Spatial Conflicts among Robot Agents

J.S.J.H. Penders

KPN Research
P.O. Box 421
2260 AK Leidschendam, The Netherlands
J.S.J.H.Penders@research.kpn.com

Abstract

In this paper we present a model to study conflicts among robot agents. Within this model we can define intentions. In the second part of the paper we discuss examples of teams behaviours, which are based on the intentions the individual agents have.

Introduction

This workshop is entitled Agents' Conflicts, the first question to be answered is what is a conflict? Our provisional answer is that two or more agents have a conflict when they aim for the same resource and thus hinder each other in using this resource. We present a model for analyzing such conflicts among robot agents.

The provisional answer contains two important clues: 'resource' and 'hinder each other'. The resource for which they are aiming is the space on the shop floor. A conflict thus can be described as the *event* of two (or more) agents aiming to occupy the same place. Below, we discuss examples of such events. It is more difficult to stipulate what it means that robot agents hinder each other. The robots perform spatial movements. Moreover, robots like software agents are autonomous agents, they are pro-active, react and interact. Thus the environment is not static, actions and reactions follow on each other. This complicates to stipulate what it means that an agent is hindered. We proceed by defining the notion *of having an intention*. An intention denotes a pattern of behaviour that shows some aim, for instance finding a goal or avoiding an obstacle. In a dynamic environment a reactive agent often cannot head straight for its goal or avoid an obstacle. Whenever we can recognize a certain characteristic in the spatial (re)actions of the agent, we say that the agent has an intention. Thus, an intention denotes the type of (re)action an agent performs. 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]. Moreover, we observe that 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.

Before starting the analysis, we need to point out some abstractions which we are making. Our robots have a limited cognitive basis [Penders, 1997], and the ontology of our robot domain is limited to two dimensional space. 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. Any of the objects may perform spatial behaviour, and in general we speak about *robot agents* including any kind of object. We treat the robot agents as so-called *free flying objects* [Latombe, 1991]. This means that the robots can move freely in the two-dimensional plane, when moving they can change speeds immediately. Moreover, the robots are considered to be *point robots* [Fujimara, 1991]: they have no spatial extension and compete in the plane about points only.

In the next section we present our interaction model. It applies to a team of robots each of which performs goal finding and obstacle avoidance. This section is quite technical and might be skipped at first reading. After that we have a less formal discussion in which we compare several different teams with respect to the team behaviour generated when the team members interact.

Interaction Model

In this section we distinguish and define interaction patterns. We start by singling out certain special situations or events which may lead to a *collision*. A team of avoiding robots is quite dynamic. Thus, during the interaction certain *incidents* may occur which force the agents to change their courses in order to prevent collisions. This is the point where interaction patterns are generated. 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 a formal definition of *having an intention.*

Time and Notations

In order to deal with the temporal aspect of interaction we introduce *points in time*. These points in time 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, a chosen action is sustained till the next point in time. Interaction patterns are extended in time and cover series of time points. However, the course of events has only one (future) direction, hence a *linear time* model is appropriate.

$\boldsymbol{D}$ is the overall domain of objects and includes all agents and robots. Each agent i is able to observe any of the other agents $j \in \boldsymbol{D}$. Thus we can speak of a subset $d \subseteq \boldsymbol{D}$ 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. 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:

$\boldsymbol{D}$	overall domain of objects and robots
d	$d \subseteq \boldsymbol{D}$ set of relevant objects
i,j,o	variables to denote objects or agents
$(\mathbb{R}\times\mathbb{R})$	the two-dimensional plane
V_{ixy}	robot or vehicle i at position (x,y), in general when no confusion is expected we leave out (x,y); hence, V_i: position of robot 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
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 1. $t \models Ob(i,o)$ **iff** $t \models \exists o\ V_o \in H_i$ 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. Horizons might differ between the agents, however for convenience of the present discussion, we assume that all agents have the same horizon. 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: $Ob(i,j) \leftrightarrow Ob(j,i)$.

Events

Definition 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 1 shows a real-collision of two robots A and B.

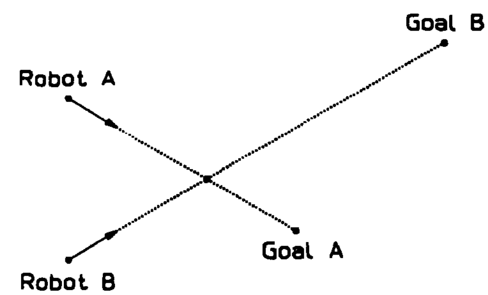

Figure 1: A real-collision situation of two robots.

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

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 2, robot A nearly passes p while robot B is due to arrive much later.

Definition 4. *Semi-collision (SC(d)).*

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

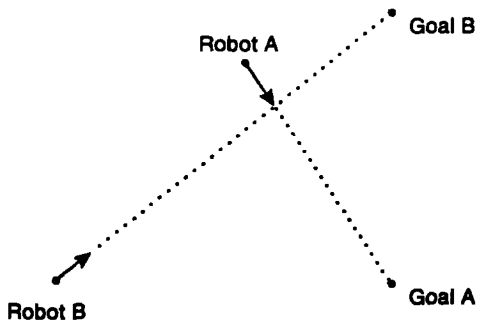

Figure 2: A semi-collision situation.

We have identified real-collisions and semi-collisions, but reactive agents start to avoid each other 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 and thus starts to deviate from its regular course.

Definition 5. *Incident (I(i,d)).*

t¦ $I(i,d)$ **iff** t¦ $_{j \in d}$ $(Ob(i,j))$.

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 I(d) all agents observe each other.

Definition 6. *Mutual incidents (I(d)).*

t¦ $I(d)$ **iff** t¦ $_{i \in d}$ $I(i,d')$ where $d'=d \setminus i$: d' is d but i is excluded.

Above in Definition 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 incidents 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 [Penders, 1991] for more examples.

Interaction Patterns

Having identified several particular events, we now 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 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 t:

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

Definition 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 8, related situations have so-called constant domains: in two related situations t and t' the same agents occur.

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 4) in related situations (Definition 8). *Conflicts* are interaction patterns where the agents mutually interfere with one another's aim to reach the (respective) destination point(s).

Definition 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.

T¦ *Conflict*(d) **iff T**¦ $I(d)$ $[t]$ 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 1). In such cases the incidents are no longer related (Definition 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 10. t¦ *Ending of Conflict*(d) **iff** t¦ $\exists\, i \in d$ $(\mathrm{dist}(V_i,G_i)=0)$.

Strictly speaking Definition 10 demands that an agent has to come to a standstill exactly at the position of the goal. In general some tolerance need to be admitted, however we will not introduce these options here, as it does not affect the general idea we want to sketch. In general, a robot involved in a conflict does not approach its destination point straightaway. As we have seen in Figure 3, 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 3 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. We capture this idea in the notion of a *targeting* series with the goal as its centre point.

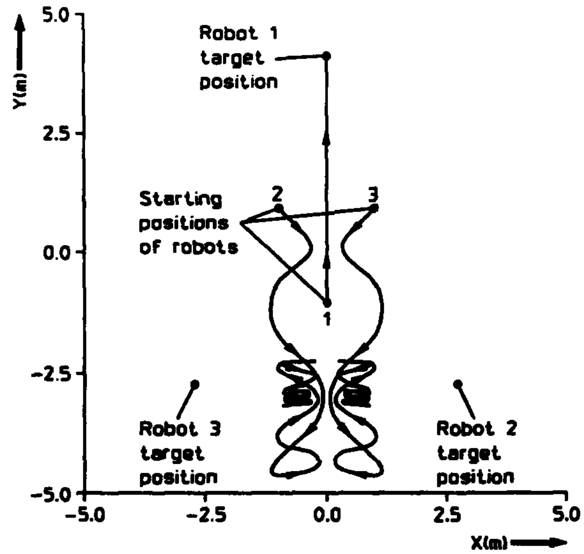

Figure 3: Trajectories of three congenial robots.

Definition 11. A time frame $\mathbf{T}=(T,<)$ contains a *Targeting series* $TS(i,G_i)$ for robot i: if

$\mathbf{T} \models \text{dist}(V_i,G_i)=c\ [t^*]$ for some $t^* \in T$

and $\forall c', c>c'>0$ it contains a time point t later than t^* such that

$\mathbf{T} \models (\text{dist}(V_i,G_i)<c')\ [t]$

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

Theorem 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 $\text{dist}(V_i,G_i) = 0$. □

Many conflicts end in the long run, though at first sight it is not clear. A targeting series, according to Theorem 1, provides a characteristic to recognise ending conflicts. However, there are conflicts that do not have an end, as already was illustrated by Figure 3. We call them balanced conflicts.

Definition 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 2. A time frame $\mathbf{T} = (T,<)$ contains a balanced *Conflict(d)* **iff**

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

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

Proof. According to Theorem 1, a balanced conflict has no end iff it contains no targeting series, and that is what Lemma 2 says. □

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 of *having an intention*. Intentions concern the behaviour of a single agent, no matter how others react. An intention is applied in a series of related situations (Definition 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 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 13a applies for goal finding, but we can also treat 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.

Defintion 13b. Agent i *avoids* object o (at position (x,y)), if there exists a diverging series T for i with o as its point of departure:

$\mathbf{T} \models \text{dist}(i,o)=c\ [t^*]$ for some $t^* \in T$

and $\forall c', c<c'$ it contains time points later than t^* such that

$\mathbf{T} \models (\text{dist}(i,o)>c')\ [t]$ for some $t \in T$, $t^*<t$.

Targeting and diverging series also can be applied when the target object o moves. Thus Definition 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 13 is an indication of what intentions are in a geometrical context.

Intentions of Agents

Above, we have defined intentions, thus we can also differentiate agents by their intentions. The basic distinction is between agents which react and those which do not. Hence in Tabel 1 we have two main groups, *reactive* agents which react and *non-reactive* agents which do not react.

Table 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 move[1]. 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 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 are also intent upon the object and aim at a *moving* target, they need another agent as their target. In other words, antagonists are chasing other agents. Antagonistic agents might differ in the sense 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.

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

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] who 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 cap abilities 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 [Penders, 1997]). 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 apply the same intentions, in terms of design they are robots provided with the same programs.

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 intentions. An overview is given in Table 2.

Table 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.

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

Table 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

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 discussed above. The notion of an autonomous robot is commonly associated with that of an avoiding agent. Our interaction model presented in the previous section applies in particular to an avoiding team.

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 Figure 4 below.

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 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.

Newcomer in a Team

The congenial teams studied above are homogeneous and this 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 3 the rows give the congenial teams of Table 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.

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 (colomns 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[1]. 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.

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.

We close the discussion with an example in Figure 4. The robots are selective non-aggressive antagonists and apply an avoidance procedure. We have seen 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. When the goal starts to move, the team follows in formation, that is without colliding with each other or the goal. Thus, 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.

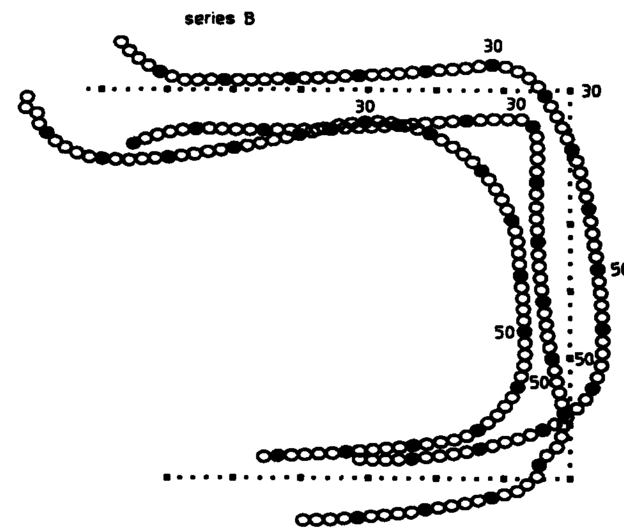

Figure 4: Robots moving in formation.

So far for the analysis of adding a newcomer to a congenial team. For non-reactive teams we conclude that on the one hand 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.

To conclude

In Section 2 we have set up an interaction model. Robots were treated in a rather abstract way. 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. In [Penders *et al*, 1994] and [Penders, 1999] mathematical analyses of team behaviour are given which are based on the presented interaction model.

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].

In section 3, we considered teams of robot agents and discussed the type of team behaviour that results. Natural

[1] 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.

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]. We believe that some clues can be found along an analysis as we have started here.

References

[Fujimara, 1991] K. Fujimara. *Motion Planning in Dynamic Environments.* Lecture Notes in Artificial Intelligence, Springer Verlag, 1991.

[Georgeff and Ingrand, 1989] M.P. Georgeff and F.F. Ingrand, Decision-Making in an Embedded Reasoning System, *Proceeding IJCAI 1989* pp. 972-978.

[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.

[Latombe, 1991] J.C. Latombe. *Robot Motion Planning.* Kluwer, Boston, 1991.

[Mataric, 1995] M.J. Mataric. Issues and approaches in the design of collective autonomous agents. *Robotics and Autonomous Systems*, 16, 1995.

[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 *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 1999] J.S.J.H. Penders, *The Practical Art of Moving Physical Objects.* PhD thesis, University of Maastricht, 1999.

[Penders and Braspenning, 1997] J.S.J.H. Penders and P.J. Braspenning. Situated Actions and Cognition. In Proceeding IJCAI-97, 1372-1377.

[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.

[Trullier and Meyer, 1997] O. Trullier and J-A Meyer, Biomimetic Navigation Models and Strategies in Animats, *AI Communications* 10 pp. 79-92 1997.

[Webb, 1995] B. Webb. Using Robots to Model Animals: A Crichet Test. *Robotics and Autonomoues Systems*, 16(2-4), 1995.