



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.

© 2009 Elsevier B.V. All rights reserved.

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 not 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@merlinsystemscoorp.co.uk (M. Norman).

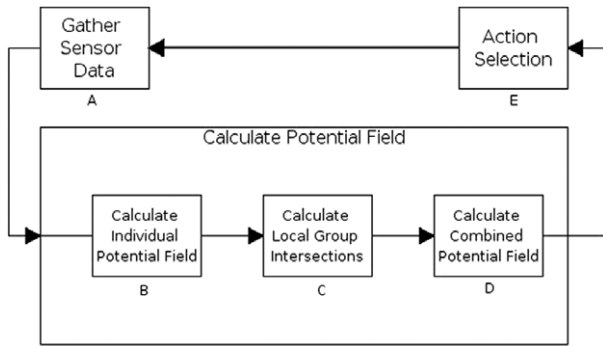


Fig. 1. Given some sensory input, individual potential fields are created. Robots within local groups share information and create Shared Potential Fields. Action selection is based upon these potential fields.

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




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.

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.

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 reflection¹ 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. Bore sight 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, f_3 , (b) The robot rotates towards the minimum force, (c) The robot moves forward (towards the minimum force).

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.

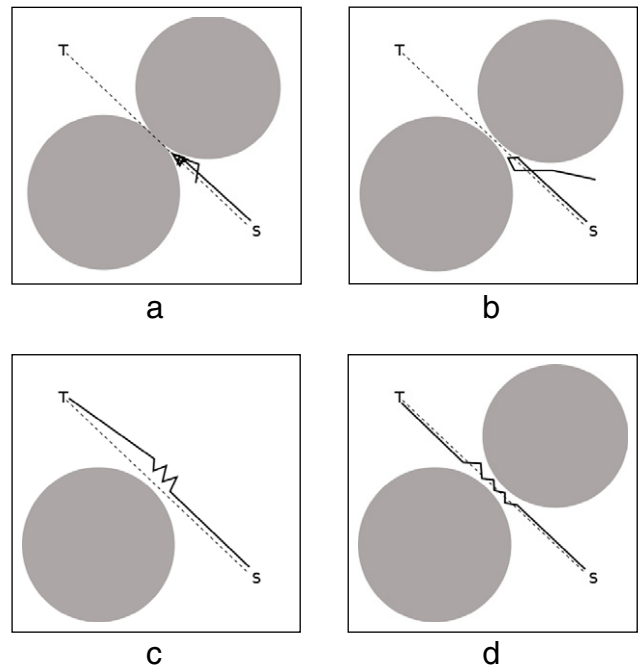


Fig. 5. (a) Trap. (b) No passage. (c) Oscillation in presence of obstacle. (d) Oscillation in narrow passage. The grey circles are obstacles. The dashed line represents the straight line path from the start location (S) to the target location (T). The solid line represents the path taken by the robot.

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

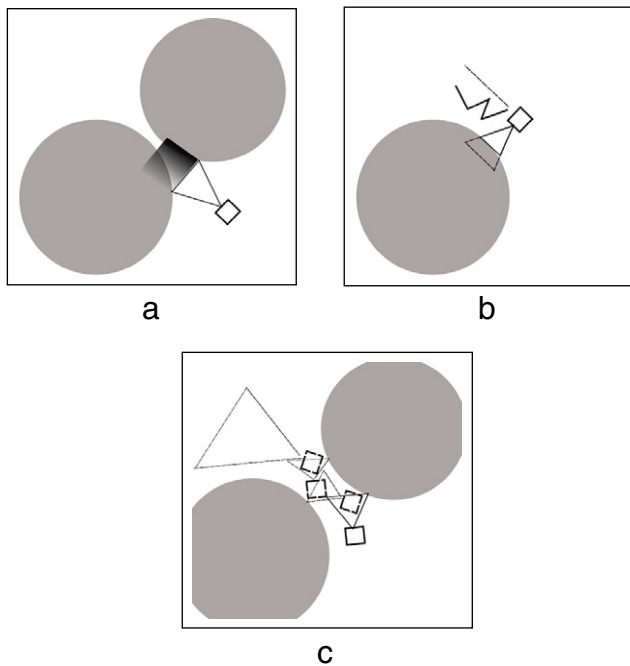


Fig. 6. (a) No passage. The shaded region is where the ultra-sonic sensor “thinks” an obstacle exists. (b) Oscillation in the presence of an obstacle. The bold line is the path taken when intermittent bad ultra-sonic echoes occur. (c) Oscillation in a narrow passage.

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

