

# Behavioural Approach for Multi-Robot Exploration

Haye Lau

ARC Centre of Excellence in Autonomous Systems (CAS)

Faculty of Engineering

University of Technology, Sydney

NSW, Australia

h.lau@cas.edu.au

## Abstract

Control strategies for mapping of unknown environments require a tradeoff between exploration and accuracy. One approach to balance these conflicting requirements is to use schema based behaviours to provide a means of exploration and manipulate the parameters controlling such schema to maintain the desired localisation and mapping accuracy. This paper provides such a behaviour based strategy for multi-robot exploration. Beyond the coordination of forces attracting agents to unexplored frontiers and repulsing against obstacles, as well as fellow agents, the proposed reactive approach is coupled with frontier based path planning to escape local minima otherwise encountered in a solely potential field based solution. The proposed behaviour based approach is designed so that it can be augmented to moderate movement based on criteria such as localisation accuracy in mapping applications. One planned application is the integration of proposed mechanism with probabilistic filtering for the mapping of indoor environments.

## 1 Introduction

This paper presents a behaviour based approach for exploration by multi-robot teams in office-like indoor environments. Starting from a known location, the robots following the behaviour disperse using and updating information in a shared grid-based map. Using potential fields to coordinate exploration behaviour, frontier based path planning was introduced to overcome the problem of robots stalling in local minima.

An exploring robot has the primary aim of covering previously unknown areas while avoiding obstacles. Employed as part of a team, a robot is also concerned with avoiding other robots while minimising duplication of effort. With additional considerations such as the need to maintain acceptable levels of accuracy or the desire to remain in range of neighbours, behaviour based systems provide a means to satisfy the potentially competing goals simultaneously [Balch and Arkin, 1998]. In the employed approach, the aims for exploration and collision avoidance are catered for in this manner, with strategies for accommodating the additional constraints

proposed. The work presented provides a basis which can be built upon to more effectively balance the overall exploration considerations of coverage, accuracy and speed [Makarenko *et al.*, 2002].

The remainder of this paper is organised as follows: Section 2 discusses related work in multi-robot mapping and Section 3 describes the proposed exploration approach. The implementation is outlined in Section 4 while Section 5 discusses future work for which this approach is designed to be a basis.

## 2 Related Work

Simultaneous Localisation and Mapping (SLAM) has been an area which attracts many researchers, with effort focussed mainly on obtaining overall best estimation for localisation and map construction [Zhang and Ghosh, 2000]. Moving beyond this passive approach, a more integrated strategy, where mapping is not the sole goal, could better facilitate the tradeoff between exploration and accuracy [Makarenko *et al.*, 2002].

The approach presented in this paper assumes the use of robots with 360-degree sensor coverage, and is inspired in part by the beacon system for the localisation of distributed Millibot robotic teams [Navarro-Serment *et al.*, 1999]. Equipped with ultrasonic sensors, the robots perform trilateration through the use of distance measurements to other Millibots. Instead of relying on landmarks provided by the environment, Millibots leap frog each other to maintain good position estimates as they traverse the unknown terrain. Similarly, Rekleitis *et al.* [2001] employed triangulation, where at any given moment a robot is stationary at a vertex of the environment while other robots divide map into triangles of free space, taking care to maintain visual contact.

The Millibots in the previous example leap-frog each other through the deliberate positioning of robots in the frame of reference of the team leader, however, behaviour based techniques for the maintenance of robot formation may be adapted for the purpose of exploration. Of the motor schemas of move-to-goal, avoid-static-obstacle, avoid-robot and maintain-formation used in one such example [Balch and Arkin, 1998], it can be seen that most reflect the aforementioned aims of exploration in a team. In this particular case, the move-to-goal schema

would need to account for the coverage of unknown areas, while maintain-formation may be adapted for the satisfaction of additional localisation accuracy or proximity constraints.

Yamauchi [1997] defined frontiers as unexplored areas bordering space known to be empty, which in turn act as points of interest for exploring robots. Extending the approach to multiple robots [Yamauchi, 1998], an acknowledged issue is the allocation of frontiers to robots. Traversal of the same path by more than one robot serves to negate the increased efficiency of multi-robot solutions. In coordinating a number of robots, Burgard *et al.* [2000] proposed a means by which the utility of an unexplored area for a robot is dependent on its projected sensor coverage by other robots. In addition to taking into account the costs of reaching target points, the solution thereby assigns target areas to robots sufficiently apart from each other to reduce exploration time.

Although vector summation is used to coordinate the competing schemas in a behavioural solution, the combination of attractive and repulsive forces can trap a robot in local minima, thereby nullifying its purpose of exploration. Local minima can manifest themselves through a potential field summation of zero magnitude, or a sequence of commands which effectively instructs a robot to steer in a small circular path. The former may occur when the robot is in an empty area with remaining frontiers beyond range of influence, the latter near multiple obstacles which serve to push a robot back and forth. In a multi-robot environment, situations may also arise where neighbouring agents act to trap a robot in this manner.

Although careful setting of schema parameters may reduce the likelihood of robots being trapped by obstacles, the nature of the problem is such that it would be impossible to exhaustively test a solution against all combinations of all robot positions in different environments. A number of techniques have been raised for the circumvention of local minima in potential field based solutions. A common method is to introduce noise or induce a random walk until the robot escapes [Arkin, 1998], while another is to develop artificial potential fields which do not exhibit local minima [Xiaoping and Ko-Cheng, 1997]. Introducing additional repulsion against areas recently visited by the robot [Arkin, 1998] may also be effective, particularly where a robot stalls in the absence of nearby attractors. Simple methods such as the switch to wall-following behaviour for the sole purpose of escaping minima have also been examined [Xiaoping and Ko-Cheng, 1997].

Makarenko *et al.* [2002] present an example of an integrated approach which moves towards balancing exploration strategy with considerations of map coverage, accuracy and exploration time, through evaluating the utilities of information gain with different destinations. Commencing with the simple behaviour presented in this paper, further work is envisioned to facilitate similar tradeoff of exploration requirements via an alternate approach.

### 3 Exploration Behaviour

The following section presents the behavioural approach employed for exploration in a grid based map.

### 3.1 Problem Statement

The aim is to arrive at a simple exploration behaviour to enable the mapping of an indoor office-like environment by one or more robots. The robots each have a 360-degree sensor (for example sonar, omnidirectional vision cameras or 180-degree lasers mounted back-to-back), and explored areas are updated to a shared grid-based map. Beyond a centralised map, in the simplest form each robot is responsible for determining its own motion without explicit inter-robot coordination.

This simple underlying exploration behaviour is to form the basis for additional mechanisms to satisfy constraints such as localisation accuracy, which shall be further discussed.

### 3.2 Architecture

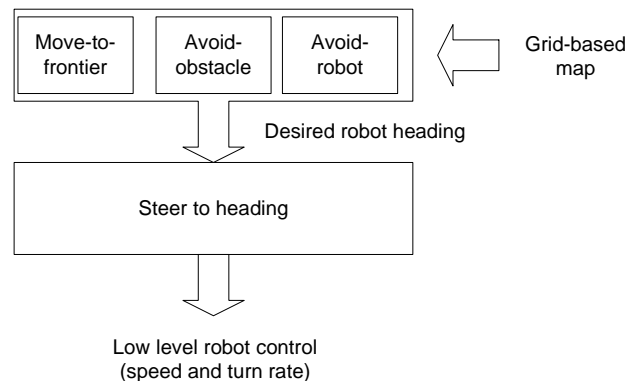


Figure 1 Structure of exploration approach

To facilitate individual robot control, the architecture is structured into two main layers. The output of the top behavioural layer represents desired robot heading as arrived at following the arbitration of motor schemas, which in turn is interpreted by the bottom layer in steering a robot. The steering mechanism at present only involves correcting a robot towards desired heading at a constant speed, with the corrections reducing in magnitude as the heading is approached. However, as shall be elaborated in discussion of future works, room exists for additional control to be exerted over the overall exploration pattern at this layer.

### 3.3 Exploration Schemas

A schema exists to take into account each of the main stated requirements of an exploring robot. As the contribution of each schema is represented as a potential field, vector summation suffices to produce a combined result. A shared grid map provides input to two of the schemas, with inter-robot distances, obtained possibly through alternate ranging means, providing a third.

#### Social Potential Fields

Unlike a ballistic move-to-goal schema [Arkin, 1998], where a fixed gain value is used for a field pointing towards a particular goal, exploration requirements clearly point towards the need for attraction to decrease with distance and repulsion to increase with proximity. To this

end, the Social Potential Fields as defined by Reif and Wang [1999] were adapted.

$$f(d) = \frac{-c_1}{d^{\sigma_1}} + \frac{-c_2}{d^{\sigma_2}} \quad (1)$$

where  $c_1, c_2 \geq 0$ ,  $\sigma_1 > \sigma_2 > 0$

Designed primarily to govern interactions between pairs of robots (hence social), the attractive and repulsive terms of the force law dominate depending on distance, culminating at an equilibrium distance where robots are content to stay in present relative position. In contrast to wavefront propagation method for navigation, potential vectors only need to be calculated locally for a restricted number of map cells surrounding the robot, constraining the required amount of computation. The tuning of the various parameters serve to adjust the clustering behaviour exhibited. As in this case only either attraction or repulsion applies for a frontier, obstacle or fellow robot, a gain and applicable range is set for each of the schemas. In case of an attractive schema, the gain set corresponds to  $c_2$ , with  $c_1$  set to zero. The converse is true for repulsive schemas. The relevant potential equation supplies the magnitude of a schema vector, with the direction provided by the bearing from robot to a given cell. Unlike the formula listed in (1),  $\sigma_1, \sigma_2 = 2$  under the proposed approach.

The overall force exerted on a robot is therefore:

$$F_{total} = \sum_{i=1}^k F_{frontier,i} + \sum_{i=1}^k F_{obstacle,i} + \sum_{j=1}^m F_{robot,i} \quad (2)$$

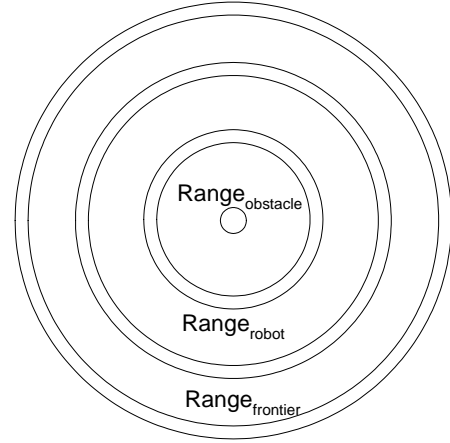
where  $k$  represents the number map grid cells and  $m$  the number of exploring robots. The remaining terms shall be discussed in the subsequent sections.

### Schema parameters

The following schemas were used in this example of multi-robot exploration. Table 1 outlines the parameters to be configured to affect overall behaviour. Please note that where a schema does not apply to a particular cell, for example, move-to-frontier for an obstacle, its corresponding force is set to be zero prior to summation in (2).

**Table 1 Parameters**

PARAMETER	DESCRIPTION
$d_{frontier}$	Range of effect of frontiers
$Gain_{frontier}$	Attractive gain exerted by each frontier
$d_{obstacle}$	Range of effect of obstacles
$Gain_{obstacle}$	Repulsive gain exerted by each obstacle
$d_{robot}$	Range of effect of other robots
$Gain_{robot}$	Repulsive gain exerted by each robot



**Figure 2 Typical ranges of influence**

### Move-to-frontier

Similar to the controlled move-to-goal schema outlined by Arkin [1998], the move-to-frontier schema instead acts as a force of attraction towards a frontier. A simplifying assumption is made from the definition by Yamauchi [1997], where in this case any unexplored grid cell adjacent to a known empty cell, without any consideration of robot size, is assumed to be a frontier. Unlike typical navigation applications where one or a small number of goals are specified, in this case the schema applies to all frontiers within a range of the robot. For example, when a robot begins in an unexplored map, frontiers all along the 360-degree sensor coverage circumference provide attraction. The attraction force of a frontier is given as:

$$F_{frontier,i} = \begin{cases} Gain_{frontier} / d_i^2, & \text{if } d_i \leq d_{frontier} \\ 0, & \text{if } d_i > d_{frontier} \end{cases} \quad (3)$$

where  $d_i$  represents the distance of  $i$ th map cell from robot. The parameters to be set in the move-to-frontier schema are the move-to-frontier attractive gain  $Gain_{frontier}$  and the frontier attraction range  $d_{frontier}$ . To be of use,  $d_{frontier}$  naturally needs to be set greater than a robot's effective sensor range to apply to relevant frontiers. The range of attraction however is constrained to an extent by the amount of computation required.

### Avoid-obstacle

Whereas the previous scheme focussed on attracting robots, this schema performs the reverse by avoiding obstacles through repulsion. The repulsive portion of the potential law has the useful property of providing for collision avoidance, as force approaches infinity as the distance to obstacle becomes zero [Reif and Wang, 1999]. The repulsion force exerted by an obstacle is given as:

$$F_{obstacle,i} = \begin{cases} -Gain_{obstacle} / d_i^2, & \text{if } d_i \leq d_{obstacle} \\ 0, & \text{if } d_i > d_{obstacle} \end{cases} \quad (4)$$

The parameters associated with this schema are the avoid-obstacle repulsive gain  $\text{Gain}_{\text{obstacle}}$  and range of effect  $d_{\text{obstacle}}$ .

Coverage of frontiers close to obstacles is improved when  $d_{\text{obstacle}}$  is significantly smaller than  $d_{\text{frontier}}$ . This ensures that robots are free to explore unimpeded without being influenced by obstacles which pose little present risk for collision. Due to the relatively small area within which obstacles exert an effect, currently all obstacle cells enclosed by  $d_{\text{obstacle}}$  are taken to contribute towards overall repulsion. This approach may need to be modified to apply only to obstacles within line of sight of a robot, neglecting the other explored side of thick walls, if  $d_{\text{obstacle}}$  is significantly increased.

### Avoid-robot

Acting to avoid collision and encourage distribution of exploration effort, the avoid-robot schema enforces separation between robots. The parameters associated with this schema are the avoid-obstacle repulsive gain  $\text{Gain}_{\text{robot}}$  and range of effect  $d_{\text{robot}}$ . As robots can approach each other at twice the speed they would obstacles, intuitively there may be a need to set  $d_{\text{robot}}$  larger than  $d_{\text{obstacle}}$ . However, this needs to be balanced against the possibility of a robot being “forced” into a static obstacle by a transient neighbour. The repulsive force exerted by one robot on another is given as:

$$F_{\text{robot},j} = \begin{cases} -\text{Gain}_{\text{robot}} / d_j^2, & \text{if } d_j \leq d_{\text{robot}} \\ 0, & \text{if } d_j > d_{\text{robot}} \end{cases} \quad (5)$$

where  $j$  represents the number of a neighbouring robot on the map. Please note that while the social potential fields as employed by Ref and Wang [1999] have a balancing attractive component to keep robots within range, a separate attractive schema shall be introduced should the need arise in future work to maintain accuracy in this fashion.

### 3.4 Exploration Algorithm

The following steps are followed by a robot in determination next direction of movement.

1. Update map with sensor input
2. For each frontier within  $d_{\text{frontier}}$  of robot, calculate  $F_{\text{frontier},i}$  for move-to-frontier schema
3. For each obstacle within  $d_{\text{obstacle}}$  of robot, calculate  $F_{\text{obstacle},i}$  for avoid-obstacle schema
4. For each neighbouring robot within  $d_{\text{robot}}$  of robot, calculate  $F_{\text{robot},j}$  for avoid-robot schema.
5. Supply angle of summed vector  $F_{\text{total}}$  as desired heading to steering layer

### 3.5 Escaping Local Minima

As previously stated, local minima in a potential field based system can manifest itself through a vector summation of zero magnitude, or a sequence of summation results which effectively instructs a robot to steer in a small circular path. The former may occur when the robot is in an empty area with remaining frontiers beyond the  $d_{\text{frontier}}$ , the latter near multiple obstacles which

serve to push a robot back and forth. In a multi-robot environment, wandering robots have the possibility of being trapped in particular configurations.

Regardless of the method employed, once a robot is detected as being in local minima, there is a need for a modification of behaviour until such time that the situation no longer applies. With unexplored map regions already primarily driving the overall exploring behaviour, it is clear that navigating to a remaining frontier presents a viable solution.

### 3.6 Path Planning

In the presented system, a robot is assumed to be trapped in local minima when the sum of all applicable schemas is zero or if the “box” area of movement traced by an exploring robot in a given period of time falls below a specified size, suggesting the lack of movement. In practice, this heuristic merely confirms what becomes apparent to an observer when a robot is stalled.

Drawing on frontier based exploration [Yamauchi, 1997], in the presented system a path is plotted to the nearest frontier within line of sight when a robot becomes trapped. Only when no such frontiers exist are other unexplored areas considered. This distinction is enforced to facilitate more comprehensive exploration of the current room by a robot prior to moving on. In the absence of such a requirement, robots trapped near walls become liable to plan paths to a room on the other side, especially if it is already being explored by another robot and hence presents numerous frontiers. With visible frontiers given preference, comparatively shorter paths are then followed. Moreover, this approach reduces the need for robots to pass through potential bottlenecks such as doorways.

The above path planning takes place through the use of the A\* algorithm. A commonly used technique optimal given admissible heuristics, in evaluation the total cost of a given map cell  $f(n)$  is equal to the sum of the cost of arriving at the cell  $g(n)$  and the heuristically estimated cost to goal  $h(n)$ . In the presented system, the cost  $g(n)$  of travelling from one cell to the next is set to be 1, but in order to arrive at paths sufficiently distant from obstacles to be safely traversable by a robot, cells within a set range of a known obstacle is assigned higher costs, decreasing with distance. The heuristic cost  $h(n)$  is defined to be the straight line distance of the cell to the goal. This approach is not optimal as the heuristic may provide an overestimate, but despite being inadmissible it provides adequate performance in practice.

To ensure frontier reachability, the A\* algorithm is only given freedom to search through the space of open map cells. While almost ensuring suboptimal paths to goal in predominantly unexplored maps, this avoids the impractical situation of charting a path through a yet to be discovered obstacle.

### 3.7 Path Following

The navigation approach employed in this case differs from other examples of frontier based exploration. In order to achieve the stated aim of developing a simple exploration behaviour, the presented system takes advantage of the interaction already present between the behavioural and steering layers. Once a path is found in the search space, the cells through which it traverses are

then marked on an individual robot's own internal map, used only for this purpose, as being unexplored. Without affecting the behaviour of other robots, this ensures that a robot is attracted to follow a given path, as the cells appear as frontiers, without the need for additional steering mechanisms for this purpose. To enable closer path following, the path cells are marked explored (open) again only when within  $d_{obstacle}$ . As  $d_{obstacle} < d_{frontier}$  and should be set to less than the sensor range, this allows the plotted path to retain more influence on the robot prior to being marked as traversed. Once the robot is within sensor range of a frontier (not necessarily the goal of its current path), the generated path is then removed from the map and normal exploration resumes. Figure 3 shows an example where a robot (left) stalled in the corner of a fully explored room is able follow plotted path back to a frontier in the corridor. Assignment of higher costs to grids near obstacles enabled path planning to provide a safe path through the doorway.




Figure 3 Robot following path to escape minma

The advantages of integration with existing motor schemas through the temporarily treatment of cells on the path as frontiers are twofold. Firstly, it reuses an existing means of reaching a goal without the need for additional control or steering. Secondly, it retains the obstacle-avoidance functionality already developed while doing so. Combined with a path given sufficient room around walls and obstacles, the approach appears effective in enabling a robot to escape local minima and resume exploration.

### 3.8 Combined Exploration Algorithm

The initially presented sequence is augmented to implement the behaviour based exploration.

1. Update map with sensor input
2. If following path and frontier is visible within sensor range, mark all internal map cells as being empty
3. Add current location to record of past movement

4. For each frontier in shared map or unexplored cell in internal map within  $d_{frontier}$  of robot, calculate vector  $F_{frontier,i}$  for move-to-frontier schema
5. For each obstacle within  $d_{obstacle}$  of robot, calculate vector  $F_{obstacle,i}$  for avoid-obstacle schema
6. For each neighbouring robot within  $d_{robot}$  of robot, calculate  $F_{robot,j}$  for avoid-robot schema
7. Sum vectors
8. If sum = 0 or if robot has not travelled sufficiently far in a past time period, find nearest frontier, with absolute preference given to frontiers within line of sight. Plot path to frontier and mark path cells as unexplored in internal map
9. Supply angle of summed vector  $F_{total}$  as desired heading to steering layer

## 4 Implementation

### 4.1 Simulation

The proposed exploration behaviour was developed to work with the Player/Stage multi-robot simulation environment [Gerkey *et al.*, 2001]. A range of environments from an empty map to a corridor with rooms were used. The behaviour was tested with simulated Pioneer II robots with two 180-degree lasers mounted back-to-back to facilitate omnidirectional sensing, each with an arbitrarily set range of 2.5 metres. The typical environment took place on a 40m x 40m map, with map grids the size of 10cm x 10cm.

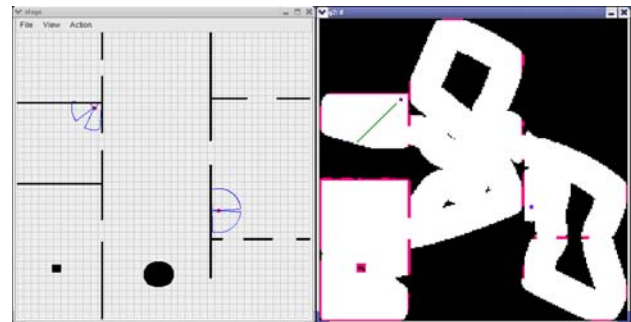


Figure 4 Simulated exploration of simple environment, actual map under Stage shown on left

### 4.2 Parameter Configuration

In determining the value of schema parameters, a simple environment was first tested with an individual simulated robot. Given the guidelines as previously stated,  $Gain_{frontier}$  is first arbitrarily set with  $d_{frontier}$  fixed to be a certain multiple of the sensing range.  $Gain_{obstacle}$  and  $d_{obstacle}$  are adjusted through successive simulations to ensure coverage of frontiers near walls without collisions.  $Gain_{robot}$  and  $d_{robot}$  are then set through testing with multiple robots, although as implemented they are also scaled from the values of their obstacle counterparts. The overall configuration is then tested in a more complex environment, with fine adjustments made.

### 4.3 Results and Limitations

As the aim of current research is to develop an underlying behaviour prior to implementing tasks such as

localisation, simulated robots are given true positional information. It is acknowledged that filtering or other approaches could similarly provide an accurate threshold when determining whether a cell in an evidence grid indicated by a laser hit is indeed an obstacle, but for the purposes of developing the behaviour false laser hits were discounted. Line of sight, taking into account map areas already detected to be obstacles, is used to determine whether or not an unexplored map cell within sensing range should be marked as open.

It is also recognised that in terms of exploration behaviour, it can be seen that a locally greedy approach to the nearest frontier, without explicit coordination with other robots, would not produce a time optimal solution.

Qualitatively, multiple robots employing this behaviour cover an initially unexplored indoor area at a much quicker rate than an individual exploring robot, as would be expected. With robots kept apart through repulsion, on occasions when more than one agent attempts to map the same area, eventually the locally greedy pursuit of frontiers lead to a divergence of paths. This effect is more apparent where deliberate path planning for escape is required, with robots then observed to spread out for effectively individual exploration until only small regions of frontiers remain. Quantitative timing results, focussed on the effects of controlling forward and turning speed on rates of exploration, shall be obtained as part of evaluation of future work.

## 5 Future Work

### 5.1 Refinements

A number of possible refinements have been identified for the underlying behaviour presented:

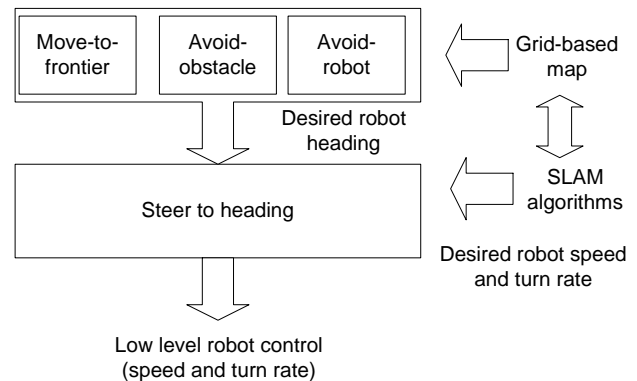
- Move beyond a simple greedy choice of frontiers to meet accuracy or time performance needs [Burgard *et al.*, 2000]. In addition to selecting a more appropriate frontier when trapped in local minima, the information utility of frontiers may also be used to modify an applicable  $\text{Gain}_{\text{frontier}}$  [Makarenko *et al.*, 2002].
- Addition of sensor-linked obstacle-avoidance subsumption behaviour to further minimise risk of collisions.
- Setting of gains with the aid of genetic algorithms as opposed to informed trial and error [Arkin, 1998].

### 5.2 Further Applications

Canvassing related work, two main applications present logical extensions. In the implemented system, multiple robots explore areas individually, influencing others only for the purpose of collision avoidance while incidentally reducing effort duplication. In teams where robots act as landmarks for one another or otherwise rely on proximity [Rekleitis *et al.*, 2002; Grabowski and Khosla, 2001], the exploration behaviour can be augmented to control rate of exploration. For example, where robots themselves present as beacons with known positional uncertainties, modifications may be made to affect speed and pose such

that visibility is maintained. The robots or smart sensors can then spread automatically into a sustainable pattern, enabling coverage of the environment while remaining close enough to provide support. Alternatively, the approach may be modified to naturally reproduce the leapfrogging or “bounding-overwatch” pattern employed by teams such as Millibots [Navarro-Serment *et al.*, 1999].

Similarly, the approach can be adapted to work in tandem with feature-based SLAM. The following architecture outlines one option currently under consideration:



**Figure 5 Mapping algorithm control of robot speed and turn rate**

As can be seen in Figure 5, the exploration behaviour proceeds as originally proposed, with the additional ability to externally set the speed and turn rate at which a robot travels at a lower level. While the potential fields are summed as before to yield a direction of movement, a filter performing the SLAM algorithm is also given the responsibility of determining a robot’s speeds. In such an example, a robot may be commanded to wait in place to continually observe surroundings should localisation accuracy fall below a certain threshold. When a robot is localised once more at an acceptable level, it can resume movement at a speed that maintains accuracy.

The feasibility of this option is currently being studied. Without any regard for localisation accuracy, allowing the robot to increase speed independent of schemas when not in close proximity to detected obstacles, as expected, reduced overall exploration time. This change was made to reflect the preference to quickly traverse open indoor areas prior to exploring confined regions with more caution. Having demonstrated that the decoupling of speed from the motor-schema derived heading can aid in optimising for time performance without disrupting behaviour, the next step is to incorporate mapping algorithms as proposed in Figure 5 to positively control speed and turn rate.

## 6 Conclusion

This paper presents a behaviour based approach for multi-robot exploration. Coupled with path planning to escape local minima where necessary, a simple mechanism is arrived at for the local planning of robot motion using a centralised map.

The approach proposed in this paper provides a basis for future research, including:

- Integration with SLAM algorithms providing speed control to complement desired heading as output from current approach.
- Introduction of schemas to maintain visibility of robots in beacon-based teams, resulting in a pattern of coverage or leap frogging behaviour.
- Association of frontier gains and choice of frontiers with information utility

It is expected that the results of the integration of the proposed approach with a SLAM algorithm will be available within the next few weeks.

## 7 Acknowledgements

The author wishes to acknowledge Professor Gaminu Dissanayake for valuable discussions and the financial support provided by the Centre of Autonomous Systems towards his studies.

## References

- [Arkin, 1998] Ronald C. Arkin. *Behavior-Based Robotics*. MIT Press, Cambridge, Massachusetts, 1998.
- [Balch and Arkin, 1998] Tucker Balch and Ronald C. Arkin. Behavior-Based Formation Control for Multirobot Teams. *IEEE Transactions On Robotics and Automation*, pages 926-939, December 1998.
- [Burgard *et al.*, 2000] Wolfram Burgard, Mark Moors, Dieter Fox, Reid Simmons, and Sebastian Thrun. In *Proceedings of the 2000 IEEE International Conference on Robotics and Automation*, pages 476-481, San Francisco, California, April 2000.
- [Gerkey *et al.*, 2001] Brian P. Gerkey, Richard T. Vaughan, Kasper Stoy, Andrew Howard, Maja J. Mataric, and Gaurav S. Sukhatme. Most Valuable Player: A Robot Device Server for Distributed Control. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1226-1231, Wailea, Hawaii, October 2001.
- [Grabowski and Khosla, 2001] Robert Grabowski, and Pradeep Khosla. Localization Techniques for a Team of Small Robots, In *Proceedings of the 2001 IEEE/RSJ International Conference on Intelligent Robots and Systems*, Maui, Hawaii, pages 1067-1072, October-November 2001.
- [Makarenko *et al.*, 2002] Alexei A. Makarenko, Stefan B. Williams, Frederic Bourgault, and Hugh F. Durrant-Whyte. An Experiment in Integrated Exploration. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 534-539, Lausanne, Switzerland, September-October, 2002.
- [Navarro-Serment *et al.*, 1999] Luis E. Navarro-Serment, Christiaan J. J. Paredis, and Pradeep K. Khosla. A Beacon System for the Localization of Distributed Robot Teams. In *Proceedings of the International Conference on Field and Service Robotics*, Pittsburgh, Pennsylvania, August 1999.
- [Reif and Wang, 1999] John H. Reif and Hongyan Wang. Social potential fields: A distributed behavioral control for autonomous robots. *Robotics and Autonomous Systems*, 27: 171-194, 1999.
- [Rekleitis *et al.*, 2002] Ioannis M. Rekleitis, Gregory Dudek, and Evangelos E. Miliotis. Multi-Robot Cooperative Localization: A Study of Trade-offs Between Efficiency and Accuracy. In *Proceedings of Intelligent Robots and Systems*, pages 2690-2695, Lausanne, Switzerland, September-October 2002.
- [Rekleitis *et al.*, 2001] Ioannis M. Rekleitis, Gregory Dudek, and Evangelos E. Miliotis. Multi-Robot Collaboration for Robust Exploration. In *Proceedings of IEEE International Conference on Robotics and Automation*, pages 3164-3169, San Francisco, California, 2001.
- [Xiaoping and Ko-Cheng, 1997] Yun Xiaoping and Tan Ko-Cheng. A wall-following method for escaping local minima in potential field based motion planning. In *Proceedings of 8th International Conference on Advanced Robotics*, Monterey, California, July 1997.
- [Yamauchi, 1998] Brian Yamauchi. Frontier-Based Exploration Using Multiple Robots. In *Proceedings of the Second International Conference on Autonomous Agents*, pages 47-53, Minneapolis, Minnesota, May 1998.
- [Yamauchi, 1997] Brian Yamauchi. A Frontier-Based Approach for Autonomous Exploration. In *IEEE International Symposium on Computational Intelligence in Robotics and Automation*, Monterey, California, 1997.
- [Zhang and Ghosh, 2000] Li Zhang and Bijoy K. Ghosh. Geometric Feature based 2½D Map Building and Planning With Laser, Sonar and Tactile Sensors. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 115-120, Takamatsu, Japan, October-November 2000.