Multi-robot collaboration and coordination in a high-risk transportation scenario

Daniela Doroftei and Eric Colon

Royal Military Academy, Department of Mechanics,
Av. De La Renaissance 30, B1000 Brussels, Belgium
{daniela.doroftei,eric.colon}@rma.ac.be
http://mecatron.rma.ac.be

Abstract. This paper discusses a decentralized multi-robot coordination strategy which aims to control and guide a team of robotic agents safely through a hostile area. The "hostility" of the environment is due to the presence of enemy forces, seeking to intercept the robotic team. In order to avoid detection and ensure global team safety, the robotic agents must carefully plan their trajectory towards a list of goal locations, while holding a defensive formation.

Key words: Multi-robot coordination, decentralized control, behavior-based control, formation control

1 Introduction

1.1 Problem Statement

Autonomous robotic agents can provide a valuable asset in the military context, where potentially dangerous missions have to be carried out. One example of such a mission is the transportation of goods from one place to another, passing through an area where enemy forces are active. An autonomous multi-agent team of robots could potentially carry out this mission and as such keep humans out of the danger zone. The transportation scenario considered here imposes that a team of robots protects a given unit, during a transport mission where (multiple) goal locations have to be attained. To defend the protected robot, the multi agent team must form a defensive formation. In this context, a circular formation was chosen to optimize the defensive capabilities of the multi-agent team. Moreover, the robotic team must aim to stay out of sight of enemy forces, avoiding contact with these enemies in order not to endanger the protected robot. Furthermore, the robots must ensure their own safety by not bumping into obstacles or tipping over due to the high slopes which are present on rough outdoor terrain.

1.2 Proposed Approach

The presented approach casts the multi-robot control problem as a behavior-based control problem. In the behavior-based spirit a complex control problem is divided into a set of simpler control problems that collectively solve the original complex control problem [1]. Section 3 of this paper describes in detail how each behavior was designed and how the behavior fusion problem was solved. The behavior-based control paradigm was chosen, because it is inherently decentralized and because it thus provides a natural and elegant way to combine the different subtasks and capabilities of each individual robot and because - unlike more traditional sense-model-plan-act approaches - it scales very well when applied to a large number of robots.

An important aspect of the presented control architecture is that it is formulated in a decentralized context. This means that the individual robots have no knowledge of any global state parameters. As a result of this, the individual robots do not have a global map; hence they cannot rely on traditional global path planning algorithms for navigation. Instead, path planning is achieved through a behavior-based control paradigm, where multiple behaviors interact together. Each behavior considers one specific navigational task (e.g. avoiding obstacles, reaching a goal, ...). Fusing all behaviors together leads to a complex global behavior, designating a path to be followed by the individual robots.

1.3 State of the Art

Control and coordination of multiple robots is a research field which has received a lot of attention during the past 2 decades, generally using some kind of exploration scenario, e.g. in the context of a mapping [2], search and rescue [3] or surveillance [4] application. An important research area is the optimization of the multi-robot formation. In nature, formation control benefits the animals that make use of it [5]. By grouping themselves, animals combine their sensors to maximize the chance of detecting predators, as such minimizing the encounters with predators [6]. Inspired by these biological examples, roboticists have applied similar formation control approaches to artificial agents [7]. In robotic systems, the control task is in general decomposed into a set of behaviors, following the paradigm for behavior-based robotics [8][9]. Also in this paper, we use a behavior-based control strategy, as described in section 3.

When reviewing the state of the art on multi-robot control, one recurring constatation is that most experiments take place in an indoor scenario. The indoor setting facilitates the control problem, as certain issues such as the roughness of the terrain do not need to be taken into account [10]. In [11], Madhavan et al. present an approach towards distributed multi-robot control for localization and mapping in an outdoor environment, taking into account uneven outdoor terrain. In this paper, we present a methodology which takes into account the roughness of the outdoor terrain, but we use it in a different context of formation control. Chen and Luh proposed in [12] a distributed approach towards formation control, showing large groups of robots moving cooperatively in various geometric formations, as we also present in this paper. However they do not take into account the problem of obstacle (or enemy) avoidance, as shown in this paper. The research works closest related to the presented approach are the ones by Balch [5] and Parker [13]. Parker simulates robots in a line-formation, navigating to waypoints while avoiding obstacles. The research performed by Balch is similar to Parker's to the extent that it includes an approach for robotic line formation maintenance [5] and extending this to three additional formation geometries. In this paper, we introduce a new formation type, more specifically the circular formation. From an architectural point of view, Parker uses the layered subsumption architecture [14], which selects behaviors competitively, meaning that only one can be active at any given time. Balch uses a motor schema approach, which enables multiple behaviors to be active concurrently. The presented architecture in this paper also permits multiple behaviors to be active simultaneously, but uses an objective function approach [1] towards the behavior definition and fusion, instead of the motor schemas used by Balch.

2 Global Strategy

Following the transportation scenario, the main goal for the robot team is not to meet any enemy forces. This scenario considers a team of robots transporting goods in a hostile environment. Enemies may be present and must be avoided. The position and movement of these enemy forces is a priori unknown and they can only be detected when they are within the sensor range of one of the team members. In the simulation, the enemy positions are initialized randomly and the enemies also follow a randomized movement strategy. While advancing to goals and avoiding enemies, the robots must maintain as well as possible an inter-robot formation. This robot formation should ensure:

- Maximum ego-visibility: The combined sensor range of all team members must be as large as possible, such that enemy forces can be detected from far away.
- Minimum exo-visibility: The visibility of the robotic team from the outside world should be as small as possible to prevent the detection of the robotic team by enemy forces

In the context of this simulation, the circular-shaped formation was chosen. A circular formation has many advantages as a defensive formation, as it ensures the all-round security of the troops. Indeed, using this type of formation, the flanks of all the individual robots involved are protected as this formation prevents the enemy from attacking any robot's flank as it is always guarded by a fellow robot. In light of these advantages, the circular formation was chosen, but the architecture is completely modular and it is trivial to implement any other formation. In practice, a perfect circular shape is impossible to attain with a discrete, limited number of robots. Therefore, the closest discrete alternative to a circular shape, a regular n-polygon, is chosen as an optimal formation, with n being the number of robots. On top of the constraints listed above, the robots must also take into consideration the traversability of the terrain. In this scenario, a large rural environment is considered, as shown on Figure 1. The black areas on this top-view map represent obstacles, whereas the gray areas indicate the height of the terrain. As can be noted from Figure 1, there are large denivelations on this terrain, as it considers a hilly area. This means that the robots need to take into account the grade of the slopes to account for the traversability of the terrain. An important aspect of the field transport scenario is the presence of

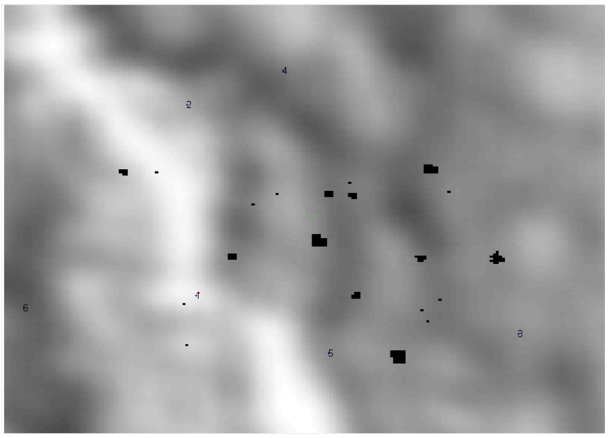

Fig. 1. Top View 2D Map of the large rural environment used for the simulation of the transportation scenario. Black areas are obstacles; gray values indicate terrain height. Six robots, numbered and represented by blue dots, are present in the environment. The goal point is represented by the green dot in the center of the map. One intruder is present, represented by a red dot, close to the first robot.

a special robot, which is placed in the center of the multi-robot team and which transports some valuable goods or persons. The other robots are protecting this robot by forming a defensive circle. The presented behavioral multi-robot coordination framework uses the same set of basic behaviors for the central and for the protecting robots, and also uses the same fusion approach for both cases. The only thing which changes is the allocation of the different behaviors to the different team members. Here, we describe the different actions behaviors acting on the different team members:

- Central robot:
  - *GoToGoal*: Directs the robot into the direction of the current goal
  - *AvoidIntruders*: Directs the robot away from intruders
  - *AvoidSlopes*: Directs the robot away from large slopes
- Protecting robots:
  - *GoToGoal*: based upon the position of the central robot, the optimal position of the current protecting robot (x_o, y_o) is calculated. This position can be estimated for each robot i as $(x_c, y_c) = (x_c + R\cos(2i\pi/N), y_c + R\sin(2i\pi/N))$, with (x_c, y_c) the position of the central robot, R the optimal formation radius and N the total number of robots. The *GoToGoal* behavior is then used to steer the robot to this optimal position
  - *AvoidIntruders*: Directs the robot away from intruders
  - *AvoidSlopes*: Directs the robot away from large slopes

As can be noted, there is no explicit formation control behavior. In fact, the *GoToGoal* behavior of the protecting robots acts as a formation control behavior, as it directs the robots into the direction of the optimal formation position. Using this strategy, the central agent seeks to attain the designated goal positions, while the other team members seek to attain the optimal defensive position around the central agent. The combination of all these behaviors enables the robot to execute the global task of the scenario, as described above. In the following, we will describe the design of each of these different behaviors more in detail. The evaluation of each behavior is based upon the evaluation of a discretized number of possible locations for a robot to move to. For each possible robot orientation, one measures the degree to which moving into this orientation would satisfy the subtask specified by the behavior formulation. This degree of goal attainment forms the output of the behavior, which can thus be seen as a one-dimensional function of the robot orientation. The approach towards behavior fusion applied in the context of this scenario is the traditional weighting method. Following this methodology, the output of each behavior is accorded a certain weight and the final output consists of a weighted average of the different individual behaviors.

3 Behavior Design & Fusion

3.1 GoToGoal

As indicated above, the design of each behavior goes out from the evaluation of a number of discretized new orientations θ where the robot could go. At each of these orientations, the amount in which the goal of the behavior is attained is measured. For the behavior making the robot advance to the goal position, the task consists of minimizing the distance to the goal point. This is achieved by calculating the direction to the goal point and by comparing this angle to each possible orientation θ. The optimal robot orientation is the one which minimizes the difference between the robot orientation θ and the angle to the goal θ_{Goal}, such that we can write an objective function for the behavior as $GoToGoal(\theta) = \frac{1}{AngleDifference(\theta,\theta_{Goal})}$. As an example of using this formulation, consider the case as sketched by Figure 1, which shows a 2D view of the environmental map, where 6 robots are present. These robots are numbered and represented by blue dots. The goal point is represented by the green dot in the center of the map. Applying the objective function formulation for the *GoToGoal* behavior, the case sketched by Figure 1 leads to an orientation diagram for each robot as shown by Figure 2. The orientation diagram of Figure

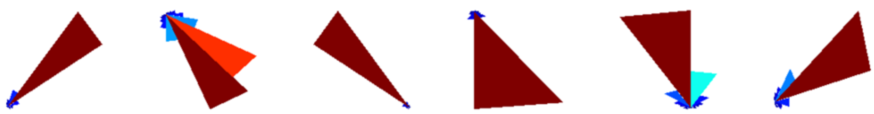

Fig. 2. Orientation Diagrams for the *GoToGoal* behavior for Robots 1 to 6.

2 indicates for each robot the direction in which it should move in order to reach the goal position. As can be noted, robots 1 and 6, which can be located in the lower left quadrant, should move in the north-east direction according to the orientation diagram. This is also the direction towards the goal point. On the other hand, robots 2 and 4, which are located in the upper left quadrant, have preferential orientations towards the south-east. To conclude, it is clear that the orientation diagrams show preferential orientations for movements which will bring the robot closer to the goal position, as required for this behavior.

3.2 AvoidIntruders

The behavior for steering the robot away from enemy forces who may want to intercept the robotic team is very similar to the behavior for guiding the robot to a goal position. Indeed, whereas the goal position provides an attractive force for the robots, enemies provide a repulsive force, but in the sense of calculations, this is quite similar. Like in the case of goal seeking, the orientation $\theta_{Intruder}$ towards the (nearest) intruder is calculated and an objective function is defined by comparing the different robot orientations to this orientation of the intruder: $AvoidIntruders(\theta) = AngleDifference(\theta, \theta_{Intruder})$. For the case of the 6 robots depicted by Figure 1, this leads to orientation diagrams as shown on Figure 3. In this situation, robots 1,5 and 6 have detected the intruder and

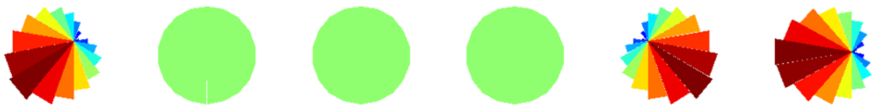

Fig. 3. Orientation Diagrams for the *AvoidIntruders* behavior for Robots 1 to 6.

are fleeing away in the opposite direction, according to their orientation diagram for this behavior. Robots 2,3 and 4 are too far from the intruder to detect the intruder and as they are unknowing about the intruder, they cannot take any action. Therefore, the orientation diagrams for these robots are uniform.

3.3 AvoidSlopes

Outdoor robots navigating on accidented terrain need to take into account the 3D properties of the terrain. Amongst others, they must avoid slopes which are too steep. In fact, any slopes should be avoided if there is an easy way round, because robots generally advance much more efficient on flat terrain. Therefore, the *AvoidSlopes* behavior takes into account the terrain elevation data, present in the height map h, as shown on Figure 1. This map is in fact a digital elevation model of the environment and the information which it holds can be directly used to define an objective function: $AvoidSlopes(\theta) = \frac{1}{abs(h(\mathbf{x})-h(\mathbf{x}+\mathbf{m}(\theta)))}$, where $\mathbf{m}(\theta)$ is a function projecting the robot to a new position as a function of a given orientation θ. Applying this objective function on the 6 robots depicted by Figure 1, leads to the orientation diagrams as shown on Figure 4.

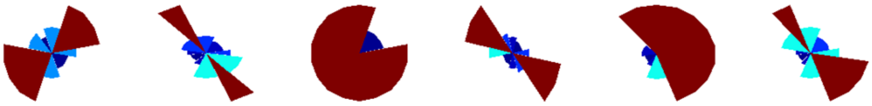

Fig. 4. Orientation Diagrams for the *AvoidSlopes* behavior for Robots 1 to 6.

3.4 Behavior Fusion

As already mentioned in the description of the global strategy, the traditional weighting method is employed here to combine the output of all behaviors. Formally, this can be described as: $\phi(\theta) = \frac{\sum_{i=1}^{4} w_i behavior_i}{4}$. The different weights w_i express the importance which is attached to each of these behaviors. Figure 5 shows for each of the 6 robots the fused orientation diagram.

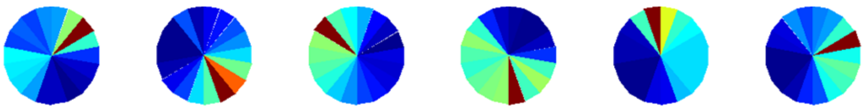

Fig. 5. Behavior Fusion using Orientation Diagrams for Robots 1 to 6.

4 Results

4.1 Quantitative Analysis

Evaluation Methodology To evaluate the presented coordination strategy, we have split up the analysis according to the different capabilities which are active in the presented behavior based multi-robot coordination architecture. First, we evaluate the performance of the Go To Goal behavior alone (GTG), then we analyze the impact of the addition of the formation control (GTG + FC) and after that we study the effect of enemy avoidance behavior (GTG + FC + IA). Finally, the full control paradigm, including the slope avoidance is analyzed (GTG + FC + IA + SA). Next to these behaviors, it must be noted that the obstacle avoidance behavior is built into the robot motion planning algorithm; as such we do not deal with it at this level. In practice, some care must be taken when validating the multi-robot coordination strategy described above, as the enemy forces can be located at any random location on the map and can wander around freely. The validation of the different metrics can therefore only lead to a valuable conclusion if each experiment e is repeated a number of times E (here, $E = 200$). This means that each data point on the following graphs is in fact a mean of data retrieved over a total 200 experiments, and in all of those experiments the initial position of the robots and the enemy forces was chosen totally random. In this series of experiments, a multi agent team consisting of 17 robots was placed in a rough country-side environment, where 2 enemy troops are wandering around.

Goal Reaching Figure 6 shows the evolution of the mean distance to the goal over a number of iterations. The mean distance to the goal is defined as:

$$\Delta(k) = \frac{1}{E}\sum_{e=1}^{E}\frac{\sum_{i=1}^{n}\sqrt{\left(x_i(k) - x_{Goal}\right)^2 + \left(y_i(k) - y_{Goal}\right)^2}}{n} \tag{1}$$

with n the number of robots (central robot + protecting robots) and k the iteration number. The asymptotic convergence of all curves shown in Figure 6 shows

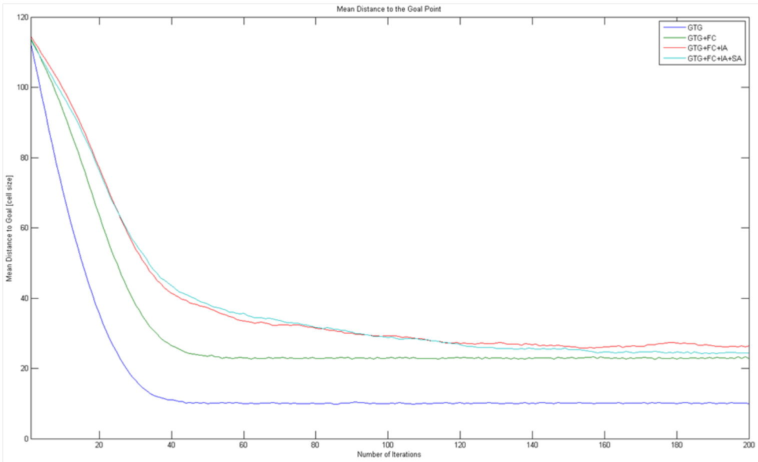

Fig. 6. Goal reaching capability.

that under all circumstances the destination point is reached and that, unsurprisingly, the fastest methodology of attaining the goal point is executing only the *GoToGoal* behavior (blue curve). This is normal, as adding extra constraints / objectives to the coordination strategy will necessarily lower the priority of the goal reaching behavior. Another remark is that the final distance from the goal is lower in the case of using only the GoToGoal behavior in comparison to all the other ones. The reason for this lies in the fact that the mean distance to goal metric, as defined above, considers all robots for calculating the distance to the goal, not only the central robot.

Formation Control Figure 7 shows the error on the formation. This metric is defined, based upon the distance between the actual protecting robot position and the ideal protecting robot position with respect to the central robot. Let (x_c, y_c) be the position of the central robot. The ideal position of any protecting robot i can then be calculated as $(x_i^*, y_i^*) = (x_c + \cos(2\pi i/n), yc + \sin(2\pi i/n))$. The formation error can then be calculated simply by calculating the distance between the actual robot position and the ideal robot position:

$$FormationError(k) = \frac{1}{E}\sum_{e=1}^{E}\frac{\sum_{i=1}^{n}\sqrt{\left(x_i\left(k\right)-x_i^*\left(k\right)\right)^2+\left(y_i\left(k\right)-y_i^*\left(k\right)\right)^2}}{n} \tag{2}$$

From Figure 7 it is evident that once the formation control is activated, the error

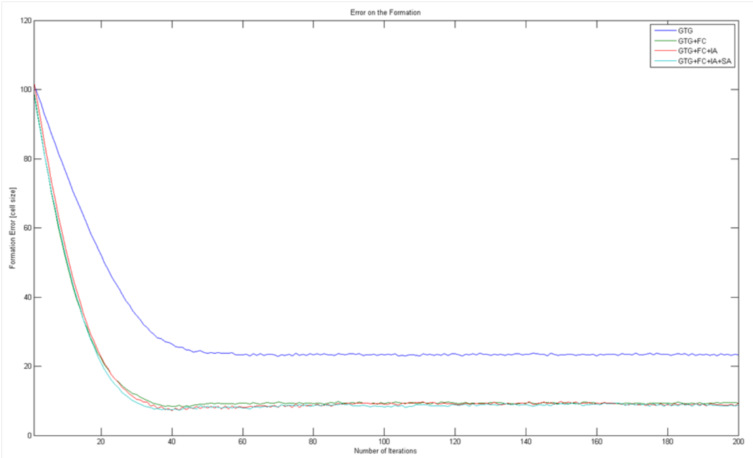

Fig. 7. Formation control capability.

on the formation is reduced drastically, indicating that the formation control behavior works very well.

Intruder Avoidance Figure 8 depicts the intruder avoidance capability of the multi-agent robotic team, by plotting the "visibility" of the robotic team by enemy forces. The visibility metric is defined by considering for each member of the enemy its field of view dov_{enemy}. Friendly robotic team members which are situated within the field of view of an enemy troop are attributed a penalty score inversely proportional to the distance to the enemy troops. Let $d_{enemy,i}$ be the distance between an enemy robot and a friendly robot i. The exo-visibility V can then be defined as:

$$V_{enemy,i} = \begin{cases} 0 & if \quad d > dov_{enemy} \\ dov_{enemy} - d_{enemy,i} & if \quad d \leq dov_{enemy} \end{cases} \tag{3}$$

This so-called exo-visibility measure per robot is divided by the number of robots to estimate a mean visibility value and this value is also averaged over a number of experiments E.

$$Visibility(k) = \frac{1}{E} \sum_{e=1}^{E} \frac{\sum_{enemy=1}^{NumberOfEnemies} \sum_{i=1}^{n} V_{enemy,i}(k)}{n} \quad (4)$$

Figure 8 clearly indicates how the inclusion of the intruder avoidance behavior

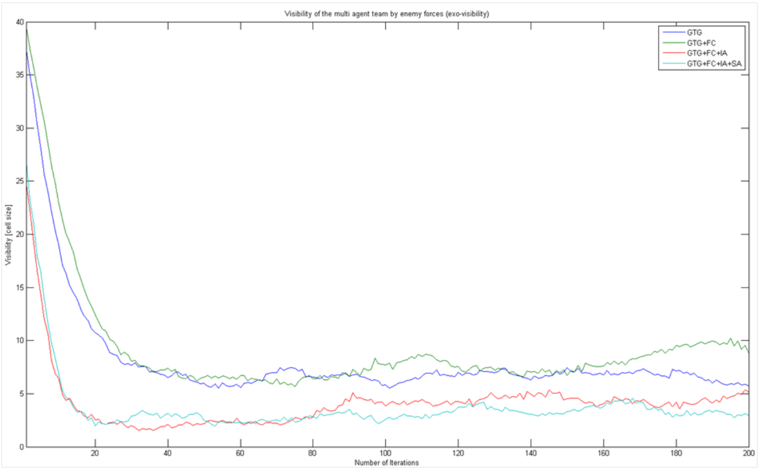

Fig. 8. Intruder avoidance capability.

(bottom 2 curves) lowers the mean visibility of the multi-agent team by enemy forces. Figure 8 shows that the visibility is minimized extremely well. It is therefore impossible to reduce the visibility to zero, as robots cannot evade enemies which they do not detect, but in general the intruder avoidance behavior succeeds to keep the visibility to within 5 times the cell size, which is very low.

4.2 Qualitative Analysis

Figure 9 shows the different steps of a simulation, indicating at different time steps the map with central robot (pink dot) and protective robots (blue dots) and enemies (red dots) and the goal position (green dot). It can be noticed from Figure 9 that the multi-robot team starts out as a totally unstructured group (initial robot positions are chosen randomly), and as time goes by, the team more and more organizes itself into a circular formation, while advancing towards the goal position and while avoiding contact with the enemy forces.

5 Conclusions

In this paper, we have presented a decentralized control strategy for multi-robot coordination for a field transport scenario. Using the presented multi-agent con-

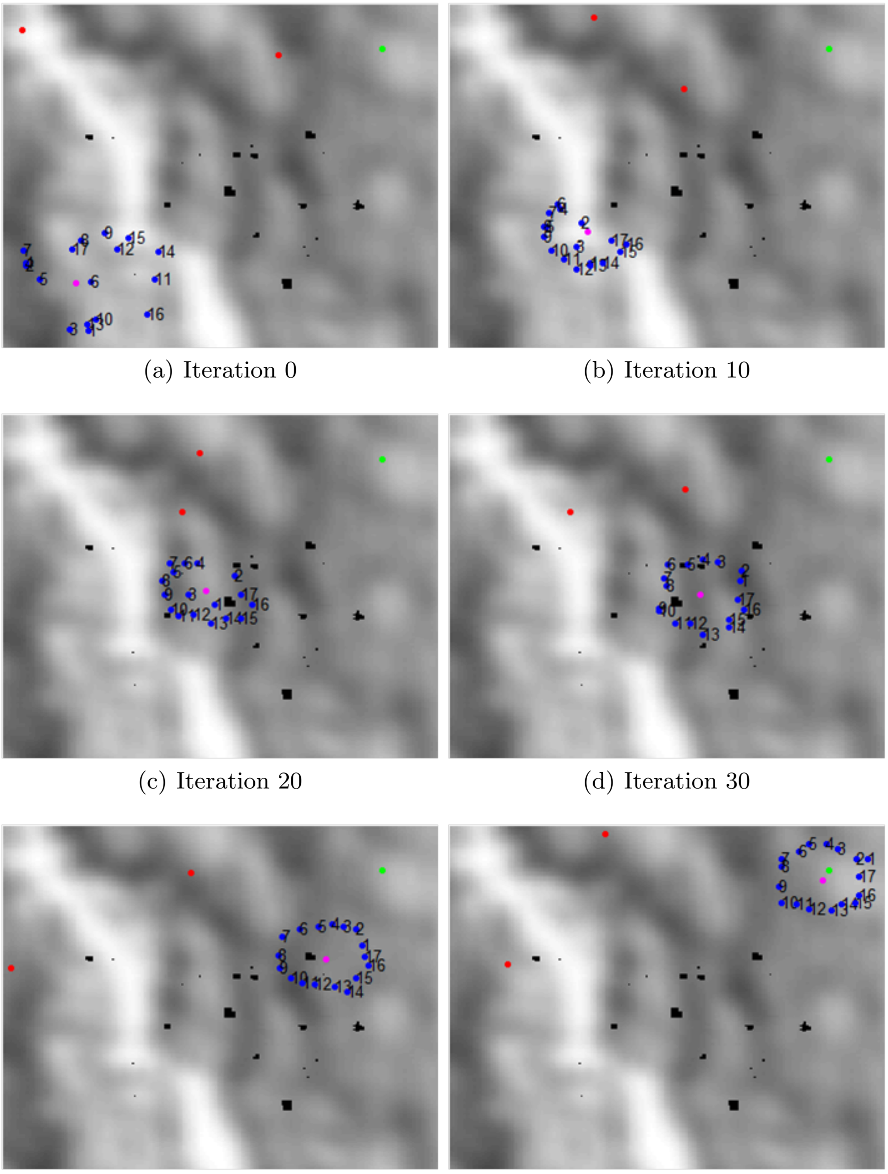

(a) Iteration 0

(b) Iteration 10

(c) Iteration 20

(d) Iteration 30

(e) Iteration 40

(f) Iteration 50

Fig. 9. Evolution of the multi-robot formation over the course of a field transport experiment with a central robot

trol architecture, it is possible to make teams of robots execute a well-defined task in a challenging environment. To achieve this, a behavior based framework was implemented, allowing multi-robot formation control in the presence of enemy forces. The algorithms presented here for multi-robot coordination were shown to achieve good results and to scale well with increasing the number of robots. The presented control strategy itself does not consider any limitations related to the number of robots, number of enemy forces or size of the simulation environment, although it is evident that increasing the problem complexity will also increase the simulation time. The analysis of Figures 6 to 9 shows that the presented multi-robot coordination strategy is capable of reaching the destination point while holding a formation and while avoiding intruders.

Acknowledgments. This research was funded by the Networked Multi-Robot Simulation (NMRS) project. The NMRS research project is funded by the European Defence Agency (EDA). Besides the Royal Military Academy, partners involved in this project are Diehl BGT Defence(Germany- project leader), Otomelara (Italy), Sener (Spain) and FGAN (Germany).

References

1. Pirjanian, P.: Behavior coordination mechanisms – state-of-the-art. T.R. IRIS-99-375, Institute of Robotics and Intelligent Systems, Univ. of South. California (1999)
2. Taylor, C.J.: Videoplus: a method for capturing the structure and appearance of immersive environments. IEEE Tr. on Vis. and Comp. Graphics 8(2):171-182 (2002)
3. Jennings, J.S., Whelan, G., Evans, W.F.: Cooperative search and rescue with a team of mobile robots. In: Proc. IEEE Int. Conf. on Adv. Robotics, pp. 193–200 (1997)
4. Feddema, J., Schoenwald, D.: Decentralized control of cooperative robotic vehicles. In: Proc. SPIE, Vol. 4364 (2001)
5. Balch, T., Arkin, R.C.: Behavior-based Formation Control for Multi-robot Teams. IEEE Transactions on Robotics and Automation 14:926-939 (1999)
6. Veherencamp, S.L.: Individual, kin, and group selection. Behavioral Neurobiology, Vol. 3: Social Behavior and Communication, pp. 354–382, Plenum Press (1987)
7. Das, A. K., Fierro, R., Kumar, V., Southall, B., Spletzer, J., Taylor, C.: Real-time vision-based control of a nonholonomic mobile robot. In: Proc. IEEE Int. Conf. on Robotics and Automation, pp. 1714-1719 (2001)
8. Mataric, M., Nilsson, M., Simsarian, K.: Cooperative multi-robot box pushing. In: IEEE/RSJ Int. Conf. on Intelligent Robots and Systems, pp. 556-561 (1995)
9. Arkin, R.:Behavior-Based Robotics. MIT Press (1998)
10. Pereira, G.A.S. et al.: Robot Navigation in Multi-terrain Outdoor Environments. The International Journal of Robotics Research 28:685–700 (2009)
11. Madhavan, R., Fregene, K., Parker, L.E.: Distributed Cooperative Outdoor Multirobot Localization and Mapping. Autonomous Robots 17:23-39 (2004)
12. Chen, Q., Luh, J.Y.S.: Coordination and control of a group of small mobile robots. In: Proc. of IEEE Int. Conf. on Robotics and Automation, pp. 2315-2320 (1994)
13. Parker, L.: Designing control laws for cooperative agent teams. In: Proc. of the 1993 IEEE International Conference on Robotics and Automation, pp. 582587 (1994)
14. Brooks, R.: A robust layered control system for a mobile robot. IEEE J. Robotics and Automation 2(1):14. (1986)