own actions (that is, on releasing the balloon), or more precisely, it has no feedback of the type which originally forced it to perform. Such a robot is reactive, but we do not want to ascribe to it any intentional action, it reminds of the tree trunk above, which indicates the age of the tree.

To do justice to the notion of intention and to avoid the balloon case above, there must exist a meaningful relationship between s_1 and v . In (ii) we required that a motor-schema performs until its inputs cause v to occur. We sharpen this and require that there needs to be the *same type of sensor connection* between s_1 and v as there was at the start between s_0 and $\neg v$. A motor-schema must have an end definable in terms of its sensor inputs. The internal state v , which characterises the motor-schema, must correspond to some feature of the future state of the environment. Thus as a standing definition replacing (ii), we define:

Definition 6.1. *Performing motor-schema.* A robot performs motor-schema v , iff in a situation where the sensor inputs cause the mechanism to be in an internal state different from v , the robot performs until the same sensors cause v to occur.

We can now formulate our definition of intention:

Definition 6.2. A motor-schema v is said to be an *implementation of intention* γ if:

- (a) in a situation s_0 where γ is not (yet) the case, the sensor inputs do trigger the mechanism to be in an internal state different from v , and
- (b) when in this state, the mechanism performs such that under normal conditions it brings about a situation s_1 in which γ holds, and
- (c) in situation s_1 where γ is the case, the same sensors cause v to occur.

With Definition 6.2 we have stipulated that a creature has an intention only if it has some notion of the results of its own action. The requirement that the creature has to have a notion of the results of its action is captured in clause c by requiring that internal state v will be reached as a result of feedback from the environment.

We summarise the requirements for having an intention. Sensor inputs bring the creature to a certain internal state, this state sets into action a mechanism that tends to bring about a certain situation (that is not yet the case). The same sensors that trigger the action also need to interrupt or stop its execution. The intention consists in aiming at this final situation. Note that we have obtained a definition of an intentional action without making any appeal to an internal representation of the situation nor any abstract concept.

The notion of having an intention being defined, we check the examples which we have encountered above. Clearly, the tree trunk does not take any action and has not intention. The robot with the balloons does satisfy clause c of Definition 6.2. The filings and the magnet example of James, discussed in the introduction and the thermostat might be considered degenerated cases. Strictly speaking, they satisfy Definition 6.2, but as James notes, the filings and the thermostat cannot explore any alternatives. Intentions become interesting when the agent can select from several actions. The robot with the potential field has alternatives, it can slowdown or turn, moreover it has two different ends, one is goal finding the other is obstacle avoidance. Because of this the robot is able to explore alternatives: in slightly different (and new) situations it performs different actions, but it retains the same intentions. In this aspect, the robot is close to the frog. The fact that there are alternatives is why we had to elaborate on Definition 6.2. For the filing and the thermostat the 'intentions' can be described univocally, but for the robot it is more complex, this is also why we can say that there is some intelligence present.

6.4 Actions, Intentions, and Imagery

Above we defined the implementation of an intention in a robot, below we extend the discussion to natural creatures. Similar to the robotics discussion, we consider perception as the first part of a process chain that results in actions and explore cognition within this operational chain. Since we consider the chain from perception actions, our approach clearly distinguishes from an abstract reasoning approach, the latter approach was criticised in Section 6.2.

6.4.1 Aristotle: Imagination, Intellect and Desire

As a start we refer to the theory of perception put forward by Aristotle. Though this theory is rather old, its original point of view is illuminating and provides a guideline through perception and action. Moreover, it shows a way to avoid the doctrine of the ghost in the machine (refer to 6.2.3). Aristotle did not presuppose a duality of mind and body. *"It seems that all the affections of the soul involve the body – passion, gentleness, fear, pity, confidence, and, further joy and both loving and hating; for at the same time as these the body is affected in a certain way"* (De Anima, 403a16).

In Aristotle's view, the faculties of perception and having an intention (he calls the latter *desire* as we discuss below) were characteristic of life. The fact that living beings have the capacity to collect information via senses and are directed to the outside world was not in question; it was obviously a (necessary) faculty for living creatures.

In Chapter 4, we signalled the problem of sensor fusion, that is the problem of how to combine signals from different sensors or senses, which originate from the same object. Aristotle notes a similar point. He speaks about the sense and its special objects, but also about the common objects (425a14: e.g. movement, rest, figure, magnitude, number, and unity) for which there is no special sense. The common objects are perceived by the senses *"(...) in so far as they form a unity"* (425a30) and a little further on: *"(...) it is like what some call a point"* (427a9). Interpretations of Aristotle differ on whether this introduces a *common sense* [Modrak 1987] or *perception in common by the individual senses* [Block, 1961; Jäger, 1968]. This point is not decisive for our discussion. In the modern view of the brain and the nervous system it is not a point of discussion that the inputs of the different senses are somehow combined. The interesting point of Aristotle's theory is that the result of combining the inputs is a *phantasma*, which is an appearance, image or phantom [Modrak, 1987]. *Imagination* (phantasia) denotes any process working on phantasmata. Imagination occurs (...) *in things which perceive and with respect to those of which there is perception (...)* (428b12). Cases of imagination are dreaming, remembering, thinking, and desire. Imagination can also invoke actions, and this is a point for our discussion below.

Intellect

Images and imagination are quite central for the relationship between perception and mind. Therefore before turning to actions, we first discuss the impact of Aristotle's theory with respect to the mind-body opposition and how the ghost in the machine is avoided. Aristotle distinguished a particular notion of mind closely related to perception, we will refer to it as *intellect*²⁸. The intellect is explained with a comparison to perception: "(...) as that which is capable of perceiving is to the objects of perception, so must be the intellect similarly to its objects." (429a17). Imagination is the basis on which the intellect works (403a3): "(...) thinking is a kind of imagination which depends on the body (...)" and (433b30): "And all imagination is either concerned with reasoning or perception. In the latter then the other animals share also."

In fact the intellect denotes the process of thinking. The difference between the intellect and the senses is that whereas the senses have their special organs (sense organs) the intellect has no special organ. (De Anima 429a20): "That part of the soul, then called intellect (and I speak of intellect as that by which the soul thinks and supposes) is actually none of existing things before it thinks." Moreover, for this faculty there is no special organ (430a17): "(...) in essence it is activity (...)" The intellect is a function only: the function to combine. Aristotle notes "And that which produces a unity is in each case the intellect." (430b5).

The idea that, in the process of perception, there must be a function that combines the input from the different senses is also found in the *Theaetetus* (a dialogue from the later period of Plato): "(...) all those things (inputs via the senses) converged on some one kind of thing, a mind or whatever one ought to call it (...)" (*Theaetetus* 184d3)²⁹. However, Plato, treats the mind as independent (*Theaetetus* 185e1): "(...) it seems to me that the mind itself, by means of itself, considers the things which apply in common to everything." From this the following is inferred (*Theaetetus* 186d2): "So knowledge is located, not in our experiences, but in our reasoning about those things we mentioned; (...)" In this way of thinking, the mind is the inner recipient of sense perception, and we are back at the ghost in the machine.

Aristotle distinguished his view of the mind from Plato's, for instance at (428a24) refer to [Hamlyn, 1968]. In contrast to Plato's view, the intellect of Aristotle is not

²⁸ We base our discussion of Aristotle mainly on his book *De Anima* [translation: Hamlyn, 1968]. The greek term νοῦς is often translated with "mind". However, to contrast it with the modern use of the term mind, we use the translation *intellect*, as used by Hamlyn [1968].

²⁹ References to Plato are given in the standard way, the page and line numbers trace back to the complete works published in 1578. We use the translation of McDowell [1973].

a separate thing (substance) but a function; “(...) *it is actually nothing before it thinks*” (429ab30). Here we are close to the view in which having a mind consists of the fact that some mental processes take place (refer to 6.2.2 above). Note that the intellect is not opposed to the body. Moreover, there is no ghost – in the form of consciousness or a rational self – which reads perception. By considering the intellect as a particular instance of imagination which works with images (phantasmata), Aristotle bridges the gap – which we note in 6.2.3 above – between reasoning and perception.

Desire

Images (phantasma) and imagination (phantasia) also play an important role in *desire* (orexis). Since Aristotle did not oppose mind to body, the question how the mental can cause actions is not his concern. Nevertheless, he notes that the fact that a creature perceives through its senses does not explain why it is driven to (re)act to perception. Just perceiving something does not trigger an action. Whenever, a creature acts because of perceiving something, there needs to be some drive: this drive is desire.

A desire is directed to an object. An image is what the senses deliver of an object. And since desire is a case of imagination and uses images, there is a link between the internal (desiring) state of a creature and an external object. (De Anima, 431b16): “*And if [a creature has] sense-perception then also [it also has] imagination and desire.*” (432b15): “(...) *for nothing which is not desiring or avoiding moves unless as the result of force (...)*”

Desire is also the basis for any connection of the higher cognitive faculties to the external world. (433a16): “*Every desire too is for the sake of something; for the object of desire is the starting-point for the practical intellect, (...)*” and (433a18): “*For the object of desire produces movement, and because of this, thought produces movement, because the object of desire is its starting-point. And when the imagination produces movement it does not so without desire. Thus there is one thing which produces movement, the faculty of desire.*” And concluding at (433b27): “*In general, therefore, as we have said, in so far as the animal is capable of desire so far it is capable of moving itself; and it is not capable of desire without imagination.*”

6.4.2 Motor-Image

The view of Aristotle provides a guideline to extend our discussion from robotics towards natural creatures. For Aristotle internal states basically have something to do with the external: sense perception generates images. Images are used by imagination; thinking and desire are particular cases of imagination. A desire is the drive behind an action and it can cause an action. Thus, cognitive faculties engage with action.

Below, we further explore cognition by analysing intentions and actions. At the basis of our analysis is the notion of motor-image. Motor-images, we introduce them below, seem close to Aristotle's theory and they will replace of his notion of image (*phantasma*). Desire plays a role comparable to that of having an intention, both generate actions. Moreover, we elaborate the idea that in thinking motor-images are processed, analogously to Aristotle's view that thinking is a case of imagination.

A motor-schema in sensory-based robotics is the analogue of the psychological concept of a motor-process [Berthoz, 1996], the latter applies to natural beings. A motor-process denotes the complex mechanism that links senses to the central nervous system and the effectors. It results in the execution of an action. Motor-processes work like cybernetic loops or servomechanisms [Berthoz, 1996, p. 102]. A motor-process associates with a motor-image. A motor image is defined as a dynamic state during which a subject mentally simulates a given action; a motor-image implies that the subject feels himself performing an action [Decety, 1996]. The approach traces back to William James' concept of *ideo-motor action*. Ideo-motor action is the experience of situations where we find ourselves thinking and acting at the same time. We think to drink our coffee and we find ourselves already holding the cup in our hands.

Annett [1996] argues that motor-imagery has "symbolic" content, we rephrase this: motor-imagery is the cognitive content of performing a motor-process. Note that a motor-image combines inputs from the senses as well as activations of the motor-processes. Thus a motor-image connects physical processes as collecting sense inputs and action with mental phenomena. A motor-image is the internal or mental part of an action producing motor-process and, by activating the motor-process, imagery can cause actions.

Motor-images are defined only for natural beings, thus we have made a step beyond robots and artificial creatures. Since a motor-image associates with a motor-process, an interesting question is whether we can distinguish something in a robot that associates with a motor-schema. We cannot answer this question, and what is more, this may be the clue to point out what is so peculiar about cognition of living beings.

6.4.3 Imagery

We further explore motor-images and motor-imagery. We argued that motor-images are needed to close the perception-action-mind chain. In this subsection, we discuss examples to show the role of motor-images and motor-imagery in the cognitive process; it will become clear that a motor-image implies a certain notion of awareness.

We start by looking at the seemingly simple concept of a flight of stairs. When does anyone experience that he is ascending or descending a flight of stairs? The essential thing is that in ascending it, one foot notably is put higher than the other, and this happens several times in succession. If the difference in height between the steps is very small, we do not refer to it as a flight of stairs. Also, if the steps themselves are very small, (a few centimetres, say) we also do not refer to it as a flight of stairs; in passing along it, we will only notice that the terrain is very uneven.

Suppose that we have a proper idea of what a flight of stairs is. Stairs can be steeply or gently inclined. Whether we conceive of stairs as steep, depends on "bodily" experiences, moreover, the impression changes depending on whether we are descending or climbing.

The neighbourhood of our former house in the south (Genoa) was built on rather steep hills. The shortest route from our house to the shops and bus stop consisted of stairs cut out in the rocks (about 30 meters high). The upper part of the stairs consisted of steps cut out in the front of the rock. On the lower part, the rock was too steep, and the stairs had been cut traversing the rock. On descending the stairs, the upper part, where one was facing the open space, seemed rather steep, much more so than the lower part. However when climbing, the lower part required much more effort than the upper part.

Understandably, the stairs are not used much. On hot afternoons one would encounter several lizards basking in the sun. They fled immediately when approached too close. We can legitimately conclude that the lizards do not consider the place as part of a flight of stairs.

Our aim is not to emphasise the stairs concept, but to show the "impression" or "image" of a flight of stairs. The image is a combination of various sense impressions and the efforts made to climb or descend the flight of stairs; we are dealing with motor-imagery. In the processing, some reductions have taken place. The final image of the flight of stairs is not a straightforward combination of all sense impressions and motor-processes. The lizards for instance, colour our experience, but they do not add to perceiving the staircase as such. Neither do the individual steps count separately. The remaining image, though based on the sense and motor data, is an abstraction from these data.

The next point is that though motor-images are abstractions, they can be recalled, we consider the following example.

At the age of four, my eldest son was speaking to his friend (slightly older) about stairs. We were living at the second floor, our stairs were straight and very steep. The friend said "you have a long stair-case" pointing with his arm sloping upwards. My son agreed. The boy himself lived at the fourth floor, and

their staircase rose in a spiral of eight separated staircases connected by landings. "Our stairs are like this ...", the boy said, spiralling down with his finger and extending the movements with his arm and body. "Yes", my son shouted, "like this ...", starting a great turning movement and while doing so, he whirled downwards, laughing [Penders, 1986].

The interesting point of the latter example is that the presentation which the small boys give of the stairs, is a pattern which is not a repetition of previously performed actions, but highlights the experiences of passing along the stairs. It is clear that what the boys present is an abstraction which cannot be resolved into the original perceptual and motor data.

In the robot discussion we noticed the problem of sensor fusion: how to conclude that the data from different sensors relate to the same object. The staircase examples indicate the ease with which the human (perceptual) mind combines sense and motor data into a (perceptual) object. It seems that we touch here on what Aristotle called the common sense, being the faculty that combines and produces unity. Pushing through this idea, thinking, particularly in the sense of imagery, is therefore the process of examining previously collected (perception and action based) images.

We have discussed motor-processes and motor-images. With a motor-process is associated a motor-image. A motor-process can be recalled to perform routine actions, also a motor-image can be recalled. A flight of stairs is an example of a concept consisting of motor-images. The characteristics which we use to define a flight of stairs, all appeal to motor-processes. Moreover, it is hard to define what a flight of stairs is without referring to bodily movements. Herewith we conclude to the view that perception and action together make up cognitive content.

Other Mental Processes

In line with this view we might speculate further, for instance on emotional experience. Emotions do have an impact on the body and the cognitive processes [Botelho and Coelho, 1997]. A sketch within the view developed here: satisfaction might be associated with a moment of absence of any drive; fear with a shortcoming of a motor process, or lack of any motor process, to deal with the situation at hand. We do not intend this as an explanation of the mentioned emotions. The point is that by proceeding an analysis in this direction, emotions prove to have by nature a bodily part. Thus it becomes understandable why – the other way around – emotions do affect the body.

6.5 Revisiting Cognitive Notions

We have argued that actions via motor-imagery are involved in cognition. In Definition 6.2 we delineated intentions for robots as a process in which the inputs from sensors result in an internal state $\neg v$; and this state triggers a motor-schema which strives for v . For a natural creature we can set up an analogous scheme: sense perception and action bring the creature into an internal state $\neg v$, which triggers a motor-process corresponding to the state v , and this motor process implies actions, but is also associated with motor-imagery.

Taking motor-imagery as an important phenomenon for living species, the connection of the mental with the body is basic. In this view, neither of them makes full sense when disconnected from the other. Adhering to such a view we have to revisit the definition of cognitive basis, introduced in Chapter 4, when considering natural beings. Actually in this definition nothing is said about actions.

In Subsection 6.2.3 we criticised the view that consciousness can be captured in a symbolic reasoning system. Below we also revisit consciousness in order to point out where it relates to action.

6.5.1 Cognitive Basis

In the robotics discussion, we introduced the notions of perception basis, knowledge basis and cognitive basis. For analysing a robot design the cognitive basis as defined in Chapter 4 seems sufficient because the robot designer provides the necessary frame-of-reference. However, when discussing *natural* (human) beings no external frame-of-reference is available, and the ability to perceive distinct objects needs to be found somewhere within the being.

We argued for the importance of motor imagery within the cognitive faculties of natural (human) beings; thus motor imagery needs to be incorporated within the definition of the cognitive basis. Moreover, motor-processes and motor imagery might provide the foothold for distinguishing objects. As a preliminary attempt we propose an extension to the definition of cognitive basis. We change Definition 4.1 in to Definition 6.3 below, by replacing sensors with senses and assuming that the senses and the nervous system map objects to motor-processes.

An agent r , has at its disposal a set A_r of input channels or logical senses and a set of motor-processes N_r . The senses (input channels) $\alpha \in A_r$ assign one or more motor-processes to an object o in the universe of discourse U .

Definition 6.3. *Sense value.* $\alpha: U \rightarrow P(N_r), \alpha(o) = \{v: v \in N_r\} \quad o \in U$.

Thus, as in Chapter 4 we can define an indiscernibility relation for a certain agent. In words the definition says that objects are distinguishable from one another only if they trigger one or more different motor-processes.

Definition 6.4. *Indiscernibility.*

$$\text{ind}(A)_r = \{(x, y) \in U \times U \mid \forall \alpha \in A_r (\alpha(x) = \alpha(y))\}.$$

Thus we can define a *motor-image basis* as the set of equivalence classes of objects which trigger different sets of motor-processes. Whenever two objects fall within the same equivalence class $[o]_{A_r}$, it means that the agent cannot distinguish one from another. The motor-image basis is the set of equivalence classes:

Definition 6.5. *Motor-image basis*³⁰. $E = \{u \subseteq U \mid u = [o]_{A_r}, \text{ for some } o \in U\}.$

In Definition 6.3 the inputs from the logical senses are taken as mapping objects directly into motor-processes or motor-images. Each motor-process incorporates parts of the perception and knowledge basis. Hence, this basis incorporates (parts of) the previously defined perception and knowledge bases. As in Chapter 4 we can define mappings between the motor-image basis and the perception, knowledge and cognitive basis. What might get lost in the mapping are those parts of the cognitive basis that do not trigger a motor-process.

We stress that Definition 6.3 is still rather preliminary. It might have some (not yet foreseen) technical flaws. For practical applications it will be a problem to discriminate between different motor-process, and even more between motor-images. A basic difficulty is that the definition starts from objects, but as we have seen, objects do not have a universal (cognitive) status among all cognitive creatures. But in fact, practitioners, for instance, in biology make reference to actions when they are comparing intelligence of one species with another one. For a biologist there are no alternatives.

What we have achieved with Definition 6.3 is that it provides a way to incorporate actions as basic into the concept of cognition (and intelligence). Moreover, it settles the principle point that objects must have an intentional significance for the creature under consideration. By requiring that a motor-process is triggered, a basic type of intention is implicitly incorporated in cognition. Moreover, the notion of a motor-image basis conforms to the test of William James (refer to Section 6.2): a more extended motor-image basis provides a creature with more choices of means to pursue future ends.

³⁰ Definition 6.4 is the same as Definition 4.2 and Definition 6.5 is the same as Definition 4.3.

According to Kant [1787] humans project their (perceptual) frames onto the world, and our perception is always within these frames. Whereas his problem was to show that his thesis is valid, we made a slight variation of this principle as leading for our discussion: an object is an object – for a creature – only if the senses and the motor-processes of the creature are sensitive to it. Herewith, we also partly answer a question posed by Dretske in his book *“Knowledge and the Flow of Information”* [1981]. Dretske states that signals and data from which to derive information are available everywhere, the puzzling fact is how we select from this abundance of data that which is useful. We cannot but select only what is or in the past has been of use, for other data we are not (yet) sensitive.

6.5.2 Awareness and Actions

We have shown with examples that bodily experiences play a role in cognition, and defined a motor-image basis. In Subsection 6.2.3, we argued that the role of mind and consciousness cannot be fully captured within systems of abstract (verbal) and symbolic reasoning. In this respect we have commented on Descartes' notion of consciousness. However, we cannot omit consciousness and awareness; perceiving and being aware of perceiving are different. We look at an example.

Implementing reasoning about actions into a robot is problematic. The problems start already with sensor fusion, as we saw in Chapter 4. Yet, a present-day robot does not have awareness. It cannot identify objects and as a consequence, it neither can identify itself. The robot simply *is* always in the environment it can “think of”, it cannot distinguish itself from the environment where it collects data. The robot with the IR sensors, described in Chapter 4, can be kept in motion by mounting an arm on it which holds a piece of white paper in the detection range of the IR sensors, as one of my colleagues once did as a joke. The robot starts driving in circles and never discovers what it is doing. The robot chases its own tail like a young cat does. However, a mature cat, but also any human being finding himself in a similar situation would fairly soon stop, becoming aware of the situation.

The examples of the flight of stairs also imply a certain awareness of the subjects. Awareness, which is included in the notion of consciousness, is not a mere epiphenomenon, it can guide and influence our actions: contrary to the robot, we soon stop chasing. Consciousness and in particular awareness arise when acting, and influence the actions. But what are they and what roles do they play when we are acting?

We criticised the Cartesians' view of consciousness. A trace of it is implicit in modern research on consciousness. For instance Chalmers [1995] notes that the “*hard problem*” of consciousness is “*how physical processes in the brain give rise to subjective experience.*” In this view, subjective experiences exist in a realm of their own, the consequence of this view are shown above. The citation states the wrong problem. In our view subjective experiences are based on physical proc-

esses; the problem to explain is how it is possible that we become aware of our experiences and that we can reason about experiences and manipulate with them.

Consciousness and awareness change our point of view and enable us to arrive at verbal and abstract reasoning. They also mediate when we apply the results of the reasoning in our actions. On the one hand the question is: how can the organism change from the experiencing mode to the abstract reasoning mode? And on the other hand, how can humans apply the results obtained with abstract reasoning to their body in order to change their physical environment? Both together, how do humans make the transitions back and forth between bodily and mental?

As a reminder we also note with Aristotle that the mental does not completely dominate the body: (De Anima, 433a1) "(...), even if the intellect enjoins us and thought tells us to avoid or pursue something, we are not moved, but we act in accordance with our wants, as the incontinent man does." The mental only influences the body.

Awareness and Physical Training

We do not have definitive answers for awareness and consciousness. Our point is that an explanation of consciousness also has to include bodily actions and the awareness of acting. Though the answers are lacking, we give below some examples that indicate what an investigation of consciousness and awareness is about. Some examples are taken from the Eastern literature on Zen Buddhism, supplemented with examples from Western lifestyle. The examples deal with the interfacing of the conscious mind with bodily actions. This point is very explicit in Zen Buddhism: practitioners try to re-establish a unity of body and mind through training. The main difference is that where we want to get hold of consciousness, Zen Buddhism tries to get rid of it: the conscious mind must become empty. Nevertheless, a comparison is useful in our context.

The religious context of Buddhism is to acquire the ultimate state of enlightenment. The *Way* of the Zen Buddhists³¹ requires profound mental and physical training, in order to manipulate the mind. History shows a rich variety of Zen masters, each with their particular style of training. What they have in common is that they all seem to have had a very strong will. Which is not really surprising, taken into account the hard training methods (studies). The training methods focus on practical (bodily) experience; study of abstract texts is considered as an addition only, and by itself completely insufficient.

³¹ Zen Buddhism is basically a Chinese variant of Buddhism. Buddhism was imported into China via the Silk Route.

First about the notion of mind in Zen Buddhism, practitioners strive to free themselves from the creations of the mind. We cite Lin Chi, who lived in China about 800 AD. He is considered the founder of what the Japanese call the Rinzai Sect. *"Followers of the Way, this thing called mind has no fixed form; it penetrates all ten directions. In the eye we call it sight, in the ear we call it hearing; in the nose it detects odours, in the mouth it speaks discourse; in the hand it grasps, in the feet it runs along. Basically it is a single bright essence, but divides itself into these six functions. And because this single mind has no fixed form, it is everywhere in a state of emancipation. Why do I tell you this? Because you followers of the Way seem to be incapable of stopping this mind that goes rushing around everywhere looking for something. So you get caught up in those idle devices of the men of old."* [Lin Chi; translation B. Watson, 1993, p. 25-26].

Takuan Soho, a Japanese Zen master who lived at the end of the Japanese period of Civil War (1600 AD) discusses mind and perception in the next citation: *"When you look at a tree, you perceive a red leaf amongst the leaves and your mind 'stops' on this leaf. As soon as this happens you see only one leaf and fail to perceive the rest of the other leaves on the tree. If, however, you simply stand in front of the tree without any thought the innumerable other leaves will come into your eyes. If your mind is identified with one leaf, you will be prevented from seeing the rest of the leaves. You can see hundreds and thousands of leaves if you are not identified with one leaf."* [Takuan; translation Hirose, 1992, p. 24].

The Zen methods are also applied to action and acting. For instance, to the lay contexts of tea ceremonies (the way of tea) and swordmanship, both of which have survived to the present. Both are mastered after long periods of training. Takuan gave the following advice to a contemporary (17th century) swordsman: *"For example, when your opponent advances to strike you, your eyes immediately catch the movements of his sword. If you are 'stopped' by this movement of his sword towards you, your mind loses its freedom and you are sure to be killed."* [Takuan; N. Hirose, 1992, p. 22].

The physical training inevitably touches on the mental. The modern heredity of swordmanship is found in *iaido*, the technique of drawing a samurai sword (katana). Practitioners have to acquire fluidity and speed. The 20th century Zen master Taisen Deshimaru remarks in a discourse for martial art practitioners: *"When the mind is calm, the body can act spontaneously and the action that results is free and easy, whereas the body's functions are impeded when the mind is working all the time"* [Taisen Deshimaru, 1982, p. 61].

Swordmanship and Zen Buddhism might seem far-fetched. However, similar points concerning mind and bodily practice, can be found in for instance the skill of juggling. Throwing and catching a single ball requires motor-processes that most people rather easily learn. When learning to juggle three balls, the novice is fully occupied putting the motor-processes in a proper sequence. Analysis shows that the time schemas are rather complicated [Beek, 1989]. After the juggler has

acquired some basic skills, spectators start to recognise some pattern in the movements of the balls. Surprisingly this pattern also incorporates the movements of the balls in the hands. Sooner or later, the juggler himself starts to recognise the patterns as well. In doing so he turns away from the rhythm of throwing and catching and takes an observer's point of view.³²

Learning to juggle five balls requires quite some practise, as is understandable because the juggler never holds more than one ball in each hand and thus he is forced to keep the other three in the air. Dancey [1994] accompanies his instructions for five-ball juggling with the remark: *"While you are learning you are trying to make yourself do it, when you can do it you watch yourself doing it instead"*. The latter remark conforms with the Zen view of the martial arts, the physical training goes together with a change in the mental stance.

We have argued that consciousness and awareness are related to perception and action. The examples highlight this relationship and show that it is a rather close relationship. Aristotle noted that the mind does not dominate the body, and anyone learning to juggle five balls will readily confirm this. But surely, the conscious mind does influence the body, otherwise no one ever would start juggling!

We started our discussion with motor-schemas which generate patterns of action and thus can be an implementation of an intention. Subsequently we analysed motor-processes and motor-images and argued that a motor-image might serve as that what is aimed at in an intention. We also noted in the examples of a flight of stairs that a motor-image implies some awareness of the subject. The motor-process approach traces back to William James's notion of *ideo-motor* action. When discussing this notion James [1890] notes that usually we do not pay attention to the action, we just perform, nevertheless on some occasion we do explicitly consider them. In fact James uses the notion of *ideo-motor* action in relation to routines which we have already acquired. We have considered Zen-Buddhism literature where it describes how to master new courses of actions and turn them into routines. It turns out that the difficulty is to overcome awareness of acting.

We have studied the art of moving physical objects. This art involves inter-subjective experiences, however we can stay away from personal and emotionally coloured experiences. Nevertheless, studying swordmanship and ball juggling reveals that quite complicated processes are involved, which also relate to consciousness. In particular Zen Buddhism makes this clear. We conclude that when trying to get hold of consciousness and awareness, the question "how physical

³² China has a long tradition of juggling. It is interesting to note that in the old days, Syrian jugglers in particular were much admired by the Chinese and many Syrian jugglers travelled along the Silk Route to end up in China [Franck and Brownstone, 1986].

processes give rise to subjective experiences" is not the issue, because in our view subjective experiences are physically based. The issue is, and this might turn out to be a hard problem, to get a hold on the different types of subjective experience involved when moving physically.

6.6 Acting in a Physical World

Above we discussed cognitive faculties of an individual creature and the view which takes the mind to be a symbolic reasoning faculty only. We showed that symbolic reasoning systems are closed; the view leads to solipsism. Some understand AI as the science of building intelligence by using computational tools only, that is they work only within the symbolic system model. We conclude, as does Traiger [1991] that this view of AI suffers from solipsism. However, the threat of solipsism is avoided with the approach which we have advocated. Moreover, the approach provides the basics for common knowledge and communication.

6.6.1 Beyond Solipsism

We have argued that adding the knowledge of one agent (with his peculiar type of sensors) to that of another might lead to inconsistency. Thus, every creature inhabits a specific cognitive world, which is determined by its particular types of senses. For instance, a simple sensory-based robot cannot identify objects as the sensor data only yield properties, not individuals. The robot designer, being a human, in general thinks in terms of objects and individuals, which is quite different.

However, individuals of the same species share the same types of senses and the same types of motor-processes (and motor-images). We have groups of agents that share much of their cognitive and motor-image bases and thus we can define *common cognitive* and *motor-image bases*. The definitions of cognitive basis and motor-image basis are indexed by the agent r ; by letting the index range over the group, the shared common bases are obtained. The common cognitive basis and the common motor-image basis are the intersections of the respective individual bases.

Also, though each species inhabits its own cognitive world, these worlds are not fully separate. All species together share the physical world, as their common playground. In this playground creatures of different species interact. Remember the staircase example above and the lizards taking a sun-bath. A robot designer/builder also breaks through the balloon of solipsism. In the common physical playground the robot designer evaluates the adequacy of his robot. Whether the robot represents the world correctly is not the point. As long as the projected actions are properly executed, the designer can and will suppose that his design

works well. In nature, the interaction between creatures of different species is physical, and the success of any creature to survive, is determined by its ability to act physically.

6.6.2 Communication and Niche

Living creatures act in a shared physical world, some species can create their even world. For example, bacteria and other low-level organisms create their own ecological niche. Creation of a (species specific) niche opens possibilities for low level interaction.

Recall that, according to information theory, the amount of information carried by a signal depends on how many alternatives the receiver knows to be possible at the source of the signal. Chaib-Draa and Levesque [1996] define communication on the basis of signs: *"the sender makes a sign which refers to some state of the environment and which has the end of letting the receiver know the same reference"*. This definition presupposes that sender and receiver share references to (potentially) the same alternatives. If, moreover, the number of alternatives is limited, communication is easily established.

The cognitive and motor-image bases contain the possible alternatives. Creatures of the same species with limited bases, have few possible but common interpretations for observations. As a consequence, they react comparably and thus interact easily. We have seen examples in the area of autonomous robots in Chapter 5. The robots were constructed as congenial agents, they can only perform a limited range of motor-schemas, and quite straightforwardly interpret actions of fellows correctly, and interact with each other.

Members of the same species basically share the same motor-images, and thus there is a mutual basis for understanding. The boy-girl example of Bateson [1973] cited above, is quite illustrative. This also confirms Annett's [1996] remark: *"It can be argued that an action representation system is of considerable value to a social species in interpreting the actions and intentions of other members of the group"*. When agents regularly act and interact one might speak of a, though very basic or primitive, social system. Some define such interaction as low-level communication: *"A structural coupling between entities resulting in co-ordination of their actions necessary for their mutual viability"* (Maturana, cited in [Stewart, 1995]).

Because of the limited but species specific view of the world, individuals of the same species can interact (or communicate) rather easily. This conclusion underlines the importance of exploring, what Sloman [1994] – referred to in the introduction – calls the *niche* of the creature. We have discussed autonomous robots and the niche within which they act. The robots are not designed to solve a specific high level problem, instead they react to the environment. They are adapted

to a specific domain of application. A description of the robot's intentions within the context of application, can be considered as an exploration and description of the niche of the robot [Sloman, 1994]. When exploring the niche in which a natural creature is acting, one is (indirectly) describing the cognitive basis and the motor-process basis of the creature, these bases and the niche go hand in hand.

6.7 Conclusions on Cognition

Perception and cognition are prerequisites for autonomous performance. Robotics, biology and philosophy provide several concepts and notions concerning perception, cognition and mind. In this chapter we attempted to order the pieces from our point of view. The point to explore was how an internal physical state of a creature, artificial or natural, relates to action and how this leads to cognition. Intentions and motor-images play central roles.

The engineering (or design) perspective of robotics, advantageously enforced us to be pragmatic, and we defined motor-schemas and intention. An intention denotes a pattern of actions performed by motor-schema such that the mechanism itself tends to bring about an internal state (which is not yet the case). However, a robot acts with intention only if it has some notion of the results of its own action. In natural beings motor-processes are the equivalent of a motor-schema. Motor-processes associate with motor-images. The motor-image is the cognitive content of the motor-process. Thus, the motor-image can be what the creature aims to realise in the intention. In the perception-action perspective, thinking can be conceived of as the process of examining previously established images.

We revisited our notion of cognitive basis and defined a motor-image basis. This basis refers only to perception-action based images and provides a first foothold for developing an epistemology which makes no reference either to an internal or to an external observer. Moreover, we implicitly defined the notion of object: an object is an object – for a creature – only if the senses and the motor-processes of the creature are sensitive to it.

We did not aim at giving definitive answers to the questions about mind, thinking and consciousness. However, we took a position against opposing mind and body. Instead, we opposed imagery to abstract reasoning. Some essential differences between the two modes of thinking were pointed out, and we argued that abstract reasoning cannot capture the full concept of cognition. We argued that consciousness cannot be captured by abstract or verbal reasoning and showed that it is a phenomenon which is empowered to manipulate the organism itself. Thus consciousness definitely has to be related to the body. We conclude that an explanation of consciousness has to explain how it is possible that we can take an abstract point of view and that we become aware of perceiving and acting. But whatever

the explanation, it has to account for the fact that conscious reasoning ultimately might result in performing actions.

6.7.1 The Problem Statements

The notions of cognitive basis and motor-image basis are the basis of interaction and communication. We started this thesis with providing a structure for ordering actions of agents into a supply chain. A common motor-image basis enables autonomous actors to operate as (more or less) congenial agents, and develop (either explicitly or implicitly) protocols for behaviour. The latter are essential for the development of a structure of operations. Thus it is shown that though the arrangement of a supply chain (problem statement 1) can be dealt with as a separate problem, the basics for any structure of the chain generates from the peculiarities of the agents involved (the latter point is questioned in our problem statement 2).

Although, we did not provide definite answers on cognition and intelligence, at least we made a profitable start by restricting them to the context of (physical) actions. To conclude, we remark that in a sense, Artificial Intelligence has just begun to assess the importance of cognitive notions for *artificial social* creatures. Our study shows that the art of moving physical objects includes some basics of intelligence.

7 Conclusions

We study *the practical art of moving physical objects*. The art is introduced in Chapter 1, it concerns *supply chains*. The subject consists of moving the objects, but also of properly *arranging* the moving activities. Arranging activities requires a global view, while in performing the acting a local view is involved. These points of view are encapsulated into two problem statements:

Problem statement 1. *How can the activities of the actors in a (large) supply chain be arranged such that a (goal) directed flow of objects results?*

Problem statement 2. *How can a set of resources be made to perform as an actor within a supply chain?*

Problem Statement 1

In the chapters 2 and 3, problem statement 1 is explored. A supply chain consists of a web of activities, and has an inherent ordering principle which forces the flow of objects into the appropriate direction. This is illustrated in Chapter 2 by surveying historical examples in the domains of warfare and mail services. Besides the inherent ordering for moving the objects, the actors performing in a supply chain are subject to an organisational model. Three basic models are discussed: the silken thread, the authority-based organisation and the contract net. From investigating the examples we conclude that none of these models is particularly designed for a supply chain.

A supply chain basically consists of a concatenation of operations; this concatenation is obtained via co-ordination. We introduce a structure of operations. It consists of a set of operations (activities), necessary for moving the objects, and a framework of conditions which gives the ordering of the operations into a chain. The historical examples reveal that the operations are performed naturally by the resources. The difficult point, however, is to set up the ordering of the operations into a chain and to maintain the chain.

In order to develop a design method, task decomposition and the design of a framework of conditions are further investigated. An operation transforms a certain start-state into an end-state. The framework of conditions orders the operations by arranging these states. Though operations may develop naturally, a systematic approach requires that operations are independently executable. Having sufficient independence among tasks, a framework of conditions can be obtained using state-based planning methods. The design method is worked out for the case of a pan-European distribution service. In practice there is always some tension between managing the supply chain and managing the resources. The balance to be sought depends on how the values of the resources are ranked with respect to the results expected from the chain. In the case of the pan-European distribution service the principle of service management applies: the actors are directly rewarded for how many objects flow through the chain. In the distribution service, the chain manager is called Logical Service Supplier. He composes and plans the framework of conditions, moreover, he guards and maintains its execution. However, actors in the distribution service can change roles, for some shipments they are the Logical Services Supplier for others they are the resource manager.

Regarding the global view implicit in problem statement 1 we show that the well functioning of a supply chain is rather independent of how actors are organised. The structure of operations, inherent to a supply chain, is separable from the organisational structures. Whether the chain performs well, depends to a large extent on the know-how of the actors and their autonomous performance.

Problem Statement 2

In the chapters 4, 5 and 6 we investigate resources (agents) performing autonomously. To perform as an actor, local agents need to know how to execute a task. Robots are typical instances of agents designed to perform an action. The discussion explores how a robot can be made to perform as an actor.

A robot reacts to its environment. In sensory-based robotics, the reactive process is captured in the concept of a motor-schema. In a motor-schema sensor information directly triggers actions. A motor-schema is typically designed to solve a particular task and implies know-how and thus exemplifies autonomous performance. Sensor data and the information provided with them are further explored. To get hold on the sensor information we define a cognitive basis and analyse the problem of sensor fusion. The later problem shows the importance of the designer. He sets the goals and aims of a motor-schema by connecting a certain sensor input to a certain reaction. The process which he thus implements is an intention. The robots are applied on a cargo terminal to perform sorting. In this example, neither the robot nor its motor-schemas have "insight" into why they are performing, and we conclude that a performing agent need not know anything about the global view to which it connects.

In Chapter 5 we consider robots operating in an environment with fellows. The focus is on the spatial and geometrical aspects of the intentions. Teams of robots are distinguished on the basis of the (spatial) intentions applied by the individual robots. Teams of reactive agents generate interaction which in turn causes dynamics in the teams. We study teams of avoiding agents and of antagonist; the first avoid others whereas the latter chase them. The teams of avoiding agents scatter and the teams of antagonists cluster together. When adding a blindly moving agent to a team of antagonist the team starts following it as if it is their leader. Teams exhibiting such connected behaviour, can be applied on a cargo terminal. Each agent performs its task autonomously, while the team behaviour ensures that these separately executed tasks fit into a whole.

In Chapter 6 we make a shift from robots towards natural beings and humans. Natural beings also perform autonomously and interact with their environment: they have senses, they do perceive, but they also have cognition. In natural beings motor-processes are the equivalent of a motor-schema. Motor-processes associate with motor-images. The motor-image is the cognitive content of the motor-process. Thus, the motor-image can be what the creature aims to realise in the intention. We define the notion of a motor-image basis. It incorporates the cognitive basis with actions and acting. Interaction and communication are founded on the cognitive basis and the motor-image basis.

The essence of the art of moving physical objects is that there exists a structure for ordering actions of agents into a (supply) chain. We searched for the origin of this structure. A common motor-image basis enables autonomous actors to operate as (more or less) congenial agents, and develop (either explicitly or implicitly) protocols for behaviour. These protocols are essential for the development of a structure of operations. Thus it is shown that though the arrangement of a supply chain (that is, problem statement 1) may be dealt with as a separate problem, the basics for any structure of the chain generates from the intention and actions of the agents involved. We show that these intentions have a local focus which needs not be changed. To arrange resources into a supply chain, that is to answer problem statement 2, the chain must be arranged such that it coincides with the local points of focus.

Appendix: Potential-Field Procedure

In this appendix we give an analysis of the interaction in a team of two robots applying potential-field avoidance. We use the interaction model developed in Section 5.3; in particular theorem 5.1. The notations used in this appendix conform with the notations introduced in Chapter 5.

In the calculations, an object is represented as a repulsive force, whereas the goal of the robot is represented as a constant force of attraction $\bar{F}$. The repulsive forces $\bar{R}$ decrease with increasing distance among the robots. At each point of time, the new velocity vector of the robot is determined as the sum of these forces; we repeat:

Definition 5.14. *Force sum.* $\bar{F}^* = \bar{F}_r + \sum_{o \in d} \bar{R}_{ro}$

in which: $\bar{F}^*$: new velocity vector of the robot r ;

$\bar{F}_r$: attraction of the goal of the robot r ;

$\bar{R}_{ro}$: repulsive force exerted by object o on the robot r .

Figure A.1 gives a graphical representation.

Notations

Robots are indicated by letters A and B ; their respective goals are indicated by G_A and G_B . The force of attraction $\bar{F}$ and the repulsive forces $\bar{R}$ are indexed as follows, $\bar{F}_A$ is the attraction exerted on A ; $\bar{R}_{AB}$ is the repulsion on A caused by B . Line segments are denoted by square brackets $[p, q]$, see Figure A.2. Note that $F_A = F_B$; their magnitude is constant for every point in the plane. Moreover, as A , B and C are congenial

$$\bar{R}_{AB} + \bar{R}_{BA} = 0.$$

Figure A.1: Potential-field summation.

Similar to the approach in Chapter 5, we look upon the robots as a two-particle system, which has a Centre of Gravity (CG). The repulsive force which the robots exert on one another are the internal forces of the system. These forces do not affect the Centre of Gravity (Newton's third law). The attractive forces of the goals on the respective robots, are external forces. The external forces cause the Centre of Gravity to move.

External forces

To start, we consider the system as internally fixed and observe where CG will move to. After the robots have reached their respective goals, CG will be in the middle between G_A and G_B . This point is denoted by ECG . To compute the trail of the point CG one should solve the equations $F_{CG} = F_A + F_B$ for each point in the plane. As the result is dependent on the positions of A and B , respectively, the resulting force F_{CG} is not unique for a point. However, the behaviour of the robot system can be clarified by the following observation for a system of two robots, A and B , which is in a real collision or semi-real collision:

Observation A.1. *End position of the Centre of Gravity. If CG is not on the baseline through G_A and G_B the force F_{CG} will make CG approach ECG .*

Proof. CG will move as a result of the external forces exerted on A and B respectively. Moreover in any real collision or semi-real collision, A and B will be in the same half-plane with respect to the base line $[G_A, G_B]$; otherwise there would not be an intersection point p as demanded by the definitions of real and semi-real collisions. $\square$

We observe the following:

Observation A.2. *Whenever A and B are in the same half-plane with respect to the baseline of $[G_A, G_B]$, the sum vector $F_{CG} = F_A + F_B$ will have a non-zero vector (component) pointing towards (some point on) $[G_A, G_B]$.*

Proof. If either A or B is on the baseline of $[G_A, G_B]$ (but not both) we obtain the same fact. It is obvious that vector $F_{CG} = F_A + F_B$ can be resolved into two components, one of which is a non-zero vector which has as its base line $[CG, ECG]$, thus CG will move towards ECG , which clarifies A.2. $\square$

Assume, that in a (semi-)real collision, contrary to the condition of observation A.1, CG is on the baseline of $[G_A, G_B]$, but not at ECG , and A and B are also on it. From observation A.1 it is clear that the robots will not reach this position in a natural course of events, they have to be forced. We call this an *artificial case*. Even more particular is the situation when on one line, one robot, which has reached its goal, is between the other robot and this robot's goal. As the forces have equal constant magnitude, in the artificial case the sum vector will be zero at all points between G_A and G_B . From this it is easy obtain:

Observation A.3. *In a balanced conflict CG will, (after some time) be in rest.*

Proof. The motion of the system (A, B) is caused by the resulting force on CG . When CG is moving observations A.1 and A.2 apply and at least one robot approaches its goal. Thus CG must be in rest, this happens when CG reaches a point where the sum of the external forces is zero. $\square$

Now we continue with non-artificial balanced conflicts and show that they can originate only from situations in which the robots are in a certain symmetric position with respect to the goals. First we need to introduce some more notations (see Figure A.2): α is the angle between the positions of the robots $[A, B]$ and $[A, G_A]$, that is the direction of the attractive force of the robot A to its goal. As positions vary with the different time points t they will be indexed as α_t , $[A_t, B_t]$ and $[A_t, G_A]$ respectively. β_t denotes the angle between $[A_t, B_t]$ and $[B_t, G_B]$. In a real or semi-real collision there exists an intersection point p_t between $[A_t, G_A]$ and $[B_t, G_B]$, by assumption. In the proofs we will show that the system of robots is turning with respect to $[G_A, G_B]$; for simplicity we evaluate the behaviour of the system with respect to the line $[\perp, ECG]$, which passes through ECG and is perpendicular to the baseline $[G_A, G_B]$. We call γ_t the angle between $[A_t, B_t]$ and $[\perp, ECG]$.

In a real collision the robots A and B have to traverse $dist(A, p)$ and $dist(B, p)$ at the same time, so we have:

Observation A.4. In a real collision $dist(A, p)/F_A = dist(B, p)/F_B$.

If $\text{dist}(A, G_A) = \text{dist}(B, G_B)$ and $\alpha = \beta$ the situation is called *symmetric*.

Lemma A.5. *The potential-field procedure preserves symmetry.*

Figure A.2: Notation.

Proof. A situation t is symmetric iff at t , $\alpha = \beta$ and $\text{dist}(A, G_A) = \text{dist}(B, G_B)$. These conditions remain valid in the next situation $t+1$ if and only if the line segment $[A_{t+1}, B_{t+1}]$ is parallel to $[A_t, B_t]$. As the repulsive forces between the robots do not influence the translation of the system AB we focus only on the attractive forces. Because $\text{dist}(A, G_A) = \text{dist}(B, G_B)$ as well as $F_A = F_B$ and $\alpha_t = \beta_t$ we have:

$$\text{dist}(A_t, [A_{t+1}, B_{t+1}]) = (F_A * \sin \alpha_t) = (F_B * \sin \beta_t) = \text{dist}(B_t, [A_{t+1}, B_{t+1}]).$$

It is clear that $\alpha_{t+1} = \beta_{t+1}$ and $\gamma_t = \gamma_{t+1} = \pi/2$, which proves A.5. $\square$

Lemma A.6. *Semi-real collisions are solved by the potential-field procedure.*

Proof. We proof A.6 by showing that the system AB is turning with respect to $[G_A, G_B]$. From this it is clear that the series contains a targeting series as required in theorem 5.1. Assume $\alpha_0 \geq \beta_0$. Note that in a collision at least one of the angles α_0 or β_0 is smaller than 90° . We have either that $90^\circ > \alpha_0 \geq \beta_0$ or $\alpha_0 > 90^\circ > \beta_0$, in the latter case, because an intersection point must exist, $\beta_0 < 180^\circ - \alpha_0$; so in both cases $\sin \alpha_0 \geq \sin \beta_0$. The attractive forces have a constant magnitude: $F_A = F_B = \text{constant}$. Note that if $\alpha_0 = \beta_0$ the situation is a real-collision. Hence, we have:

$$\text{dist}(A_0, [A_1, B_1]) = (F_A * \sin \alpha_0) > (F_B * \sin \beta_0) = \text{dist}(B_0, [A_1, B_1])$$

which implies $\gamma_0 > \gamma_1$, if $90^\circ > \alpha_0 \geq \beta_0$, and $\gamma_0 < \gamma_1$, if $\alpha_0 > 90^\circ > \beta_0$. It is not difficult to show that at t_1 $\alpha_1 > \beta_1$. Thus the γ 's form a strictly decreasing (or increasing) se-

ries. This means that the system is turning (in a constant direction) around its CG with respect to $[G_A, G_B]$; this proves A.6. $\square$

Figure A.3: Right: a Symmetric real collision; left evolution of a semi-real collision.

Lemma A.7. *Non-symmetric real-collisions are transformed into semi-real collisions at the next point of time.*

Proof. A non-symmetric real collision is a real collision (so observation A.4 holds), but the situation is not symmetric, thus either $\alpha \neq \beta$, or $dist(A, G_A) \neq dist(B, G_B)$, or both inequalities hold. It is obvious that a non-symmetric real-collision occurs only if $dist(A, p)$ is equal to $dist(B, p)$, so in the triangle (ABp) , $\alpha = \beta$ and thus it must be the case that $dist(A, G_A) \neq dist(B, G_B)$. It follows that if at time t_0 we have a real collision the equation holds:

$$(i) \quad (F_A * \sin \alpha_0) = (F_B * \sin \beta_0)$$

Equation (i) implies that if at time t_0 the robots are on the line $[A_0, B_0]$, then at the next point of time t_1 they must be on the line $[A_1, B_1]$, parallel to $[A_0, B_0]$. If at time t_1 equation (i) does not hold, we have a semi-real collision which according to lemma A.6 is solved. So, we need only to understand what happens at t_1 .

We will call the end-points of the attractive forces, acting on the robots, A' and B' respectively (see Figure A.4). It is clear that they are the intersection points of $[A_1, B_1]$ with $[A_0, G_A]$ and $[B_0, G_B]$. As (i) states, $[A', B']$ is parallel to $[A_0, B_0]$. Due to the repulsion between them, robots A and B will translate in opposite directions along the base line of $[A', B']$ and take the positions A_1 and B_1 .

Now we look at the triangles $(G_A A' A_1)$ and $(G_B B' B_1)$, see Figure A.4. Since $\alpha_0 = \beta_0$, the angle $(G_A A' A_1)$ is equal to the angle $(G_B B' B_1)$ and $dist(A', A_1) = dist(B', B_1)$.

However, as $dist(A_0, G_A) \neq dist(B_0, G_B)$ and because $dist(A_0, A')$ is equal to $dist(B_0, B')$, $dist(A', G_A)$ is not equal to $dist(B', G_B)$. It is obvious, that at $t_1, \alpha_1 \neq \beta_1$ and we have a semi-real collision. This proves A.7. $\square$

Corollary A.8. *Two consecutive situations are real-collisions iff both are symmetric.*

Proof. Obvious from lemma A.5 and A.7. $\square$

Theorem A.9. *Series of symmetric real-collisions are invariant for the potential-field procedure.*

Figure A.4: A non-symmetric real collision.

Proof. That the artificial case is invariant for the procedure is clear, so we continue with the other cases. Symmetry is preserved, as it was proved in A.5, thus series of symmetric real-collisions are invariant. $\square$

From the above we obtain the main proposition:

Proposition A.10. *For the potential-field procedure, when dealing with constant forces of equal magnitude, balanced conflicts coincide with a series of real-collisions.*

Summary

We have studied *the practical art of moving physical objects*, it is concerned with *supply chains*. Supply chains are probably as old as human society. A comprehensive example is the historical Silk Route, but we also met a supply chain in the metaphor of an ant colony. In these examples many active agents together form a supply chain. The motive for our research was to find out what binds all the activities into a chain.

Though *moving* the objects is the kernel activity, having the objects moving is not simple: a supply chain is a subtle linking of relatively small physical activities. The art of moving physical objects consists of both moving the objects, and properly arranging the process pieces. We encapsulated these two aspects in two problem statements, one expressing the global point of view and the other the local point of view.

Problem statement 1. *How can the activities of the actors in a (large) supply chain be arranged such that a (goal) directed flow of objects results?*

Problem statement 2. *How can a set of resources be made to perform as an actor within a supply chain?*

Problem Statement 1

In Chapter 2 we distinguished three organisational models: silken thread, authority-based organisation, and the contract net, and surveyed historical examples of supply chains in warfare and in the mail services.

The warfare examples showed the importance of a well-functioning supply chain: proper supply of the army is decisive in war. As to the organisation model we noted that an army is typically authority-based, but that this did not guarantee proper supplies. In the history of mail services we encountered all three organisational models. The mail connection Amsterdam-Hamburg was an example of what we have labelled a silken thread. The French La Poste had an authority-based or-

ganisation, similar to the military organisation. The mail services that developed within the Low Countries were organised by a contract net. In each of these cases the supply chain worked, thus proving that the organisational model is not essential; a supply chain is not characterised by its organisational structure. Essentially a supply chain consists of a web of activities, and its correct functioning depends on a proper arrangement of the activities.

A silken thread results when actors happen to link to each other. The existence of the thread depends on each of the actors individually. The authority-based organisation is developed from the power of certain actors. The French La Poste derived its existence from the king. The advantage of an authority-based organisation is that the authority could make sure that actors and their resources are available. A contract-net model in a supply chain presupposes that the chain is clearly ordered and transparent. In the background some gremium is required that establishes "standards". Given this, actors can participate in an open competition. However, the establishment of the gremium requires a long lasting process of evolution and development.

To reveal the fabric of the supply chain we introduced in Chapter 3 our notion of a structure of operations, consisting of a framework of conditions and a set of operations. An *operation* is an activity that changes some pre-conditions into certain post-conditions. The post-conditions of one operation can be matched to the pre-conditions of another one. A supply chain consists of an arrangement and ordering of pre- and post-conditions, that is a framework of conditions. A complicated operation is achieved as a result of accurately planning the co-ordinates of actions performed by single agents.

We applied our analysis to the case of Logistic Service Suppliers who intended to set up a pan-European distribution service. We defined the role of chain manager, also called Logical Service Supplier, who has to compose, guard, and maintain the framework of conditions. The structure of operations is the co-ordination framework among the partners in a service. Partners collaborate simply by co-ordinating within the structure of operations. To execute operations, resources and in particular agents are needed. The *resource manager* prepares the resources with which he participates in the supply chain.

With respect to problem statement 1 we have shown that the correct functioning of a supply chain is quite independent of how actors are organised. A supply chain has an inherent structure of operations, which can be separated from the organisational model. The performance of a chain depends to a large extent on the know-how of the actors and their autonomous execution of the operations. Autonomy is not so much a nice to have characteristic, it is inevitably necessary in a supply chain. Autonomy implies that actors and agents know how to perform, this is captured in our problem statement 2.

Problem Statement 2

In Chapters 4, 5 and 6 we investigated resources (agents) performing autonomously. To perform as an actor, the local agents need to know how to execute a task. We studied robots as resources in the context of a chain of actions. A robot is often called an autonomous agent, whereas in fact the robot only has know-how and performs autonomously. Perception and cognition are prerequisites for autonomous performance. An analysis of how to make a robot perform answers problem statement 2, and involves a local point of view.

In Chapter 4 we discussed the design approach called sensory-based robotics, and explored how sensor data are used to enable autonomous performance. The approach works with *motor-schemas*, by which the robot reacts to its environment. Motor-schemas are triggered by their sensors. Once a schema is triggered, it results invariably in an action. The motor-schemas are incorporated into a control architecture. While each motor-schema reacts to its sensor inputs, compound behaviours generate when more schemas are triggered. The control architecture ensures that only one schema at a time will be executed. The difficulty is to ensure in advance that the best motor-schema is triggered and selected.

Clearly, sensors are at the basis of any action. In order to understand how sensor data determine the view of the environment, we introduced the notion of a sensor basis. Furthermore, we investigated refinement of sensor data, which can be obtained by combining data from different sensors using sensor-data fusion techniques. The sensor basis is quite helpful in analysing the inherent problems of sensor data fusion. A central problem is that whereas actions in general involve the manipulation of objects, sensors do not specify objects but only provide data on attributes of objects. The keys to drawing the conclusion that an object exists must be provided by the robot designer. The definition of sensor basis is extended to obtain the *cognitive* basis. And we conclude that the cognitive basis of a robot might differ considerably from the designer's: objects exist in the designer's cognitive basis but not in the cognitive basis of the robot.

In Chapter 5 we considered a robot operating in an environment with fellow robots, with the focus on the spatial and geometrical interaction of the robots and the resulting team behaviour. Robots are instances of spatially performing agents. In the teams the robots act and interact autonomously, without communicating. We defined a notion of intention in terms of patterns of spatial actions. Our robots apply only two types of intentions: goal finding and obstacle avoidance.

Intentions and the resulting behaviour of the individuals can conflict in a team, thus team behaviour evolves. Congenial teams are teams in which the members apply the same intentions. We discussed team behaviour of several congenial teams. Teams of reactive agents generate interaction, and we distinguished avoiding teams and antagonistic teams. The main team characteristics disclosed are the scattering behaviour of avoiding teams and the clustering behaviour of antagonis-

tic teams. The scattering behaviour of congenial avoiding teams is further analysed. Central is the notion of conflict. A balanced conflict prevents a team from making any progress. The designer of a team, of course, seeks to avoid balanced conflicts. We studied two types of avoiding teams mathematically, one team applying flee-opposite avoidance and the other team basing avoidance on an artificial potential-field. The flee-opposite procedure generates series of repetitive movements, and is attractive because of its easy implementation: it requires hardly any processing. The potential-field procedure generates a smoother path. Moreover, in potential-field-based avoidance, goal finding, and obstacle avoidance are integrated in one calculation, so that the approach is mathematically transparent and can be handled analytically.

In the analysis of a team applying potential-field avoidance, we concluded that balanced conflicts form a closed category. If the team starts in a situation that may lead to a balanced conflict, all subsequent situations will be of this kind. However, a team starting in any different situation never encounters a balanced conflict. This conclusion is proven completely for a system of two robots; for a three-robot system it is proven that balanced conflicts exist, that a team may turn around an axis, and that two robots might push the third one forwards or backwards (lemmas 5.3, 5.4 and 5.5). A further and thorough mathematical analysis should prove that these cases indeed fully cover the behaviour of a three-robot team.

Our mathematical approach provided insights in team development, and revealed that the basics of team behaviour can be traced back to the intentions of the team members. Team behaviour evolves without an overall plan being available. Thus, interaction is a starting point for arranging a chain of actions as part of a supply chain.

In Chapter 6 we argued from robots towards natural beings and humans. Natural beings also perform autonomously and interact. In our robotics discussion the underlying view of cognition is action-driven. This view is also elaborated for natural creatures and resembles the view held by Aristotle. The discussion focused on perception and intentions and explored how cognition of a creature relates to action. Intentions and motor-images play central roles. We briefly reviewed distinct notions concerning the mind: abstract reasoning, consciousness and intellect. We showed with examples that an abstract reasoning approach cannot explain intentions as the drive for actions and thus cannot capture the full concept of cognition.

The engineering perspective of robotics forced us to be pragmatic, and we defined under what conditions a motor-schema implements an intention without appealing to any abstract nor mental concept. However, in Chapter 4 we noted the problem that objects exist in the frame-of-reference of the designer, but not automatically in the cognitive basis of his robot. Intentions denote patterns of actions. In a motor-schema sensors (perception) directly cause actions. When performing a motor-

schema the mechanism itself tends to bring about an internal state. A motor-schema implements an intention if there is feedback on the results of the action.

In natural beings motor-processes are the equivalent of a motor-schema. In psychology, the notion of motor-process refers to a pattern of actions and includes a (mental) *motor-image* as well. The motor-image is the cognitive content of the motor-process. We approached the faculties of thought and thinking as processing motor-images. In the perception-action perspective, thinking, particularly in the sense of imagery, can be conceived of as the process of examining previously established images. Thus, the motor-image can be what the creature aims to realise in the intention. Therefore, we returned to our notion of cognitive basis and defined a *motor-image basis*. This basis refers only to perception-action based images and provides a first foothold for developing an epistemology which makes no reference either to an internal or to an external observer.

We also reconsidered the role of consciousness and in particular awareness in action. Consciousness and awareness enable humans to transform abstract notions into actions. For this, we looked briefly at the training methods in Zen Buddhism and the eastern disciplines of martial arts. We did not aim to give definitive answers to questions about mind, consciousness and awareness. However, we took a position against opposing mind to body, and concluded that an explanation of consciousness only suffices if it can explain how it is possible that human beings can take an abstract point of view and that they become aware of perceiving and acting. Moreover, it has to explain how reasoning about actions ultimately might result into performing actions. We concluded the discussion by showing that a motor-image basis provides a basis for mutual understanding and thus is useful foundation for studying interaction and communication.

Samenvatting

De kunst van het verplaatsen van fysieke objecten is een pragmatische vaardigheid en heeft betrekking op toeleveringsketens. Toeleveringsketen bestaan al zo lang als de menselijke samenleving. Een veelomvattend voorbeeld is de historische Zijderoute, daarnaast vinden we toeleveringsketens in de metafoor van een mierenkolonie. In deze voorbeelden zijn vele agenten actief die samen een toeleveringsketen vormen. Het motief voor ons onderzoek was om na te gaan wat al de activiteiten bindt tot een keten.

Hoewel het verplaatsen van objecten de kernactiviteit is, is het niet eenvoudig om de objecten in beweging te houden: een toeleveringsketen is een subtiële aaneenschakeling van relatief kleine activiteiten. De kunst van het verplaatsen van fysieke objecten bestaat zowel uit het verplaatsen van de objecten als uit het aaneenschakelen van de processtukjes. Beide aspecten zijn vervat in twee probleemstellingen. Waar de eerste uitgaat van een globaal gezichtspunt, gaat de tweede uit van een lokaal gezichtspunt.

Probleemstelling 1. Hoe kunnen activiteiten van de actoren in een (lange) toeleveringsketen zodanig geordend worden dat een (doel)gerichte stroom van objecten ontstaat?

Probleemstelling 2. Hoe kunnen de verzamelingen hulpmiddelen en -bronnen worden aangewend om op te treden als actoren in een toeleveringsketen?

Probleemstelling 1

In hoofdstuk 2 onderscheiden we organisatorische modellen: een zijdedraad, een autoriteit-gebaseerde organisatie en een contractnet en verkennen we historische voorbeelden van toeleveringsketens in de krijgskunst en in de postale dienstverlening. De voorbeelden uit de krijgskunst laten het belang van een goed functionerende keten zien. Voldoende bevoorrading van een leger is beslissend in de oorlogsvoering. Wat de organisatie modellen betreft merken we op dat het leger een typisch op autoriteit gebaseerde organisatie heeft. In de geschiedenis van de post

komen we alledrie de organisatiemodellen tegen. The postale verbinding tussen Amsterdam en Hamburg is een voorbeeld van wat wij een zijdedraad hebben genoemd. Het Franse La Poste had een op autoriteit gebaseerde organisatie, vergelijkbaar met de militaire organisatie. The postdienst die ontstond in de Lage Landen was georganiseerd volgens een contractnet. In elk van deze gevallen werkte de toeleveringsketen, daarmee bewijzend dat het organisatiemodel niet essentieel is en niet karakteristiek is voor een toeleveringsketen. In essentie bestaat een toeleveringsketen uit een web van activiteiten. Of de keten goed functioneert is afhankelijk van een goede ordening van de activiteiten.

Een zijdedraad ontstaat als een toevallige aaneenschakeling van actoren. Het bestaan van de draad hangt af van elke individuele actor. Een op autoriteit gebaseerde organisatie ontstaat uit de macht van bepaalde actoren. Het Franse La Poste ontleende zijn bestaan aan de koning. Het voordeel van de op autoriteit gebaseerde organisatie is dat de autoriteit kan verzekeren dat de actoren en de hulpmiddelen en -bronnen beschikbaar zijn. Een contractnetmodel in een toeleveringsketen veronderstelt dat de keten duidelijk geordend en transparant is. Op de achtergrond is een gremium nodig dat standaarden bestendigt. Zodra dit tot stand is gebracht, kunnen actoren meedingen in een open competitie. Het instellen van het gremium vergt echter een langdurig ontstaans- en ontwikkelingsproces.

Om het weefpatroon van een toeleveringsketen bloot te leggen introduceren we in hoofdstuk 3 ons concept van een structuur van operaties. Deze structuur bestaat uit een raamwerk van condities en een verzameling operaties. Een operatie is een activiteit die bepaalde pre-condities verandert in post-condities. De post-condities van de ene operatie kunnen worden gepast op de pre-condities van een andere. Een toeleveringsketen bestaat als een regeling en ordening van pre-condities en post-condities. Een complexe operatie wordt bereikt via een accurate planning van de coördinaten van door agenten uitgevoerde activiteiten.

Onze analyse is toegepast op de casus van logistieke dienstverleners die van plan zijn om een pan-Europese distributiedienst op te zetten. We hebben de rol van een ketenmanager gedefinieerd, ook wel logische dienstverlener genoemd. Deze componeert, bewaakt en onderhoudt het raamwerk van condities. De structuur van operaties is het coördinatieraamwerk tussen de partners in de dienstverlening. Partners werken eenvoudig samen door zich te schikken in de structuur van operaties. Om de operaties uit te voeren zijn hulpbronnen en in het bijzonder agenten nodig. De manager van de hulpbronnen prepareert de hulpbronnen waarmee hij deelneemt in de toeleveringsketen.

Wat probleemstelling 1 betreft hebben we laten zien dat een goed functionerende toeleveringsketen vrijwel onafhankelijk is van de manier waarop de actoren zijn georganiseerd. Een toeleveringsketen heeft een inherente structuur van operaties die te onderscheiden is van de organisatorische modellen. De prestatie van de keten hangt voor een groot deel af van de know-how van de actoren en de door hen

autonoom uitgevoerde operaties. Autonomie is onvermijdelijk en noodzakelijk in een toeleveringsketen. Autonomie impliceert dat actoren en agenten weten hoe ze een operatie moeten uitvoeren. Dit aspect is vervat in onze tweede probleemstelling.

Probleemstelling 2

In de hoofdstukken 3, 4 en 5 onderzoeken we hulpbronnen (agenten) die autonoom opdrachten uitvoeren. Om als actor op te kunnen treden, moet een lokale agent weten hoe de taak uit te voeren. We bestuderen robots als hulpbronnen in de context van een keten van acties. Robots worden vaak autonome agenten genoemd, waar de robot in feite slechts autonoom opdrachten uitvoert. Perceptie en cognitie zijn een eerste vereiste om autonoom op te kunnen treden. Een analyse ten behoeve van Probleemstelling 2 geeft antwoord op de vraag hoe robots autonoom kunnen worden gemaakt en betreft een lokaal gezichtspunt.

In hoofdstuk 4 bespreken we de zogenoemde sensorgebaseerde robotica-aanpak en verkennen hoe sensorgegevens worden gebruikt ten behoeve van autonoom optreden. Deze aanpak werkt met motorschema's waarmee de robot reageert op zijn omgeving. Motorschemas worden in werking gezet door de sensoren. Is een schema eenmaal in werking gezet dan resulteert dit onvermijdelijk in een actie. De motorschema's worden verenigd in een controle-architectuur. Terwijl elk motorschema reageert op zijn sensoren, ontstaat samengesteld gedrag als er meerdere tegelijk in werking treden. De controle-architectuur zorgt dat slechts één schema tegelijk wordt uitgevoerd. De moeilijkheid is om vooraf te zorgen dat het meest geschikte motorschema in werking wordt gezet en wordt uitgekozen.

Het is duidelijk dat sensoren aan de basis liggen van elke actie. Om te bepalen hoe sensoren de kijk op de omgeving bepalen, introduceren we de notie van een sensorbasis. Verder onderzoeken we het verfijnen van sensorgegevens. Verfijning kan verkregen worden door gegevens van verschillende sensoren samen te voegen, waarbij gebruik gemaakt wordt van specifieke fusietechnieken. De sensorbasis is een hulpmiddel om de inherente problemen van sensorfusie-technieken te analyseren. Een centraal probleem is dat de acties betrekking hebben op het manipuleren van objecten, terwijl sensoren geen objecten specificeren maar uitsluitend data aanleveren over kenmerken van objecten. De sleutel om de conclusie te trekken dat een object bestaat, moet aangereikt worden door de robotontwerper. De definitie van de sensorbasis is uitgebreid om de cognitieve basis te verkrijgen. We concluderen dat de cognitieve basis van de robot aanzienlijk kan verschillen van die van de ontwerper: objecten bestaan in de cognitieve basis van de ontwerper maar niet in de cognitieve basis van de robot.

In hoofdstuk 5 beschouwen we het gedrag van een robot in een omgeving met gelijken. De focus ligt daarbij op de ruimtelijke en geometrische interactie van de robots en het resulteerde teamgedrag. In de teams opereren en interacteren de ro-

bots autonoom, zonder te communiceren. We definiëren de notie van intentie in termen van ruimtelijke actiepatronen. Onze robots passen slechts twee intenties toe: 'doel zoeken' en 'obstakel ontwijken.'

Intenties en het resulterende gedrag van de individuen kunnen conflicteren, zodat teamgedrag ontstaat. In teams van (geest)verwanten passen de leden dezelfde intenties toe. We bespreken het gedrag van verscheidene verwantenteams. Teams van reactieve agenten genereren interactie en we onderscheiden ontwijkende teams en antagonistische teams. De belangrijkste karakteristieken die we blootleggen zijn het verspreidingsgedrag van ontwijkende teams en het samenklittende gedrag van antagonistische teams. Het verspreidingsgedrag van de verwante en ontwijkende teams is verder geanalyseerd. Centraal is de notie van conflict. Een gebalanceerd conflict weerhoudt een team ervan voortgang te maken. De ontwerper van een team probeert natuurlijk om gebalanceerde conflicten te vermijden. We bestuderen twee typen ontwijkende teams met wiskundige middelen. Het ene team ontwijkt door te vluchten in tegenovergestelde richting en het andere team baseert het ontwijken op een kunstmatig potentiaalveld. De 'vlucht tegenovergesteld' procedure genereert series van herhaalde bewegingen en is aantrekkelijk omdat het eenvoudig te implementeren is: het vergt nauwelijks rekenkracht. De 'potentiaalveldprocedure' genereert gladde banen. In de potentiaalveldprocedure worden 'doel zoeken' en 'obstakel ontwijken' geïntegreerd in één berekening. Daardoor is deze aanpak mathematisch transparant en kan analytisch worden bestudeerd. Uit de analyse van een team dat de potentiaalveldprocedure toepast, concluderen we dat gebalanceerde conflicten een gesloten categorie vormen. Als een team begint in een situatie die tot een gebalanceerd conflict leidt, dan zijn alle opvolgende situaties ook van dit type. Start het team daarentegen in enige andere situatie, dan komt het ook nooit in een gebalanceerd conflict. Deze conclusie is volledig bewezen voor een systeem van twee robots; voor een drie-robotsysteem is bewezen dat er gebalanceerde conflicten bestaan, dat het team om een denkbeeldige as draait en dat twee robots de derde voor zich uit of achter zich drukken (lemma's 5.3, 5.4 en 5.5). Een verdere en omvattende mathematische analyse zal moeten bewijzen dat deze gevallen inderdaad het gedrag van een team van drie robots volledig dekken.

Onze mathematische aanpak levert inzicht in teamontwikkeling en laat zien dat het elementaire teamgedrag terug te voeren is op de intenties van de teamleden. Teamgedrag ontstaat zonder dat een omvattend plan beschikbaar is. Interactie is daarom een startpunt om een keten van acties te arrangeren als onderdeel van een toeleveringsketen.

In hoofdstuk 6 redeneren we beginnend met robots naar natuurlijke wezens en mensen. Ook natuurlijke wezens opereren autonoom en interacteren. In onze bespreking van robotica speelt het onderliggende gezichtspunt dat cognitie door actie wordt gedreven. Dit gezichtspunt is ook uitgewerkt voor natuurlijke wezens en lijkt op dat van Aristoteles. We richten ons op perceptie en intenties en verkennen hoe cognitie van een wezen gerelateerd is aan actie. Intenties en motorbeelden

spelen een centrale rol. We behandelen kort verschillende noties betreffende geest (*mind*): abstract redeneren, bewustzijn en intellect. We laten met voorbeelden zien dat een benadering vanuit abstract redeneren niet verklaart hoe intenties aanzetten tot actie, zodat deze benadering geen volledig beeld geeft van het cognitieconcept.

Het constructieperspectief van de robotica dwingt ons om pragmatisch te zijn, en we definiëren onder welke condities een motorschema een intentie implementeert, zonder daarbij te refereren aan enig abstract of mentaal concept. Echter, reeds in hoofdstuk 4 is het probleem gesignaleerd dat objecten bestaan in het referentieraamwerk van de ontwerper, maar daarmee nog niet in de cognitieve basis van de robot. Intenties duiden actiepatronen aan. In een motorschema veroorzaken sensors rechtstreeks acties. Terwijl een motorschema wordt uitgevoerd, streeft het mechanisme ernaar om een interne toestand teweeg te brengen. Een motorschema implementeert een intentie als er een terugkoppeling is van de resultaten van de acties.

Bij natuurlijke wezens zijn motorprocessen het equivalent van een motorschema. In de psychologie refereert de notie van motorproces aan een patroon van acties en het bevat tevens een motorbeeld. Een motorbeeld is de cognitieve inhoud van een motorproces. Wij benaderen de faculteiten denken en gedachten als het werken met motorbeelden. In het perceptie-actie perspectief kan het denken in de zin van verbeelding worden opgevat als het proces van het onderzoeken of doorzoeken van eerder gevormde beelden. Aldus kan een motorbeeld datgene zijn waarnaar in de intentie van een schepsel wordt gestreefd. Daarom keren we terug naar de notie van cognitieve basis en definiëren we een motorbeeld-basis. Deze basis maakt uitsluitend gebruik van perceptie-actie gebaseerde beelden en biedt een eerste houvast voor het ontwikkelen van een kennisleer die niet refereert aan interne of externe waarnemers.

We heroverwegen de rol van het bewustzijn en het bewustworden van acties en handelingen. Bewustzijn en bewustworden stellen mensen in staat om abstracte noties om te zetten in acties. We staan kort stil bij Zen Boeddhisme en oosterse Martiale kunsten waar dit punt uitgebreid aan de orde komt. Het streven is niet geweest om definitieve antwoorden te geven op vraagstukken betreffende geest, bewustzijn en bewustworden. We kiezen echter wel positie tegen de lichaam-geesttegenstelling. We concluderen dat een verklaring van bewustzijn alleen voldoet als deze inderdaad verklaart hoe het mogelijk is dat mensen een abstract gezichtspunt kunnen innemen en dat ze zich bewust worden van waarnemingen en handelingen. Bovendien moet verklaard worden hoe het redeneren over handelingen uiteindelijk resulteert in het uitvoeren van handelingen. We besluiten de verhandeling met te laten zien dat de motorbeeld-basis tevens een basis legt voor wederzijds begrijpen en dus een bruikbaar fundament is om interactie en communicatie te bestuderen.

References

- [Aertsen *et al.*, 1996] F. Aertsen, A.A.Th. de Schepper and G.C.J.M. Vos. *Dynamiek in Logistiek* (in Dutch), Samson, Alphen and den Rijn, 1996.
- [Allen, 1984] J.F. Allen. Towards a General Theory of action and time. *Artificial Intelligence*, 23:123-154, 1984.
- [Annett, 1996] J. Annett. On knowing how to do things: a theory of motor imagery. *Cognitive Brain Research*, 3:65-69, 1996.
- [Arkin, 1989] R.C. Arkin. Motor schema-based mobile robot navigation. *Int. J. of Robotics Research*, 8(4):92-112, 1989.
- [Arkin, 1990] R.C. Arkin. Integrating Behavioral, Perceptual, and World Knowledge in Reactive Navigation. *Robotics and Autonomous Systems*, 6:105-122, 1990.
- [Arkin, 1992] R.C. Arkin. Cooperation without Communication: Multi Schema-Based Robot Navigation. *J. of Robotic Systems*, 9(3):351-364, 1992.
- [Badreddin and Holenstein, 1991] E. Badreddin and A.A. Holenstein. Reflexive collision avoidance in a recursive architecture. *Robotics and Autonomous Systems*, 8(1991):177-189, 1992.
- [Bateson, 1973] G. Bateson. *Steps to an Ecology of Mind*. Paladin, Granada Publishing Limited, 1973.
- [Bay, 1995] J.S. Bay. Design of the "Army Ant" Cooperative Lifting Robot. *IEEE Robotics & Automation Magazine*, March 1995.
- [Beek, 1989] P.J. Beek. *Juggling Dynamics*, PhD-thesis Free University Press, Amsterdam 1989.
- [Berthoz, 1996] A. Berthoz. The role of inhibition in the hierarchical gating of executed and imagined movements. *Cognitive Brain Research*, 3:101-113, 1996.
- [Block, 1961] I. Block, The Order of Aristotle's Psychological Writings, *American Journal of Philology*, 1961 (82) pp 50-77.

- [Bomhof and Penders, 1992] F.W. Bomhof and J.S.J.H. Penders. A Vehicle that uses Acoustic Radar for Collision Avoidance. In *Proc. 25th ISATA Dedicated Conf. on Road Transport Informatics/Intelligent Vehicle Highway Systems*, 293-299, Florence, Italy, June 1992.
- [Bond and Gasser, 1988] A.H. Bond and L. Gasser. *Readings in distributed Artificial Intelligence*. Morgan Kaufmann Publ., San Mateo, California, 1988.
- [Borenstein and Koren, 1991] J. Borenstein and Y. Koren. Histogramic In-Motion Mapping for Mobile Robot Obstacle Avoidance. *IEEE Trans. on Robotics and Automation*, 7(4), 1991.
- [Botelho and Coelho, 1997] L.M. Botelho and H. Coelho. Emotion-Based Attention Shift in Autonomous Agents, *Lecture Notes in Artificial Intelligence*, 1193, Springer Verlag 1997, pp. 277-291.
- [Botter, 1967] C. Botter. *Industrie en Organisatie, een verkenningstocht* (in Dutch). Philips Eindhoven, 1967.
- [Bradshaw, 1997] J.M. Bradshaw (ed.) *Software Agents*, MIT Press 1997.
- [Braspenning, 1997] P.J. Braspenning. Agenten en Multi-Agentsystemen, (in Dutch) *Informatie*, 1997.
- [ten Brink, 1950] E.A.B.J. ten Brink. *Geschiedenis van het Nederlandse postwezen van 1795-1810* (in Dutch). PTT, Den Haag, Netherlands, 1950.
- [Brooks, 1986] R.A. Brooks. A Robust Layered Control System for a Mobile Robot. *IEEE*, 1986.
- [Brooks, 1990] R.A. Brooks. Elephants Don't Play Chess. *Robotics and Autonomous Systems*, 6:3-15, 1990.
- [Brown, 1987] F.M. Brown, *The Frame Problem in Artificial Intelligence*, Morgan and Kaufmann, 1987.
- [Bruijs, 1980] P.A. Bruijs, Afstemmingsproblemen rond het toekomstige postale vervoerbedrijf (in Dutch), *Het PTT-Bedrijf Deel XXII nr. 1-2*, PTT Den Haag, 1980.
- [Budenske and Gini, 1997] J. Budenske and M. Gini, Logical Sensor/actuator: Knowledge-based Robotic Plan Execution. *Journal of Experimental and Theoretical Artificial Intelligence* 9:361-374, 1997.
- [Cesta et al., 1996] A. Cesta, M. Miceli and P. Rizzo. Effects of Different Interaction Attitudes on a Multi-Agent System Performance, Agents Breaking Away. In *Proc. MAAWAW 1996*, Van de Velde, Perram (eds.), Springer Verlag, 1996.
- [Chaib-Draa and Levesque, 1996] B. Chaib-Draa and P. Levesque. Hierarchical Model and Communication by Signs, Signals and Symbols in Multiagent environments. *Journal of Experimental and Theoretical Artificial Intelligence* 8, 1992.

- [Chalmers, 1995] D.J. Chalmers. The Puzzle of Conscious Experience. *Scientific American*, December 1995.
- [Chattergy, 1985] R. Chattergy. Some Heuristics for the Navigation of a Robot. *Int. J. of Robotics Research*, 4:59–66, 1985.
- [Connell, 1987] J.H. Connell. Creature Design with the Subsumption Architecture. In *Proc. IJCAI-87*, 1124–1126, Morgan Kaufmann 1987.
- [van Creveld, 1977] M. van Creveld. *Supplying War, Logistics from Wallenstein to Patton*. Cambridge University Press, 1977.
- [Crick and Mitchison, 1983] F. Crick and G. Mitchison. The function of dreamsleep. *Nature*, 304:111–114, 1983.
- [Culbertson, 1963] J.T. Culbertson. *The Minds of Robots*. Univ. of Illinois Press, 1963.
- [Culbertson, 1976] J.T. Culbertson. *Sensations, Memories and the flow of Time*. The Cromwell Press, Santa Margarita, California, 1976.
- [van Dalen et al., 1987] D. van Dalen, H.C. Doets and H. de Swart. *Sets: Naive, axiomatic and applied*. Pergamon Press, Oxford, 1987.
- [Damen, 1994] J.T.W. Damen. *Flexible Mail Processing — An Application of Integral Postal Engineering*. PhD thesis, University of Technology Delft, April 1994.
- [Damen, 1995] J.T.W. Damen, *Theoretical model*, IDEAS report R&D-IN-96-200, KPN Research Leidschendam, 1995.
- [Decety, 1996] J. Decety. Do imagined and executed actions share the same neural substrate?. *Cognitive Brain Research*, 3:87–93, 1996.
- [DelCanto, 1995] *Proceedings of the 2nd IDEAS Workshop, October 1995*, Elsag Bailey Genoa 1995; also available as KPN Research R&D-IN-95-1409.
- [Dennett, 1987] D.C. Dennett. *The Intentional Stance*. The MIT Press, Cambridge, 1987.
- [Descartes, 1642] R. Descartes. *Meditations on first philosophy*. Translation: E.S. Haldane and G.R.T. Ross, Cambridge University Press, 1968.
- [Dolev et al., 1987] D. Dolev, C. Dwork and L. Stockmeyer. On the Minimal Synchronism Needed for Distributed Consensus. *J. of the Ass. for Computing Machinery*, 34(1):77–97, January 1987.
- [Dretske, 1981] F.I. Dretske. *Knowledge and the Flow of Information*. Basil Blackwell, Oxford, 1981.
- [Flanagan and Güzeldere, 1997] O. Flanagan and G. Güzeldere. Consciousness: a Philosophical Tour, in: *Cognition, Computation and Consciousness*, K.Ito, Y. Miyashita and E. Rolls (eds), Oxford University Press, 1997.
- [Franck and Brownstone, 1986] I.R. Franck and D.M. Brownstone. *The Silk Road: a history*, Facts On File Publications. New York, Oxford, 1986.

- [Fujimara, 1991] K. Fujimara. *Motion Planning in Dynamic Environments*. Lecture Notes in Artificial Intelligence, Springer Verlag, 1991.
- [Gallop, 1975] D. Gallop (translator). *Plato Phaedo*. Translated with Notes, Clarendon Plato series, Oxford, 1975.
- [Gat et al., 1994] E. Gat, R. Desai, R. Ivlev, J. Loch and D. Miller. Behaviour Control for Robotic Exploration of Planetary Surfaces. *IEEE Trans. On Robotics and Automation*, 10(4), 1994.
- [Genesereth et al, 1986] M.R. Genesereth, M.L. Ginsberg and J.S. Rosenschein, Cooperation without communication, *Proceedings AAAI, 1986*.
- [George, 1961] F.H. George. *The Brain as a Computer*. Pergamon Press, Oxford, 1961.
- [Georgeff and Ingrand, 1989] M.P. Georgeff and F.F. Ingrand, Decision-Making in an Embedded Reasoning System, *Proceeding IJCAI 1989* pp. 972-978.
- [van Goor et al., 1990] A.R. van Goor, A.H.L.M. Kruijtzter and G.W. Esmeijer. *Goederenstroombesturing, voorraadbeheer en materials handling* (in Dutch). Stenfert Kroese, Leiden/Antwerpen, 1990.
- [de Graaf, 1996] R.P. de Graaf. *Oorlog om Holland 1000-1375* (in Dutch). PhD thesis, Verloren Hilversum, 1996.
- [Guarino, 1998] N. Guarino, Formal Ontology and Information Systems, in N. Guarino (ed.) *Formal Ontology in Information Systems*, IOS Press 1998.
- [Hamlyn, 1968] D.W. Hamlyn (transl). *Aristotle's De Anima books II,III*. Translated with Introduction and Notes, Clarendon Press, Oxford, 1968.
- [Harman, 1973] G. Harman. *Thought*. Princeton University Press, 1973.
- [Hartman and Roos, 1988] W. Hartman and J. Roos. *Methodiek voor systeemonderzoek, 1* (in Dutch). Kluwer, Deventer, 1988.
- [Henderson, 1990] T.C. Henderson. Logical Behaviors. *Journal of Robotic Systems*, 7(3), 1990.
- [Hestenes, 1986] D. Hestenes. *New Foundations for Classical Mechanics*. Reidel Publishing Company, Dordrecht, 1986.
- [Heyes, 1998] C.M. Heyes, Theory of Mind in Nonhuman Primates, *Behavioral and Brain Sciences* (21), 1998.
- [Hoogeweegen, 1997] M.R. Hoogeweegen, *Modular Network Design, assessing the impact of EDI*, PhD-thesis, Erasmus University Rotterdam, 1997.
- [Houzelle et al., 1994] S. Houzelle and G. Giraudon, Contribution to Multisensor Fusion Formalisation, *Robotics and Autonomous Systems* (13) pp. 69-85, 1994.
- [Hughes and Cresswell, 1984] G.E. Hughes and M.J. Cresswell, *A companion to modal logic*, Methuen, London and New York 1984.

- [Hwang and Ahuja, 1992] Y.K. Hwang and N. Ahuja. Gross Motion Planning – A survey. *ACM Computing Surveys*, 24(3), 1992.
- [Jager, 1968] M. Jager, master thesis, Rijks Universiteit Leiden, september 1968 (in Dutch).
- [James, 1890] William James. *The principles of psychology*. Dover Publications, Inc., 1950.
- [Jomini, 1855] Baron de Jomini. *The Art of War*. Translation of: *Precis de l'art de la Guerre* (1855), Greenwood Press, West port, Connecticut, 1971.
- [Kant, 1787] E. Kant. *Kritik der Reine Vernunft*. 1787.
- [Kanzaki, 1996] R. Kanzali. Behavioural and Neural Basis of Instinctive Behavior in Insects: Odor-source searching strategies without memory and learning. *Robotics and Autonomous Systems*, 18, 1996.
- [Khatib, 1986] O. Khatib. Real-Time Obstacle Avoidance for Manipulators and Mobile Robots. *Int. J. of Robotics Research*, 5, 1986.
- [Koomen, 1996] J. Koomen. *Finding boxes in a heap, using a sonar image*, IDEAS deliverable 4.3.5, R&D-IN-96-435, KPN Research Leidschendam, 1996.
- [Kube and Zhang, 1994] C.R. Kube and H. Zhang. Collective Robotics, From Social Insects to Robots. *Adaptive Behavior*, 2(2), The MIT Press, Cambridge, 1994.
- [Latombe, 1991] J.C. Latombe. *Robot Motion Planning*. Kluwer, Boston, 1991.
- [Lewis and Papadimitiou, 1981] H.R. Lewis and C.H. Papadimitriou, *Elements of the Theory of Computation*, Prentice-Hall, 1981.
- [Lyons and Arbib, 1989] D.M. Lyons and M.A. Arbib. A Formal Model of Computation for Sensory Based Robotics. *IEEE Trans. on Robotics and Automation*, 5(3), June 1989.
- [Malone, 1986] T.W. Malone. *Modelling coordination in organisations and market*. Reprinted in [Bond and Gasser, 1988].
- [Mataric, 1995] M.J. Mataric. Issues and approaches in the design of collective autonomous agents. *Robotics and Autonomous Systems*, 16, 1995.
- [McDowell, 1973] J. McDowell (translator). *Plato Theaetetus*. Translated with Notes, Clarendon Plato Series, Oxford, 1973.
- [Modrak, 1987] D.K.W. Modrak, *Aristotle, the Power of Perception*, University of Chicago Press, Chicago 1987.
- [Nikonoriv, 1997] V.P. Nikonorov, *The armies of Bactria 700 BC-450 AD*, Montvert Publications 1997.
- [Noreils, 1992] F.R. Noreils. Multi-Robot Coordination for Battlefield Strategies. In *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, 1777-1784, 1992.

- [Ogilvie, 1990] R.C. Ogilvie. *Krijgen is een kunst, Omtrent krijgskunde en ondernemingsstrategie*, (in Dutch) Stenfert Kroese, Leiden/Antwerpen, 1992.
- [van Overvoorde, 1900] J.C. van Overvoorde. De centralisatie van het hollands postwezen in het midden der 18-e eeuw, Bijdragen voor vaderlandse geschiedenis en oudheidkunde (in Dutch). Iste reeks, deel 1, 1900.
- [van Overvoorde, 1902] J.C. van Overvoorde. *Geschiedenis van het postwezen in Nederland voor 1795*. Leiden, 1902 (in Dutch).
- [Parker, 1979] G. Parker, *Spain and The Netherlands, 1559-1659*, Collins, London 1979.
- [Peddie, 1994] J. Peddie, *The Roman War Machine*, Sutton publishing, Gloucestershire, 1994.
- [Pelavin and Allen, 1986] R. Pelavin and J.F. Allen. A Formal Logic of Plans in Temporally Rich Domains. In *Proc. of the IEEE*, 74(10):1364-1382, October 1986.
- [Penders, 1986] J.S.J.H. Penders. *Intentionaliteit*, Master thesis University of Amsterdam 1986 (in Dutch).
- [Penders, 1988] J.S.J.H. Penders. *On the application of AI-techniques in a robotworkcell* (in Dutch). Mem. 1766 DNL/88, KPN Research Leidschendam, 1988.
- [Penders, 1989] J.S.J.H. Penders. Errorrecovery in a robot system. In *Proc. IAS-2*, 929-936, Amsterdam, 1989.
- [Penders, 1991] J.S.J.H. Penders. Autonomous Vehicles Avoiding One Another: A case study of Obstacle avoidance Procedures. In *Proc. Computing Science in the Netherlands*, Stichting Mathematisch Centrum Amsterdam, 439-452, 1991.
- [Penders, 1994a] J.S.J.H. Penders. *Conflicts in the Spatial Interaction of Autonomous Agents*. PTT Research pre-publication, NT-PU-93-1728, 1994, (abstract appeared in Proceedings KI-94, Springer Verlag, 1994).
- [Penders, 1994b] J.S.J.H. Penders. *Global Reference Model*. IDEAS deliverables 1.1-1.3.2, R&D-IN-94-1203, KPN Research Leidschendam, 1994
- [Penders, 1995a] J.S.J.H. Penders. Spare parts delivery and Reusable Collection services, using Railway and Road transport. In *Proc. ISATA-95*, Stuttgart, 1995, Automotive Automation Ltd, Croydon, England.
- [Penders, 1995b] J.S.J.H. Penders. IDEAS-Informationsvernetzung in Distributionsketten. *Tagungsband zu den 13. Dortmunder Gesprächen*, Verlag Praxiswissen, Dortmund, 1995.
- [Penders, 1995c] J.S.J.H. Penders. *A note on the history of mail services*. R&D-PU-95-966, KPN Research Leidschendam, 1995.
- [Penders et al., 1994] J.S.J.H. Penders, L. Alboul (Netchitailova) and P.J. Braspenning. The Interaction of congenial autonomous Robots: Obstacle avoidance using Artificial Potential fields. In *Proceedings ECAI-94*, 694-698, 1994.

- [Penders and Braspenning, 1997] J.S.J.H. Penders and P.J. Braspenning. Situated Actions and Cognition. In *Proceeding IJCAI-97*, 1372-1377.
- [Perry, 1977] J. Perry. Frege on Demonstratives. *Philosophical Review*, 86:474-497, 1977.
- [Perry, 1979] J. Perry. The Essential Indexical. *Nous*, 13:3-21, 1979.
- [Pfeifer, 1995] R. Pfeifer. Cognition - Perspectives from autonomous agents. *Robotics and Autonomous Systems*, 47-70, 1995.
- [Pfeifer and Scheier 1997] R. Pfeifer and C. Scheier, Sensory-Motor Coordination: The Metaphor and Beyond. *Robotics and Autonomous Systems*, 157-178, 1997.
- [Pollack et al., 1987] M.E. Pollack, D.J. Israel and M.E. Bratman. *Towards an Architecture for Resource-Bounded Agents*. CSLI Report, No. CSLI-87-104, Stanford, 1987.
- [Ramsay and Barrett, 1987] A. Ramsay and R. Barrett. *AI in Practice: Examples in POP11*. John Wiley & Sons, New York, 1987.
- [Rao, 1996] Rao, A.S., Georgeff, M.P., Intelligent Real-Time network management, in: *Proc. ECAI Workshop on Intelligent Agents for Telecom Applications (IATA'96)*, Budapest, 1996.
- [Rauszer, 1992] C.M. Rauszer. Knowledge representation systems for groups of Agents. In *Proc. of Logic at work*, Amsterdam 1992.
- [Rimon, 1992] Elon Rimon. Exact Robot Navigation Using Artificial Potential Functions. *IEEE Trans. on Robotics and Automation*, 8(5), 1992.
- [Rosenschein and Kaelbling, 1986] S.J. Rosenschein and L.P. Kaelbling. The synthesis of Digital Machines with provable Epistemic Properties. In J. Halpern (ed.), In *Proc. of the Conf. on Theoretical aspects of Reasoning about Knowledge* (Morgan Kaufmann, 1986), pp. 83-98.
- [Rosenschein and Kaelbling, 1995] S.J. Rosenschein and L.P. Kaelbling. A situated view of representation and control, *Artificial Intelligence*, 73(1-2) 1995.
- [Rosenschein, 1985] S.J. Rosenschein. Formal Theories of knowledge in AI and Robotics. *New Generation Computing*, 3(4):345-357, 1985.
- [Rosenschein, 1989] S.J. Rosenschein. Synthesizing Information-Tracking Automata from Environmental Descriptions. In *Proc. of the 1st Int. Conf. on Principles of Knowledge Representation and Reasoning*, Morgan Kaufmann, 1989.
- [Rosenschein, J. and Zlotkin, 1994] J.S. Rosenschein and G. Zlotkin. Designing Conventions for Automated Negotiation. *AI magazine*, Fall 1994.
- [Ross and Schoman, 1977] D.T. Ross and K.E. Schoman. Structured Analysis for Requirements Definition. *IEEE Trans. on Software Engineering*, SE-3(1):6-15, 1977.
- [Ross, 1977] D.T. Ross. Structured Analysis (SA): A language for Communicating Ideas. *IEEE Trans. on Software Engineering*, SE-3(1):16-34, 1977.

- [Ryle, 1949] G. Ryle. *The Concept of Mind*, Barnes & Noble 1949.
- [Schreiber, 1996] *Proceedings of the 3rd IDEAS Workshop* 13-14 november, IML Dortmund 1996.
- [Searle, 1980] J. Searle. *Minds, Brains and Programs. The Behavioral and the Brain Sciences*, 3, 1980.
- [Seghrouchni and Haddad, 1996] A. El Fallah Seghrouchni and S. Haddad. A Coordination Algorithm for Multi-Agent Planning. *Workshop MAAMAW* 1996, Morgan and Kaufmann 1996.
- [Shannon, 1950] C.E. Shannon. *The Lattice Theory of Information*. Collected Works, 180-183, IEEE Press, New York, 1950.
- [Shoham, 1988] Y. Shoham. Chronological Ignorance: Eperiments in Nonmonotonic Temporal Reasoning. *Artificial Intelligence*, 36:279-331, 1988.
- [Singh, 1994] M.P. Singh. Multiagent Systems, *A theoretical Framework for Intentions, Know-How, and Communications*, Lecture Notes in Artificial Intelligence, Springer Verlag, 1994.
- [Slicher van Bath, 1960] B. Slicher van Bath. *De agrarische geschiedenis van West-Europa, 500-1850* (in Dutch). Spectrum, Utrecht/Antwerpen, 1960.
- [Sloman, 1994] A. Sloman. Explorations in Design Space. In Cohn, *Proc. ECAI 94*, 578-582, 1994.
- [Stalnaker, 1984] R. Stalnaker. *Inquiry*. Basil Blackwell, Oxford, 1984.
- [Steels, 1994] L. Steels. Equilibrium Analysis of Behavior Systems. In Cohn, *Proc. ECAI 94*, 714-718, 1994.
- [Stewart, 1995] J. Stewart. The implications for understanding high-level cognition of a grounding in elementary adaptive systems. *Robotics and Autonomous systems*, 107-116, 1995.
- [Svesta and Overmars, 1998] P. Svesta and M.H. Overmars, Coordinated path planning for multiple robots, *Robotics and Autonomous Systems* 23 125-152, 1998.
- [de Swart and Rauszer, 1994] H.C.M. de Swart and C.M. Rauszer. Different approaches to knowledge, common knowledge and Aumann's theorem. In *Proc. Workshop on Epistemic Logics*, Saarbruecken 1994, Springer Verlag, 1994.
- [Thorpe, 1917] G.C. Thorpe. *Pure Logistics: the science of War Preparation*. Published 1917. New edition: National Defence University Press, Washington DC, 1986.
- [Toth, 1995] J.A. Toth, Book review of *Reasoning Agents in a Dynamic World: The Frame problem* (K. M Ford and P.J. Hayes, eds.), *Artificial Intelligence*, 73(1-2) 1995.
- [Traiger, 1991] S. Traiger. Solipsism, Individualism and Cognitive Science. *J. Expt. Theor. Artif. Intell.*, 3:163-170, 1991.

- [Trullier and Meyer 1997] O. Trullier and J-A Meyer, Biomimetic Navigation Models and Strategies in Animats, *AI Communications* 10 pp. 79-92 1997.
- [Tychonievich, *et al.*, 1989] L. Tychonievich, D. Zaret, J. Mantegna, R. Evans, E. Muehle and S. Martin. A Maneuvering-Board Approach to Path Planning with Moving Obstacles. *IJCAI89*, 1017-1021, 1989.
- [Veltman, 1996] F. Veltman. Defaults in Update Semantics. *J. of Philosophical Logic*, 25(3), 1996.
- [Vermunt, 1996] A.J.M. Vermunt. *Wegen naar Logistieke grondvormen*. (in Dutch; inaugural lecture), Tilburg University Press, 1996.
- [Wagenaar, 1997] R.W. Wagenaar. *De virtuele koopman—fictie of werkelijkheid*. (in Dutch, Inaugural lecture), University of Amsterdam, March 1997.
- [Wang, 1989] P.K.C. Wang. Interaction Dynamics of Multiple Mobile Robots with Simple Navigation Strategies. *J. of Robotic Syst.*, 6(1):77-101, 1989.
- [Wang, 1991] P.K.C. Wang. Navigation Strategies for Multiple Autonomous Mobile Robots Moving in Formation. *J. of Robotic Syst.*, 8(2):177-195, 1991.
- [Webb, 1995] B. Webb. Using Robots to Model Animals: A Cricket Test. *Robotics and Autonomous Systems*, 16(2-4), 1995.
- [Webb, 1996] B. Webb. A Cricket Robot. *Scientific American*, December 1996.
- [Webster's, 1990] The new lexicon Webster's dictionary of the English language 1990.
- [Wernersson *et al.*, 1989] A. Wernersson, U. Wiklund, U. Andersson and K. Hyypä. Vehicle Navigation using Image Information. In *Proc. Intelligent Autonomous Systems*, 814-822, Amsterdam, 1989.
- [Wolters, 1995] W. Wolters. Kooperation contra Chaos. *Tagungsband zu den 13. Dortmunder Gesprächen*, Verlag Praxiswissen, Dortmund, 1995.
- [Wooldridge and Jennings, 1995] Agent Theories, Architectures, and Languages: a Survey, in Wooldridge and Jennings (eds), *Intelligent Agents: ECAI-94 workshop on Agent Theories, Architectures, and Languages*, Springer Verlag, 1-39, Berlin 1995.
- [Zeghal and Ferber, 1993] Karim Zeghal and Jacques Ferber. Craash: a Coordinated Avoidance System, Modelling and Simulation 1993. In *Proc. of the 1993 European Simulation Multiconference*, ed. A. Pave, Lyon, 1993.
- [Zelinsky, 1992] A. Zelinsky. A Mobile Robot Exploration Algorithm. *IEEE Trans. on Robotics and Automation*, 8(6), 1992.
- [Zhao *et al.*, 1992] Zhao Jingsong, Wang Fengchen and Wu Zhao-lin. The development of Ship Collision Avoidance Automation. *The Journal of Navigation*, 45(1), January 1992.

Other references

- [Bhagavad Gita, 500 B.C.] Translated by J. Mascaro, Penguin Books London 1962.
- [ten Brink, 1969] E.A.B.J. ten Brink. *Geschiedenis van het postvervoer*. (in Dutch) Fibula-van Dishoek, Bussum, 1969.
- [Clavell, 1981] J. Clavell. *Foreword to The Art of War by Sun Tzu*, Hodder and Stoughton, London 1981.
- [Dancey, 1994] C. Dancy, *Encyclopaedia of Ball Juggling*, Butterfingers, Bath UK 1994.
- [Hildebrand, 1839] Hildebrand. *Camera Obscura* (in Dutch).
- [Hirose, 1992] N. Hirose (translator). *Immovable Wisdom: The Art of Zen Strategy: The Teachings of Takuan Soho*. Element books Ltd., Shaftesbury, Dorset, 1992.
- [Hogesteger, 1989] G. Hogesteger. *Van lopende bode tot telematica* (in Dutch). Koninklijke PTT Nederland NV, 1989.
- [Krans, 1963] G.H.A. Krans. *Boden en wegen van middeleeuwse kooplieden, Vereeniging tot beoefening van Overijssels Regt en Gechiedenis*. (in Dutch) Deventer, 1963.
- [van der Laan, 1994] G.M. van der Laan. Vijftig jaar Technische Troepen, hoe nu verder? in Symposium 30 mei 1994, uitg. RTT-50 jaar, KMA Breda 1994.
- [Lauffer 1954] O. Lauffer. Der Laufende Bote im Nachrichtenwesen der früheren Jahrhunderte, *Beitrage zur Deutschen Volks- und Altertumskunde* 1954 (in German).
- [Lin Chi, †866] *The Teachings of Master Lin-Chi*, translated by B. Watson [1993].
- [Sun Tzu, 490 B.C.] *The art of war*; translated by P.F. Calthrop in [Clavell, 1981].
- [Suzuki, 1959] D.T. Suzuki. *Zen and Japanese Culture*. Bollingen/Priceton, New York 1959.
- [Taisen Deshimaru 1982] Taisen Deshimaru. *The Zen Way to the Martial Arts*, Arkana, Penguin Books 1982.
- [Takuan ±1600] *Immovable Wisdom: The Art of Zen Strategy: The Teachings of Takuan Soho Teachings* translated by Hirose [1992].
- [Watson, 1993] B. Watson (translator). *The Zen Teachings of Master Lin-Chi*. Shambala, Boston and London, 1993.

Curriculum Vitae

Jacques Penders was born in the Noord Oost Polder on the 20th of June 1955. He started his secondary school education in Stevensbeek and finished it at the Christelijke Scholen Gemeenschap in Emmeloord. After having worked as a woodcutter and gardener for a year, he resumed his studies. He was educated as a Garden and Landscaping architect, at the Rijks Hogere School voor Tuin- en Landschapsinrichting in Boskoop. He continued his studies at the University of Amsterdam. Starting with Philosophy, with a specialisation in Analytical Philosophy and Philosophical Logic, he extended his studies to Mathematics with a specialisation in Mathematical Logic as well.

During his studies Jacques had a job as a teacher of Mathematics. After finishing his studies in Philosophy in 1986, he received a job as Toegevoegd docent (Supply teacher) in Logic at the University of Amsterdam. In 1987 he joined the robotics group at KPN Research. His research concerned the application of artificial intelligence techniques in robotics, culminating in the design of a portable autonomous vehicle which his project team built. Meanwhile, in 1988 he also finished his studies in Mathematics. In the beginning of the nineties, his work shifted to logistics. With pleasure he remembers defining and managing the IDEAS project, a three year European project which focused on logistics. Two years ago, Jacques returned to the field of artificial intelligence, and currently does research on intelligent agents and their interaction.

Jacques' hobbies include juggling and iaido. Around 1984, he taught himself juggle with three balls. At that time his son was about two years old and the son's amusement about the dropping balls provided the initial stimulus to practice. After an interval of about ten years he took up this old hobby, and is now able to juggle five balls. Iaido, the other hobby, is a Japanese martial art, aiming at fluidly drawing a samurai sword. Jacques holds a second dan (degree) in this art.

SIKS Dissertatiereeks

In 1999 zijn de volgende SIKS-dissertaties verschenen.

- 99-1 Mark Sloof (VU)
 *Physiology of Quality Change Modelling; Automated Model-
 ling of Quality Change of Agricultural Products*
Promotor: Prof.dr. J. Treur
Co-promotor: Dr.ir. M. Willems
Promotie: 11 mei 1999
- 99-2 Rob Potharst (EUR)
 Classification Using Decision Trees and Neural Nets
Promotor: Prof.dr. A. de Bruin
Co-promotor: Dr. C. Bioch
Promotie: 4 juni 1999
- 99-3 Don Beal (Queen Mary and Westfield College)
 The Nature of Minimax Search
Promotor: Prof.dr. H.J. van den Herik
Promotie: 11 juni 1999
- 99-4 Jacques Penders (KPN Research)
 The Practical Art of Moving Physical Objects
Promotor: Prof.dr. H.J. van den Herik
Co-promotor: Dr. P.J. Braspenning
Promotie: 11 juni 1999

ISBN 90-72125-64-9