Robotics and Autonomous Systems 57 (2009) 1048–1055

Contents lists available at ScienceDirect

Robotics and Autonomous Systems

journal homepage: www.elsevier.com/locate/robot

Shared Potential Fields and their place in a multi-robot co-ordination taxonomy

Joseph L. Baxter [a,*], Edmund K. Burke [a], Jonathan M. Garibaldi [a], Mark Norman [b]

[a] Automated Scheduling, Optimisation and Planning Group, School of Computer Science, University of Nottingham, Nottingham, NG8 1BB, England, United Kingdom
[b] Merlin Systems Corp. Ltd., Plymouth International Business Park, Plymouth, PL6 5WR, England, United Kingdom

ARTICLE INFO

Article history:
Available online 6 August 2009

Keywords:
Potential field
Multi-robot systems
Search and rescue

ABSTRACT

Previously our novel Shared Potential Field (SPF) method has been introduced and compared against a non-sharing control in both simulation and laboratory settings. In this paper, extended from a paper presented at the CIRAS 2008 conference, we compare the SPF against an existing type of robot architecture, a hybrid robotic system. The SPF method is compared to the traditional potential field method, and it is shown that the SPF is less susceptible to the traditional limitations of potential fields. The SPF method's position in Farinelli's multi-robot taxonomy is also discussed, and it is shown that rather than being placed in one category it encompasses two, corresponding to the two levels of control within the architecture. In experiments, the multi-robot systems are given the task of traversing an unknown environment, in an attempt to locate a target. The metric of performance for this task was the time taken to find the target. Experiments show that the hybrid system showed similar performance to the non-sharing control. In contrast, our Pessimistic variant of the SPF outperformed the hybrid system in the cluttered environment, and the Optimistic SPF variant outperformed the hybrid system in the sparse environment. We conclude that the SPF method reacts more robustly to the dynamic nature of the real world, and so performed significantly better throughout the experimentation.

1. Introduction

In the previous work, we introduced a novel form of multi-robot co-ordination architecture, which we termed the 'Shared Potential Field' (SPF) architecture [2,3]. Two variations of the system were introduced, which differ in terms of their degree of belief in sensor readings: an *optimistic* variant of the SPF, which is more inclined to believe in the *absence* of obstacles, and a *pessimistic* variant, which is more inclined to believe in the *presence* of obstacles. From now on these are referred to as the 'Optimistic SPF' and the 'Pessimistic SPF', respectively. The Optimistic SPF and the Pessimistic SPF were compared against a non-sharing control in simulation, over a number of search problems varying in object density and the number of targets. Results showed that the SPF approaches significantly outperformed the non-sharing system. The fact that no difference was observed between the two SPF systems was a result of the lack of noise within the simulated environment.

In [4] the SPF approach was transferred onto a real robotic system and the single target search problem was repeated in a laboratory environment. Results showed that unlike in the simulation, only the Pessimistic SPF significantly outperformed the non-sharing control. This was attributed to the noisy nature of the real world, resulting in pessimism being the best practice. The SPF can be defined as a reactive robotic system. Therefore, it was decided to compare the SPF against a different type of robotic architecture. This paper is an extension to the previous work by Baxter et al. presented at CIRAS 2008 in which experiments comparing the SPF method against a hybrid system were discussed [5].

Hybrid systems were pioneered in the early work of Arkin [1], and are an attempt to merge the benefits of both deliberative and reactive systems. Deliberative systems plan all the actions they are going to take before performing them; as such they rely heavily on complete and accurate information about the environment. Reactive systems meanwhile do no perform and planning; each action taken is a consequence of sensory information. However, due to the reactionary nature of their actions reactive systems often find sub-optimal solutions to problems. In a hybrid system the high level goals of the system are achieved through a deliberative component, whilst the low level goals are achieved through a reactive component. For example, the deliberative component will plan the shortest path across a room taking into account static obstacles. Whilst traversing this path if an unexpected obstacle is discovered the reactive component will take over until the obstacle has been successfully passed and the original path rediscovered.

* Corresponding author.
E-mail addresses: jlb@cs.nott.ac.uk (J.L. Baxter), ekb@cs.nott.ac.uk (E.K. Burke), jmg@cs.nott.ac.uk (J.M. Garibaldi), marknorman@merlinsystemscorp.co.uk (M. Norman).

0921-8890/$ – see front matter © 2009 Elsevier B.V. All rights reserved.
doi:10.1016/j.robot.2009.07.023

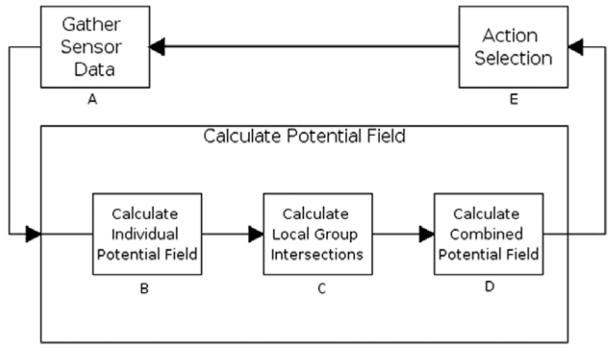

Fig. 1. Given some sensory input, individual potential fields are created. Robots within local groups share information and create Shared Potential Fields. Action selection is based upon these potential fields.

This paper is structured as follows: In Section 2 the SPF method will be described in detail. Two variants will be introduced a *pessimistic* and an *optimistic* approach in terms of belief in sensor values. The SPF method will then be compared to the traditional potential field method in terms of susceptibility to the known limitations of potential fields. The method's place within Farinelli's multi-robot taxonomy will also be discussed. Section 3 will describe the hybrid system that we will compare our SPF method against in this paper. The hybrid system consists of two modules, a deliberative path planner, and a reactive motor controller. In Section 4 the experimental setup will be explained in detail, including a description of the robots used and of the laboratory environment. Section 5 will present the results obtained, and will show that the SPF method significantly outperformed the hybrid system in almost all cases. Section 6 will conclude this paper, with a discussion of the results obtained and some proposals of possible future areas of research.

2. Shared Potential Field (SPF)

The outline of the processes involved in the SPF method is as follows: Individual robots construct potential fields from available sensor data (in the current system only ultra-sonic data is used), see Fig. 1A–B. Robots that are assigned to the same group then calculate local group intersections and share relevant potential field information, Fig. 1C. Each individual robot then creates a combined potential field using the shared information (Fig. 1D) and then makes the relevant action selection, (Fig. 1E). The non-sharing system that is used as a control in the following experiments only contains the processes A, B and E in Fig. 1.

The basis of our SPF is Coulomb's law of electrostatic force, as given by

$$F = k_C \frac{q_1 q_2}{r^2}, \tag{1}$$

where F is the force, q_1 represents the unit charge of the robot and q_2 the unit charge of an obstacle detected by an ultra-sonic sensor. r is the sensor reading from the relevant ultra-sonic sensor, i.e. the distance of the closest object. For simplicity, all obstacles and robots are given a unit charge of 1 and k_C (the electrostatic constant) we ignore (set to 1). Therefore, the calculation is now the inverse square law:

$$F = \frac{1}{r^2}. \tag{2}$$

So the smaller the reading from the ultra-sonic sensor the greater the force. The reason why we have simplified Coulomb's law in this manner, is that as we are only using ultra-sonic sensors to create the potential field, there is no way of distinguishing between objects, as such all objects have to be assigned the same unit value. Therefore the unit charge of the robot is also set the same unit value, in order to get the appropriate behaviour (attracted to regions of low resistance). The electrostatic constant is ignored as it is just a constant and has not direct affect in this application. As the system was designed with the Miabot Pro class of robot in mind, we calculate eight forces per robot corresponding to the fact that each robot has eight ultra-sonic modules in an array.

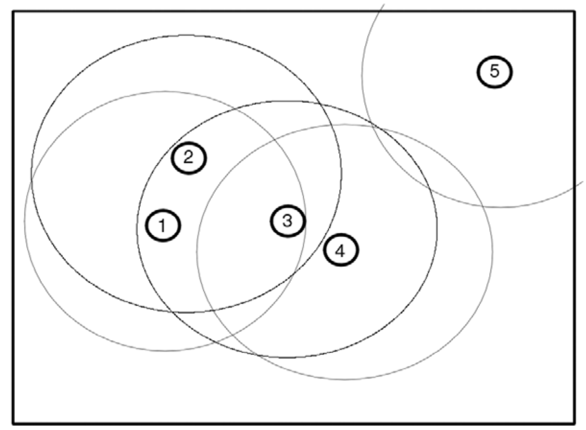

Fig. 2. Local groups: Robots 1, 2 and 3 form group A. Robots 2, 1 and 3 form group B. Robots 3, 1, 2 and 4 form group C. Robots 4 and 3 form group D. Finally, robot 5 is in a group of its own.

After each robot has calculated its own local potential field, local groups of robots are formed (see Fig. 2), which share this information. The process is similar in concept to "dynamic robot networks" in Clark et al. [6] but instead of sharing trajectories and plans, pose and potential field information is shared. Local groups are formed by sequentially going through the robots within the environment and assigning other robots to their group if they are within a given distance of that robot. Robots that move out of range are no longer considered a part of the local group. This process enables robots to join and leave multiple local groups throughout the life-time of the task. A distance of 70 cm was used in the laboratory as this is approximately double the boresight reflection [1] of the ultra-sonic sensors.

Once local groups have been assigned, pose and potential field information is shared. This is achieved through the following process. To provide a brief example we will use Fig. 3:

1. The robots are modelled within two-dimensional space. Ultra-sonic ranges are represented as lines on a plane and the sensory limits of a robot are represented as circles on a plane.
2. All ultra-sonic ranges that intersect the line A (the radical line) are available to share information. These lines are marked with a $*$ in Fig. 3.
3. Information is shared between these ultra-sonic and any of the other robot's ultra-sonic ranges that they intersect. It shares information with 3 ultra-sonic ranges (thick lines). In such cases, calculations are completed on intersections sequentially. In the example, the force of B would be compared with the force of 1, then the force of 2, then 3.

We have implemented two versions of the SPF approach, referred to as *Pessimistic* and *Optimistic* in terms of sensor noise. In the Pessimistic SPF, when two lines intersect their forces are compared and the highest value is used. In the Optimistic SPF the lowest value is used. In our example (Fig. 3), robot 3's ultra-sonic sensor detects the obstacle X. However, robot 4 detects obstacle Y. In the Pessimistic SPF, robot 4's force value, being the greater of the two, is used by robots 3 and 4, hence robot 3

[1] The boresight of an ultra-sonic sensor is the angle to which it is pointing. Boresight reflection refers to the response received by the ultra-sonic sensor from the ground.

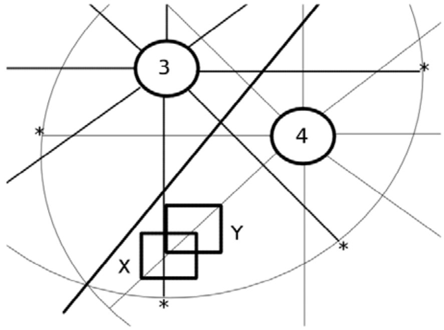

Fig. 3. Obstacle detection: Robot 3 detects obstacle X, robot 4 detects obstacle Y.

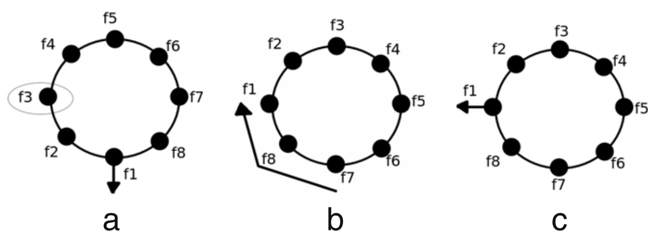

Fig. 4. Action selection: (a) The robot calculates the minimum force, f3, (b) The robot rotates towards the minimum force, (c) The robot moves forward (towards the minimum force).

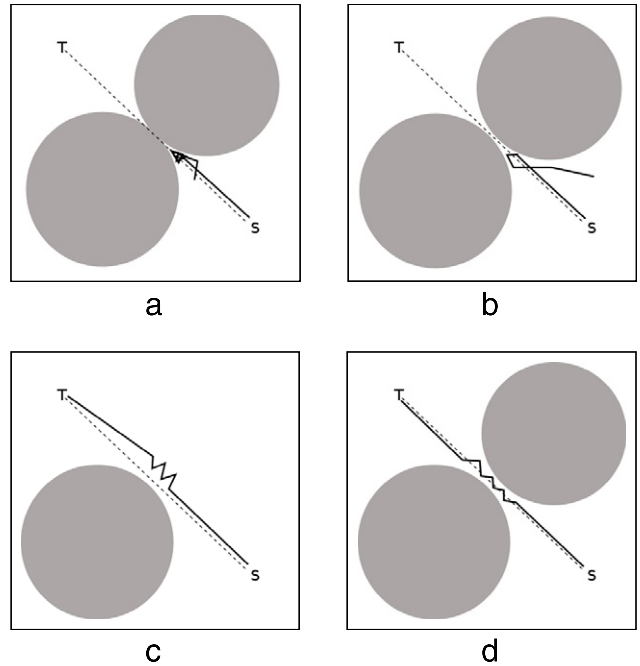

Fig. 5. (a) Trap. (b) No passage. (c) Oscillation in presence of obstacle. (d) Oscillation in narrow passage. The grey circles are obstacles. The dashed line represents the straight line path from the start location (S) to the target location (T). The solid line represents the path taken by the robot.

would detect an obstacle closer than would have been previously possible. Conversely, in the Optimistic SPF, robots 3 and 4 would use the smaller force value from robot 3 and so robot 4 would detect an obstacle further away than previously. As you can see the sequence of comparisons can make a huge difference to the resulting potential field, the difference being the scale to which the system is pessimistic or optimistic. Currently no further research has been done on the effects of the scale of optimism or pessimism. However, it would be of interest to see if the scale had any significant effect.

The desired advantage of the Pessimistic SPF is that it will be less vulnerable to false negatives (not detecting obstacles that are there). However, it will be more susceptible to false positives (detecting obstacles that are not there). The desired advantage of the Optimistic SPF is that it will be less vulnerable to false positives. However, it will be more susceptible to false negatives.

Once the combined potential field has been calculated (or not in the case of the non-sharing control), the minimum force $f_{\min}$ is discovered (Fig. 4(a)). The robot has a default forward motion within the environment unless it comes across an obstacle in its path (a force value of 25 or less, even though this is approximately 20 cm from an obstacle, due to lag in communications, the robot will not avoid the obstacle until a distance of about 5 cm), in which case the robot rotates towards $f_{\min}$ (Fig. 4(b)). Once the forward orientation of the robot equals the direction of $f_{\min}$, the robot resumes its forward motion (Fig. 4(c)).

2.1. Comparison to traditional potential field method

Koren et al. [10] identified the four major limitations of the potential field method:

1. *Trap situations*: These are situations in which cyclic behaviour between local minima occurs, as shown in Fig. 5(a). Trap situations occur when robots run into dead ends e.g. U shaped obstacles, where the robot's range sensors pick up objects to the front and sides of the robots. The potential field generated in such a case forces the robot into a cyclic behaviour. In these situations, it is often necessary for some global recovery mechanism to intervene, which often results in a sub-optimum solution, but at least the robot is no longer trapped.
2. *No passage*: Closely spaced obstacles bar passage, as shown in Fig. 5(b). No passage situations occur when a robot comes across two obstacles that are close together. However, there is enough physical space between the objects for the robot to navigate safely through. The potential field that is generated forces the robot away from the passage, so that in essence, the two objects are treated as one large object. Again the results are sub-optimal.
3. *Oscillations in the presence of obstacles*: This type of oscillation occurs when a robot passes an object — the resultant potential field generated forces unstable motion whilst passing the object. This results in a sub-optimal solution. An example is shown in Fig. 5(c).
4. *Oscillations in narrow passages*: This type of oscillation occurs when a robot is travelling down a narrow passage. The resultant potential field causes the robot to oscillate from one side to the other; this is a result of the robot being repulsed by the objects on either side of it sequentially, until the robot has escaped the passage. Again, the results are sub-optimal. An example is given in Fig. 5(d).

As described above the potential field method has a number of limitations. The SPF method presented in this paper reacts to these limitations as follows:

1. *Trap situations*: This is still a limitation within this system. It is also worth noting that currently no global trap recovery is employed. Hence, once a robot enters a U shaped object it is not guaranteed that it will ever escape.
2. *No passage*: As the system employs a default forward motion to all robots until they meet an object being directly in their path, it is not affected by this limitation directly. However, it is indirectly affected when the ultra-sonic sensors do not differentiate between objects. An example is shown in Fig. 6(a), where the robot (the square) is unable to distinguish between the two objects (grey circles) and hence creates a ghost object (the shaded region).
3. *Oscillation in the presence of obstacles*: Again the system is not directly affected by this limitation. Only if the ultra-sonic

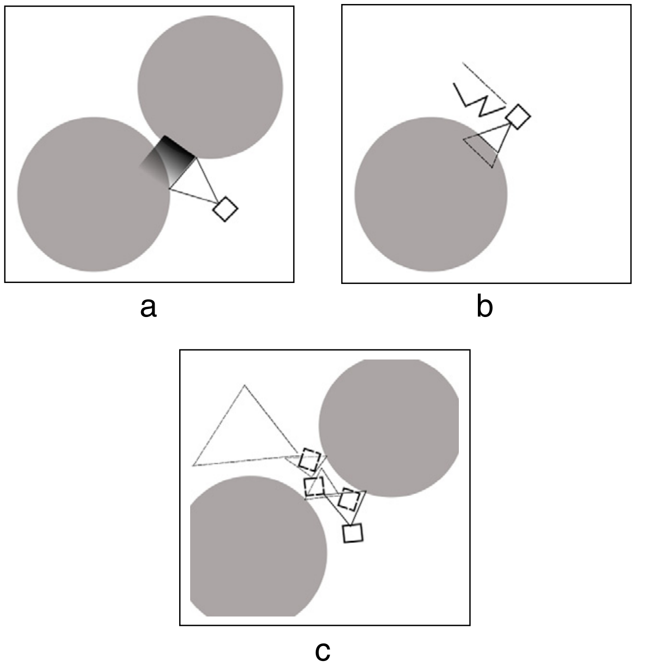

Fig. 6. (a) No passage. The shaded region is where the ultra-sonic sensor "thinks" an obstacle exists. (b) Oscillation in the presence of an obstacle. The bold line is the path taken when intermittent bad ultra-sonic echoes occur. (c) Oscillation in a narrow passage.

sensors produce bad echo data is the system affected. That is, the sensed distance to an object keeps changing, so as to make the robot believe that it would be advantageous to turn towards the object. An example is shown in Fig. 6(b), where the robot has a bad sensor reading (the dashed triangle) which results in a sub-optimal path being taken (bold line). The straight line is the desired optimal path.

4. *Oscillation in narrow passages*: If a robot meets the passage head on, as in the *No passage* case, it is not directly affected by this limitation. However, if a robot enters a passage at an angle, then oscillation can occur until the robot's orientation matches the orientation of the passage or the robot exits the passage. An example is shown in Fig. 6(c), where the dashed squares and triangles are the position of the robot and its forward ultra-sonic range in future time steps.

2.2. Place in Farinelli's multi-robot taxonomy

Farinelli's multi-robot taxonomy categorises multi-robot system based on the amount of communication and co-ordination within the system [7]. Unlike previous multi-robot systems, the SPF method presented in this paper does not fit neatly into one category, as the system possesses different characteristics at different levels. Each robot within the multi-robot system can move freely between the different levels of co-ordination throughout the given task.

The SPF method is *unaware* at the global level. The relevant section of the multi-robot taxonomy hierarchy is shown in Fig. 7. The global level is everything outside of the local group radius of an individual robot. At this level, robots are not assigned to local groups. The only source of sensory input to the potential fields is the ultra-sonic array, which only provides range readings to the nearest object. It is not known what this object is, as robots have no concept of team members or indeed other robots outside of their local groups. A blob-finder module is used to detect the target. However, it is not used for team member recognition. As shown in Fig. 8 groups A, B and C are *unaware* systems.

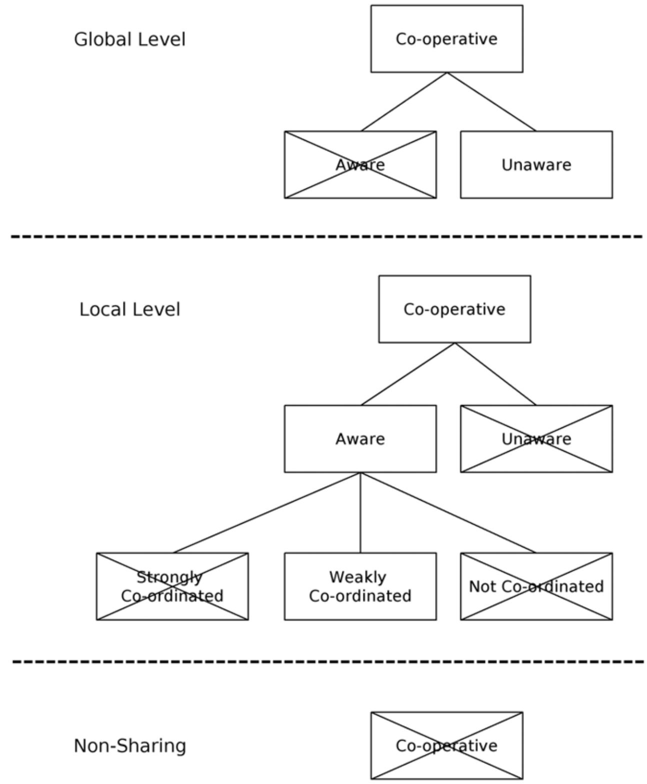

Fig. 7. Hierarchical view of the sharing potential field method's place within Farinelli's taxonomy. Crossed out categories are not implemented.

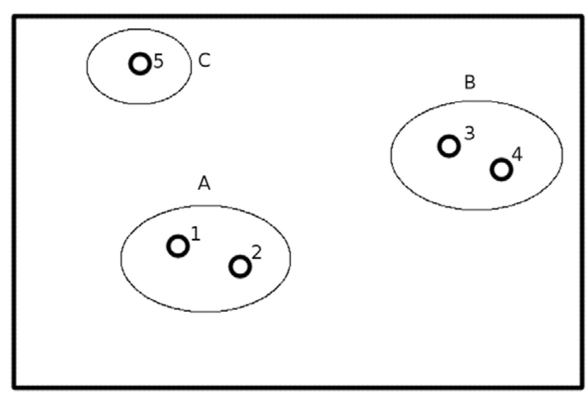

Fig. 8. Groups A, B and C are all *unaware* of each other. Robots 1 and 2 are *weakly co-ordinated*. Robots 3 and 4 are *weakly co-ordinated*. Robot 5 reverts to non-sharing behaviour.

At the local level, the system is considered *weakly co-ordinated*. The relevant section of the multi-robot taxonomy hierarchy is shown in Fig. 7. The local level is everything within the local group radius of an individual robot. At this level, robots are assigned to local groups—the members co-ordinate implicitly through the use of Shared Potential Fields. However, no explicit co-ordination occurs as robots do not have the ability to distinguish between team members from other objects within the environment as previously discussed. As shown in Fig. 8, robots 1 and 2 are *weakly co-ordinated*, as are robots 3 and 4, whilst robot 5 reverts to single-robot system, as depicted in the multi-robot taxonomy hierarchy shown in Fig. 7.

It is worthy to note that the system could converge to an entirely *weakly co-ordinated* system. This is more likely within small environments or in cases where the local group radius has been set to be arbitrarily large. Conversely, the system could also

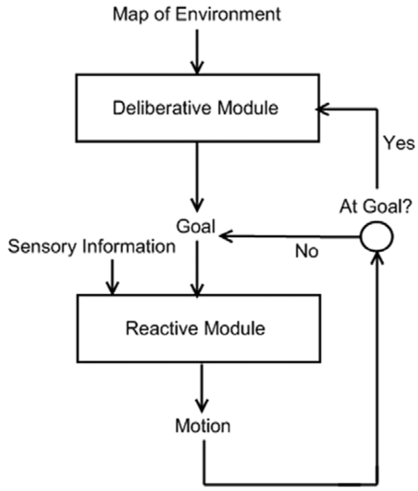

Fig. 9. Hybrid system architecture.

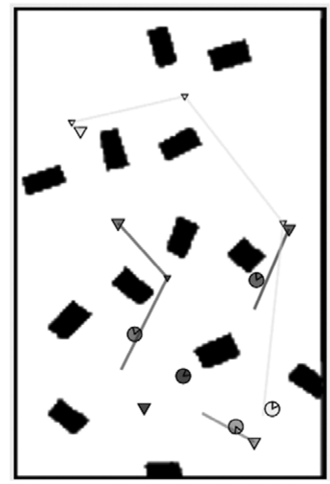

Fig. 10. Example path planning by the Wavefront propagation path planner.

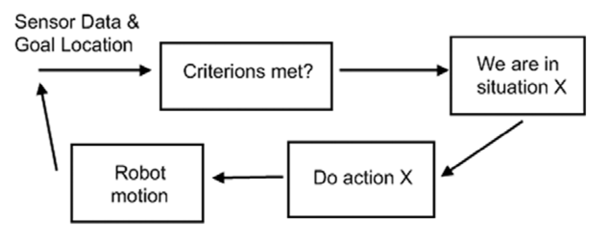

Fig. 11. Basics of the ND algorithm.

disperse to an entirely *unaware* system. This is more likely in large environments or in cases where the local group radius has been set to be arbitrarily small. It is believed that the hybrid of the two system types will be beneficial to the system, as co-ordination only occurs when needed. However, this co-ordination is not a necessity to task completion.

It is also envisaged that more tightly coupled tasks such as search and rescue may require a higher level of co-ordination than the *weakly co-ordinated* local level can provide. In such cases, the local level could be substituted with a *strongly co-ordinated* system, without disrupting the global level. As such, it would still be possible to design a system with low levels of communication and co-ordination during the "search" aspect of the task, only relying upon communication and co-ordination when necessary. That is, the "rescue" sub-task.

3. The hybrid system

The hybrid system we implemented was comprised of two modules. A diagram of the complete hybrid system architecture is given in Fig. 9. The first module, a deliberative module, was the Wavefront propagation path planner, which given a map of the environment calculated the shortest path to randomly generated targets. Random targets were generated using a lagged Fibonacci pseudo-random number generator over a uniform distribution. The targets x and y positions and θ orientation were all generated separately. This provided a visually nice spread of targets across the environments. It should be noted that the target generation did not take into account obstacles within the environment, as such if a target was generated within an obstacle, the Wavefront algorithm would request a new goal.

The robots in the system were unaware of the other robots. However, the current goals of all robots taking part in the experiment were globally available, and so the current goals of other robots were taken into account when generating new goals for individual robots (new goals were forbidden to be generated within a 20 cm radius of current goals), this is shown in Fig. 10.

The second module which was reactive in nature, was the Nearness Diagram (ND) algorithm developed by Minguez et al. [11,12], which enables the Miabot to avoid non-mapped obstacles. The basic algorithm is as described in Fig. 11. Given a goal location and current sensor readings, the algorithm attempts to match the current situation, to one of a list of pre-defined situations. These situations described the obstacle layout of the environment. Each of the pre-defined situations had a series of actions associated with it, which the robot attempted to complete. This process was continually looped until either the goal was reached or the robot "gave up" (determined by a pre-defined threshold value) i.e. the robot did not make enough progress towards the target in terms of distance or time. The main advantage of the ND algorithm over other conventional reactive controllers is its ability to actively avoid oscillation scenarios, which are prevalent in potential field based approaches [10]. 'U' shaped obstacles were actively avoided by the robot.

Both the Wavefront propagation algorithm and the ND algorithm were already implemented in Player [2] which made the deployment of the hybrid system within our robotics laboratory a simple process.

4. Experimental setup

The type of robot used in the experimentation was a Merlin Miabot Pro, see Fig. 12. Each Miabot was approximately 18 cm × 8 cm × 8 cm and was equipped with the following sensors/actuators:

- *Differential drive:* With an optical encoder resolution of 0.04 mm and a maximum reported speed of 3.5 m/s, this is limited to 1 m/s in the experimentation in order to allow accurate tracking by the overhead camera system. The Miabot is non-holonomic.
- *Ultra-sonic array:* With a range of approximately 3 cm–2 m and a field-of-view of approximately 360°. Small gaps of approximately 15° existed between each individual sensor.
- *Blob-finder:* A camera with a fixed forward orientation and 30° field-of-view. An on-board blob-finder algorithm tracked the environment for a single blob defined by an RGB value sampled from a single frame grabbed from a camera approximately 30 cm away from the target.

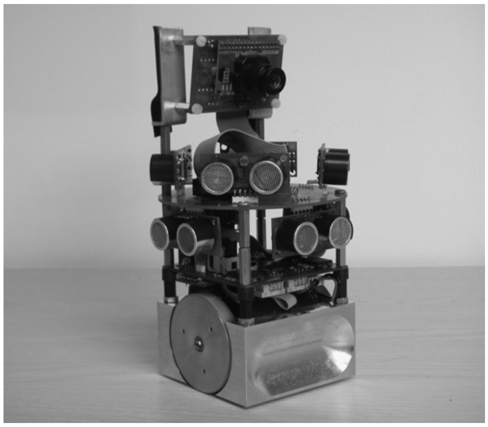

Fig. 12. Miabot Pro.

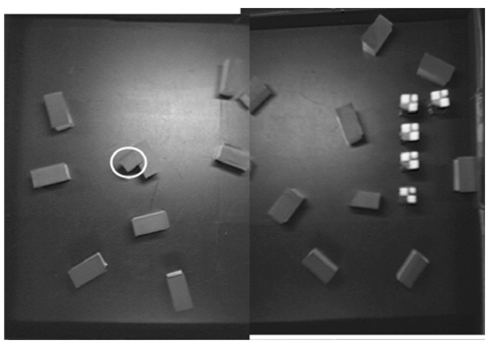

Fig. 13. Cluttered environment.

A global vision system tracked the position and orientation of the Miabots within the environment. This is due to limitations within the Miabot that rule out any reliance upon odometric readings.

The hybrid system completed the task on the same environments as the SPF methods. These environments differed from each other, in the number of obstacles present, from a sparse environment (20% obstacle coverage) to a cluttered environment (40% obstacle coverage), see Figs. 13–15. The positions of obstacles, the target and the Miabot starting location in the cluttered environment were generated randomly. To create the sparser environments four obstacles were removed randomly (one from each image quarter) to create environment 2 and another four to make environment 3. Each environment was approximately 1.7 m × 2.5 m and was enclosed within 15 cm high walls. The target was a pink can approximately 12 cm × 7 cm × 7 cm. Obstacles were rectangular boxes approximately 16 cm × 12 cm × 8 cm.

During the hybrid experiments the Miabots had a maximum velocity of approximately 0.1 m/s or 10 deg/s and a minimum velocity of 0.02 m/s or 5 deg/s. The on-board blob-finder algorithm in the camera was used to detect the colour of the target. The Miabot(s) came to a permanent halt once the target was found. As with previous experiments the Miabots were controlled through Player servers. The hybrid architecture was deployed as a Player client.

5. Results

Table 1 shows the average time taken, in seconds, for the hybrid system to complete the search task with a given number of Miabots. The best results for a given number of Miabots, in a given environment, are shown in **bold**. The hybrid system clearly performs best in the sparse environment (environment 3). However, this is still markedly worse than the two potential field sharing systems. Further detailed statistical analysis follows.

[2] An open source robot controller. More information about the Player project can be found at http://playerstage.sourceforge.net/.

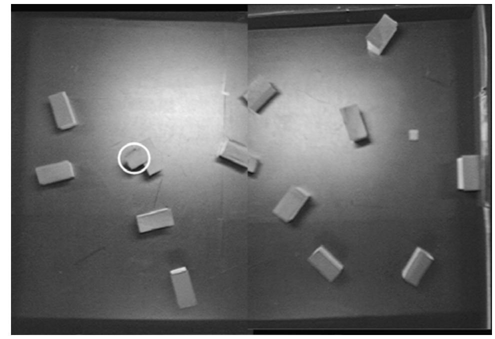

Fig. 14. Normal environment.

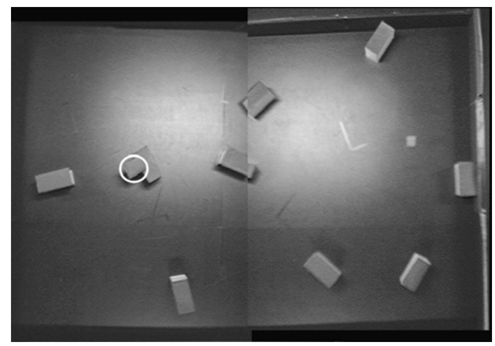

Fig. 15. Sparse environment.

5.1. Comparison between systems

The Kruskal–Wallis test [9] was used to compare the performance of the hybrid system against the three systems previously defined (in [2–4]), over the different environments. The parameters were set as follows, $k = 4$ (the number of samples), $N = 80$ (the total sample size) and $z = 2.4$ (to 1 d.p.), where the *critical z value* is for a level of significance of 0.1. Table 2 shows the differences between the means of ranks for the three potential field systems and the hybrid system, significant differences are in **bold**, non-significant differences are omitted where possible for brevity. As only values of 17.6 (to 1 d.p.) or more are significant it can be seen that in environment 1, the hybrid system performs as good as the non-sharing system, but worse than the Pessimistic SPF with 5 or less and 8 Miabots. The Optimistic SPF also performs better with 5 and 8 Miabots. In environment 2, the hybrid system performs as well as all of the potential field systems, in all but one case. In environment 3, the hybrid system has the same performance as the non-sharing and pessimistic systems. The Optimistic SPF performs better with 5 to 6 Miabots.

5.2. Comparison between environments

Again we employed the Kruskal–Wallis test to compare system performance across the three different environments. This time however, using the parameters $k = 3, N = 60$ and $z = 2.1$ to 1 d.p. As only values above, 11.8 (to 1 d.p.) are significant, it can be seen that the hybrid system clearly performs better in environment 3 than in environment 1. The performance across the other environments are mixed, with the hybrid system performing better in environment 2 than in environment 1 with 4 and 7 or more Miabots. Whereas, the hybrid system performs better in environment 3 than in environment 2 in just two cases (3 and 6 Miabots) (see Table 3).

Table 1
Mean completion (seconds) for each system in each environment for 2–8 Miabots, to 1 d.p.

	2	3	4	5	6	7	8
Env 1							
Non-sharing	266.0	**220.6**	216.3	232.8	219.3	236.1	172.4
Pessimistic	273.7	228.9	**177.9**	**174.6**	**187.3**	**212.9**	**167.2**
Optimistic	291.0	243.2	221.8	197.1	194.8	225.3	169.2
Hybrid	**264.6**	267.9	248.4	254.0	209.7	246.6	232.3
Env 2							
Non-sharing	219.4	184.4	160.1	219.4	180.9	184.5	167.4
Pessimistic	**182.5**	**161.8**	140.2	**115.4**	**164.4**	189.0	**153.7**
Optimistic	216.2	212.1	**134.5**	183.9	167.4	187.8	178.2
Hybrid	210.8	230.4	168.9	200.6	219.5	**155.0**	161.4
Env 3							
Non-sharing	189.4	161.6	143.9	154.9	111.8	**84.9**	102.1
Pessimistic	216.8	177.4	171.7	143.9	107.1	102.2	**100.1**
Optimistic	207.0	163.1	**117.0**	**97.3**	**65.0**	119.7	101.2
Hybrid	**156.7**	**142.5**	128.6	176.8	126.5	115.7	117.9

Table 2
Significant differences between the potential field systems (non-sharing (R_1), pessimistic (R_2) and optimistic (R_3)) and the hybrid (R_4) system (to 1 d.p.).

	4	5	6	7	8
Env 1					
$\overline{R_1} - \overline{R_4}$	−10.9	−6.1	3.1	−2.9	−17.5
$\overline{R_2} - \overline{R_4}$	**−23.2**	**−23.5**	−6.8	−9.9	**−19.9**
$\overline{R_3} - \overline{R_4}$	−8.3	**−18.2**	−4.3	−5.7	**−19.5**
Env 2					
$\overline{R_1} - \overline{R_4}$	−7.8	4.5	−10.8	6.0	0.6
$\overline{R_2} - \overline{R_4}$	−9.6	**−24.3**	−15.8	9.1	−4.2
$\overline{R_3} - \overline{R_4}$	−16.8	−5.4	−14.2	8.4	3.9
Env 3					
$\overline{R_1} - \overline{R_4}$	−0.5	−5.9	−10.6	−17.1	−11.8
$\overline{R_2} - \overline{R_4}$	6.9	−11.0	−11.2	−8.5	−13.1
$\overline{R_3} - \overline{R_4}$	−6.3	**−20.4**	**−22.1**	−4.4	−14.5

Table 3
Significant differences between means, hybrid system, to 1 d.p.

	2	3	4	5	6	7	8
Env 1 − Env 2	9.5	6.3	**14.3**	9.2	−1.2	**18.0**	**15.3**
Env 1 − Env 3	**18.6**	**21.8**	**23.9**	**14.0**	**18.3**	**26.3**	**24.5**
Env 2 − Env 3	9.2	**15.5**	9.5	4.8	**19.4**	8.3	9.2

5.3. Comparison between group sizes

The Friedman test [9] was used to compare the performance of each system with differing group sizes. With the following parameters; $k = 20$, $n = 7$ and $z = 2.8$ to 1 d.p. As such all values above 38.6 (to 1 d.p.), were significant. However, the number of Miabots in the experiment had no significant effect on the performance of the hybrid system in any of the environments. Therefore, the results are not shown.

6. Discussion

In this paper we have presented the findings from comparing the SPF against a hybrid system, in a laboratory setting.

The results section clearly shows that the hybrid system performed best in the sparse environment (environment 3), this is due to two main reasons. Firstly, the sparser the environment, the more valid paths/goals the Wavefront algorithm could plan, which lead to a greater area of the environment being explored. Secondly, as there were less obstacles in the environment the ND algorithm could move the Miabot at its maximum velocity more frequently.

The hybrid system performed as well as the non-sharing system but worse than both variants of the SPF. The Optimistic SPF performed better with 5 to 6 Miabots in the sparse environment (environment 3), the Pessimistic SPF performed better in the cluttered environment (environment 1) with 5 or less and 8 Miabots whilst the optimistic system performed better with 5 or 8 Miabots. This is because, whereas the hybrid system plotted a path of least resistance to a goal, the SPF reacted to the least resistance in the field, which was more adaptable to environmental changes than the Wavefront algorithm i.e. other Miabots moving in the environment. The embodied behaviour of the SPF to spread out (robots repulse each other), enabled the systems to cover a larger area of the environment more frequently than the hybrid system which relied upon a good random spread of target locations.

The number of Miabots within the hybrid system had no bearing on its performance. This was due to the relatively small increase in group size, which did not lead to any substantial increase in probability that the global planner would select target locations near the actual position of the target within the environment.

A comparison to the traditional potential field method was conducted in terms of susceptibility of the SPF method to the known limitations of the potential field method. It has been discussed that although the system is still susceptible to *trap situations* and *oscillations in narrow passages*, it is not directly susceptible to *no passage* and *oscillations in the presence of obstacles*. However, the more noise that is introduced to the environment, the more affected by the latter limitations the system becomes.

The SPF method's position in Farinelli's multi-robot taxonomy was also discussed. It was shown that the SPF method is *unaware* at the global level, yet *weakly co-ordinated* at the local level. This is beneficial to the system as co-ordination occurs when needed, but the method does not rely upon this co-ordination to complete a task.

6.1. Future work

In the current system, only ultra-sonic data is embodied within the SPF. Our architecture could be extended to allow the input of multiple sensors. In particular, blob-finder data could be fused within the SPF in order to encourage robots to move towards areas that possibly include targets. The blob-finder could also be used to aid robot dispersal, as currently other robots are treated as obstacles, rather than the special case obstacle which they are.

Currently, the potential fields are combined by either taking the maximum or minimum of two intersecting fields. It would be interesting to investigate the effect of using other methods to combine the two. For example, the mean of the two potential fields may be used, or a method based on the comparison of two robots internal beliefs [8].

We simplify Coulomb's law to an inverse square law, it would also be interesting to see what effect implementing Coulomb's law fully would have on system performance. This would enable us to assign different types of obstacles different unit values. For example, if other robots had a higher unit charge, the repulsion rate would increase, which may help the robots to disperse more evenly throughout the environment.

Finally, in the current system only the "search" aspect of search and rescue is attempted. It would be interesting to extend the system to implement the complete search and rescue problem; once the target has been found it needs to be rescued (taken to a designated position within the environment). This rescue operation may require multiple robots to complete it. That is, if the target's weight is greater than a single robots servo limitations, multiple robots will be needed to "rescue" it. Such a task would require the local level of the SPF to be changed to a *strongly co-ordinated* system. This adaptation of the system would not impede the current benefit of the architecture—only relying on communication and co-ordination when necessary, as the global level would not be effected by a change to the local level.

References

[1] Ronald C. Arkin, Reactive robotic systems, in: M.A. Arbib (Ed.), The Handbook of Brain Theory and Neural Networks, MIT Press, Cambridge, MA, 1995, pp. 793–796.

[2] J.L. Baxter, E.K. Burke, J.M. Garibaldi, M. Norman, The effect of potential field sharing in multi-agent systems, in: The 3rd International Conference on Autonomous Robots and Agents (ICARA 2006), December 2006, pp. 33–38.

[3] J.L. Baxter, E.K. Burke, J.M. Garibaldi, M. Norman, Multi-robot search and rescue: A potential field based approach, in: Subhas Mukhopadhyay, Gourab Sen Gupta (Eds.), Autonomous Robots and Agents, in: Studies in Computational Intelligence, vol. 76, Springer-Verlag, 2007, pp. 9–16.

[4] J.L. Baxter, E.K. Burke, J.M. Garibaldi, M. Norman, Multi-robot co-ordination using shared potential fields, Autonomous Robots, June, 2009 (submitted for Review).

[5] J.L. Baxter, E.K. Burke, J.M. Garibaldi, M. Norman, Real-world evaluation of a novel potential field sharing method, in: Proceedings of 5th International Conference on Computational Intelligence, Robotics and Autonomous Systems, CIRAS 2008, June 2008, pp. 79–85.

[6] C.M. Clark, S.M. Rock, J.-C. Latombe, Motion planning for multiple mobile robot systems using dynamic networks, in: Proc. IEEE Int. Conf. on Robotics and Automation, 2003, pp. 4222–4227.

[7] A. Farinelli, L. Iocchi, D. Nardi, Multirobot systems: A classification focused on coordination, IEEE Transactions on Systems, Man and Cybernetics B 34 (5) (2004) 2015–2028.

[8] D. Fox, W. Burgard, H. Kruppa, S. Thrun, A probabilistic approach to collaborative multi-robot localization, in: Autonomous Robots, in: Computer Science, vol. 8, Springer, Netherlands, 2000, pp. 325–344.

[9] J.D. Gibbons, Nonparametric Methods for Quantitative Analysis, third edition, American Sciences Press, Columbus, Ohio, 1997.

[10] Y. Koren, J. Borenstein, Potential field methods and their inherent limitations for mobile robot navigation. in: IEEE International Conference on Robotics and Automation, vol. 2, Sacramento, CA, 1991, pp. 1398–1404.

[11] Javier Minguez, Luis Montano, Nearness Diagram (ND) navigation: Collision avoidance in troublesome scenarios, IEEE Transactions on Robotics and Automation 20 (1) (2004) 45–59.

[12] Javier Minguez, Javier Osuna, Luis Montano, A Divide and Conquer Strategy based on Situations to Achieve Reactive Collision Avoidance in Troublesome Scenarios, New Orleans, 2004, pp. 3855–3862.

Joseph L. Baxter received a B.Sc. degree in e-commerce and digital business from the University of Nottingham, UK, in 2004. He is a researcher with the Automated Scheduling, Optimisation and Planning (ASAP) research group in the School of Computer Science at the University of Nottingham, UK. His primary research interest is in multi-robot systems.

Edmund K. Burke is Dean of the Faculty of Science at the University of Nottingham and he leads the Automated Scheduling, Optimisation and Planning (ASAP) Research Group in the School of Computer Science. He is a member of the EPSRC Strategic Advisory Team for Mathematics. He is a Fellow of the Operational Research Society and the British Computer Society and he is a member of the UK Computing Research Committee (UKCRC). Prof. Burke is the Editor-in-chief of the Journal of Scheduling, Area Editor (for Combinatorial Optimisation) of the Journal of Heuristics, Associate Editor of the INFORMS Journal on Computing, Associate Editor of the IEEE Transactions on Evolutionary Computation and a member of the Editorial Board of Memetic Computing. He is also the Research Director of EventMAP Ltd. and a Director of Aptia Solutions Ltd, both of which are spin out companies from the ASAP group. Prof. Burke has played a leading role in the organisation of several major international conferences in his research field in the last few years. He has edited/authored 14 books and has published over 180 refereed papers. He has been awarded 47 externally funded grants worth over £11M from a variety of sources including EPSRC, ESRC, BBSRC, EU, Research Council of Norway, East Midlands Development Agency, HEFCE, Teaching Company Directorate, Joint Information Systems Committee of the HEFCs and commercial organisations. This funding portfolio includes being the Principal Investigator on a recently awarded EPSRC Science and Innovation award of £2M, an EPSRC grant of £2.6M to investigate the automation of the heuristic design process and an EPSRC platform grant worth £423K.

Jonathan M. Garibaldi is an associate professor at the University of Nottingham, and is a member of the Intelligent Modelling and Analysis (IMA) research group. His main research interest is in the modelling of human decision making, primarily in the context of medical applications. His work to date has concentrated on utilising fuzzy logic to model the imprecision and uncertainty inherent in medical knowledge representation and decision making.

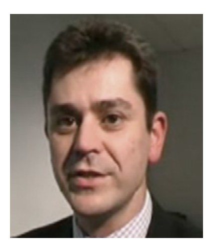

Mark Norman is the Managing Director of Merlin Systems Corporation Ltd. The Manufacturers of the Miabots and the global tracking system used in the experimentation dicussed in this paper.