Experiments in sensing and communication for robot convoy navigation

G. Dudek[1], M. Jenkin[2], E. Milios[2] and D. Wilkes[3]

[1]Centre for Intelligent Machines, McGill University,
Montreal, Quebec, Canada
[2]Department of Computer Science, York University,
North York, Ontario, Canada
[3]Ontario Hydro Technologies
Etobicoke, Ontario, Canada

December 8, 1995

Abstract

This paper deals with coordinating behaviour in a multi-autonomous robot system. When two or more autonomous robots must interact in order to accomplish some common goal, communication between the robots is essential. Different inter-robot communications strategies give rise to different overall system performance and reliability. After a brief consideration of some theoretical approaches to multiple-robot collections, we present concrete implementations of different strategies for convoy-like behaviour. The convoy system is based around two RWI B12 mobile robots and uses only passive visual sensing for inter-robot communication. The issues related to different communication strategies are considered.

1 Introduction and Motivation

Convoy behaviour is applicable in numerous application contexts. Although most existing mobile robotic systems involve a single robot operating alone in its environment, a wide range of potential applications would be natural contexts for *multiple* robots working in concert. Collectives of simple robots may be simpler in terms of individual physical design than a larger, more complex robot, and thus the ensuing system can be more economical, more scalable and less susceptible to overall failure.

Communication is a key design issue for multi-robot systems. This communication may take place directly via an explicit communication channel or indirectly through one robot sensing a change in other robots or its environment. Intra-collective communication presents difficulties in terms of collective efficiency, fault tolerance, and cost. The communication mechanism utilized by the collective is critical to the collective's practicality, efficiency and reliability. Sophisticated reliable inter-robot communication will be required for certain tasks yet this assumption can lead to reliability difficulties. If there are fixed communication topologies (e.g. [13]) or controller robots (e.g. [7]), or other fragile communication mechanisms then failure of these fixed links in the communications network will cause the entire collective to fail. In order to maximize the reliability of the collective, the communication mechanism between elements of the collective must survive the worst possible destruction of collective elements. Communication, like action, should be distributed throughout the collective.

Many different collective models have been proposed in the literature. The behaviour based control strategy put forward by Brooks [3] has been applied to collections of simple independent robots, usually for simple tasks. Other authors have also considered how a collection of simple robots can be used to solve complex problems, often in simulation only. Ueyama *et. al.* [13] propose a scheme whereby complex robots can be organized in tree-like hierarchies with communication between robots limited to the structure of the hierarchy. Hackwood and Beni [7] propose a model in which the robots are particularly simple but act under the influence of "signpost robots". These signposts can modify the internal state of the collective units as they pass by. Under the action of the signposts, the entire collective acts as a unit to carry out complex behaviors. Parker has proposed an architecture for control of multiple cooperating robots, aiming at fault tolerance and adaptive action selection for mission completion [10]. A learning version of the architecture allows the adaptation of the parameters with experience.

Mataric [8] describes experiments with a homogeneous population of actual robots acting under different communication constraints. The robots either act in ignorance of one another, informed by one another, or intelligently (cooperating) with one another. As inter-robot communication improves, more and more complex behaviors are possible. In the limit, in which all of the robots have complete communication, then the robots can be considered as appendages of a single larger robot (or robotic "intelligence"). Balch and Arkin have explicitly addressed these issues by empirically evaluating the advantages of different types of communication for a pair of simulated robots [2].

One major goal of a robotic collective is to distribute not only the sensing (and possibly actions) of the robots, but also the intelligence. What sort of processing can be accomplished by a collection of robots that cannot be accomplished by a single one? What effects do limits on communications and unit processing capabilities have on the potential actions of the collective? How do we compare the structure of various possible collectives?

The information processing ability of a collective is dependent upon a large number of factors including the number of units, their sensing abilities, their communication mechanisms, etc.[1, 9]. In order more fully to understand the properties of various collective designs, it is instructive to group collectives into classes and to determine the processing ability of each class. It may be the case that certain collective organizations have more potential processing ability than others, and that some collective organizations may be similar to existing parallel models of computation.

In this paper, we develop autonomous robotic collectives which utilize only inter-robot communication based on passive on-board visual sensing. These approaches are validated experimentally. The approaches described here are focused primarily on non-isotropic communications: i.e. robot teams where there is an explicit "leader" that performs (or receives) navigation information, and one or more "followers" that must act accordingly.

Given the variety of possible designs of groups of mobile robots, we have suggested that it is useful to organize these concepts along taxonomic axes. The objective is to both clarify the strengths, constraints and tradeoffs of various designs, and also to highlight various design alternatives. We suggest that there are several natural dimensions along which robotic swarms can be classified [5]. With respect to this taxonomy, the key characteristics of the work presented here are as follows. 1) The behaviour we have examined is applicable to collections of two or more robots. 2) These robots can communicate with one another at distance distances. 3) Since we are dealing explicitly with convoy behaviour, we consider only a strictly linear hierarchical communications topology. 4) The bandwidth for communications is low.

2 The Convoy Task

There is considerable research interest in the task of having one autonomous vehicle follow another (c.f. [4, 11]). This task is usually implemented as only a single robot following some other autonomous agent. In practice, a variety of strategies are available for implementing this type of inter-robot collaboration. In previous applications it is assumed that the target to be followed does not actively aid in the process of being followed but rather that the follower must attempt to track the leader as the leader undergoes possibly rapid random course changes. Although a leader might make rapid random course changes in a typical convoy application, another possibility would be to have the leader aid the follower. By communicating the leader's intentions to the follower, simpler, more reliable convoy behaviours are possible. There are several natural design alternatives for this communication.

Figure 1: Robots used in the experiments. The leader robot (Rosie), is on the left while the follower (Agamemnon) is on the right. The upper part of Rosie has been covered with a pattern which is used to communicate to Agamemnon.

Two-way communication The leader and the follower are in constant two-way communication. This is the strategy used in precision aircraft flying in which the leader telegraphs his intentions to the other aircraft via radio.

Explicit one-way communication
The leader signals the follower(s) through some behavior which can be sensed by the followers. In a truck convoy, for example, communication is accomplished using indicator signals for turns and brake lights for deceleration.

Completely implicit communication
The classic convoy model in which the leader ignores the follower(s).

Each of these alternatives is affected by the communication channel bandwidth and the mechanisms used to communicate between the elements of the convoy. In addition to the question of communication between the two robots, there is a separate question of how the robots can determine one another's relative positions. Even with continuous high-bandwidth communication, the absence of external sensors will lead to dead reckoning errors that result in large uncertainties in the displacement estimate between the two robots.

In this work, we examine the strategies for implementing convoy behaviour using visual sensing (i.e.

video data) and no direct communication. This sensing modality has several advantages. In particular, it is passive and hence inconspicuous, has a long range of operation, and many robots can simultaneously use it (i.e. it avoids interference problems that can occur with active sensors). Using direct communication would slightly simplify the coordination problem but we believe the key issue is one of sensing. Furthermore, direct communication (eg. by radio) can be problematic and costly for large numbers of robots.

3 Followers and leaders

In the following, we assume that the robots motion consists of piecewise linear trajectories. Although this is not necessary, it was the case in the cited experiments in order to minimize errors in dead reckoning.

The convoy problem is defined as a path following problem for two or more robots in R^n. We consider a pair of robots where a follower's path is specified relative to a leader's as a function of time t. The problem is thus defined such that a follower robot R_2 moving on a trajectory $\mathbf{r}_2(t)$ is to follow a lead robot R_1 moving on a trajectory $\mathbf{r}_1(t)$, and that

$$\mathbf{r}_2(t) = \mathbf{r}_1(t + \tau(t)) \tag{1}$$

where $\tau(t)$ is small and positive and expresses the delay between the leader and the follower. Although $\tau(t)$ may, in practice, vary over time, in the interests of notational simplicity we will assume in the rest of this paper that it is constant. In order to avoid collisions, it is necessary to include the side condition

$$\forall t \quad ||\mathbf{r}_2(t) - \mathbf{r}_1(t)|| > \epsilon \tag{2}$$

where ϵ is the sum of the radii of the two robots (assumed to be cylindrical). For a larger group of robots $R_0 ... R_N$ the straightforward generalization is $\mathbf{r}_i(t) = \mathbf{r}_{i-1}(t+\tau_i(t))$ for $i = 1..N$. A more pragmatic variation we will call the *lenient convoy problem* is obtained by replacing equation (1) by

$$||\mathbf{r}_2(t) - \mathbf{r}_1(t + \tau(t))|| < \delta \tag{3}$$

where δ is a constant.

We assume that the robots can move at a translational velocity of v and can rotate at an angular velocity of ω. Rather than account for issues of acceleration and other such complications, we assume that each time the robot stops is remains where it is for a minimum time γ. The time required to acquire and process a video frame is given by T.

3.1 Simple following without communication

The simplest and most efficient strategy for the lead robot is to perform no explicit communication and simply perform its actions in the most efficient manner possible. In this context, the follower(s) must observe the leader without explicit "assistance".

The simplest possible strategy is for the follower robot to continuously move towards the leader while the leader performs its actions oblivious to the follower's behaviour. In general, this leads to a violation of condition 2 and a collision since the follower may continue to head for the leader even when it is stationary or even when it heads in the direction of the follower itself. A further complication is that tracking the leader while the follower is, itself, moving presupposes additional abilities on the part of the follower and precludes using the camera for collision avoidance. To avoid the risk of collision while maintaining assured linkage of the convoy involves an element of cooperation between the leader and follower.

Convoy behaviour using video tracking without communication from the leaders entails the following behaviour on the part of the follower:

1. detecting (and observing) the lead robot at $\mathbf{p} = \mathbf{r}_1(t_i)$,
2. waiting for the leader to start moving,
3. moving to a safe position *en route* to $\mathbf{p}$ if the leader is moving towards the follower,
4. moving to the leader's former position such that $\mathbf{r}_2(t_i + \tau) = \mathbf{p}$, otherwise.

To carry this out, the leader, at each step, cannot move until sufficient time has elapsed that its follower is assured of having completed steps (1).

Two fundamentally different strategies are available to implement even this simple behaviour:

1a. the leader can move in small steps to assure its continued detectability over time (using tracking), or

1b. the leader can move in arbitrary (linear) steps and the follower can perform an explicit search for the leader (for example by panning its cameras about the scene) as part of step (1).

It is not only time consuming, but also error prone to have to explicitly detect (and perhaps search for) the lead robot, especially when the distance between robots is not well constrained (for example, if the leader takes large steps). For each step a delay of $O(\omega T)$ to acquire data from different possible orientations is required to identify the position of the lead robot. The probability of error after κ steps is $(1 - e)^\kappa$ when the probability of mis-identification on any step is e. Since this has to be repeated for each step the overall behaviour is likely to be limited by this process and the likelihood of failure may be substantial (especially if there are several robots operating at once). A panoramic sensor such as [14] would permit a more efficient solution to this problem.

In contrast, behavior (1a) has a much lower probability of failure since tracking can be performed very robustly however the rate at which the leader can move becomes highly constrained. Under the assumption that the follower is always exactly one step behind, the leader can move at an optimal rate determined by the field of view of the follower's camera, the degree of clutter in the environment, the image processing rate, and the time needed to assure that the follower has arrived at the most recent position. For a camera with a field of view of 2θ degrees in an obstacle-free convex environment (see Fig. 2), this implies that the leader may be able to move in steps no larger than of $||\mathbf{r}_1(t) - \mathbf{r}_2(t)|| \sin\theta$. At each step it must wait for the follower to move from its prior position (the leader's own position two steps ago), to the leader's most recent position before starting to move.

At each step the robot must wait for the follower to move from its prior position (the leader's own position two steps ago) $r_1(t-2\tau)$, to the leader's most recent position $r_1(t)$ by way of the leader's position one step ago. If the robot moves with velocity v, then the time for the robot to make this motion is bounded by $||r_1(t) - r_1(t-\tau|| \sin\theta/v + ||r_1(t-\tau) - r_1(t-2\tau)||/v$. Including the overhead of acquiring the data and allocating the acceleration/deceleration costs bounds the velocity of the leader by

$$s_{net} =$$

$$\frac{||\mathbf{r}_1(t) - \mathbf{r}_1^{\tau}|| \sin\theta}{T + 2\gamma + \frac{1}{v}(||\mathbf{r}_1(t) - \mathbf{r}_1^{\tau}|| \sin\theta + ||\mathbf{r}_1^{\tau} - \mathbf{r}_1^{2\tau}||)}$$

where $\mathbf{r}_1^{\tau}$ denotes $\mathbf{r}_1(t-\tau)$ and $\mathbf{r}_1^{2\tau}$ denotes $\mathbf{r}_1(t-2\tau)$.

If the leader moves no faster than this then the total convoy speed will be limited by this value as well. Note that the numerator of this expression is a pessimistic value for the longest distance that the robot can move and not leave the field of view of a follower. In the worst case, this implies a reduction in the maximum net speed for the convoy to value no higher than one-half of the (lowest) individual speed: in practice, it may be even lower than this. Under certain circumstances the leader can move considerably faster than s_{net} and still ensure that the follower remains in the field of view; this depends of the direction of travel relative to the line of sight of the follower. For example, suppose that the leader is moving in a straight line and knows that the follower is directly behind. Then the leader can move considerably faster than s_{net} directly away from the follower and still be sure that it will be in the follower's field of view.

3.2 Following with explicit hints

A key difficulty in the strategy described above is that the following robot is unable to infer the target position, or even the direction, of the leader and hence it must observe it continuously. The performance of the convoy can be enhanced if the leader is able to communicate its intended orientation for the next motion before it starts moving. Using this cue, even if the follower does not track the leader when it gets to the leader's prior position it needs only to point its camera in the correct direction of reacquire the leader. This can be especially important in a cluttered environment when the use of continuous tracking would impose a severe constraint on the maximum velocity due to the requirement for a small step size.

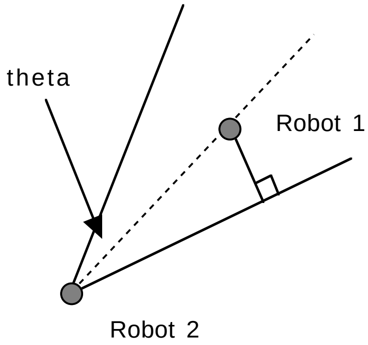

Figure 2: Schematic illustrating the view from the follower to the leader robot.

Using this added information, the leader and follower robots can both be in motion simultaneously. The leader need only pause occasionally to allow the follower to observe the new heading before starting the next step. This allows a substantially higher net speed roughly proportional to the variation in length and orientation between successive steps.

4 Experimental verification

4.1 Hardware

Two RWI (Real World Interface) B12 mobile robots were used in the experiments reported here (see Figure 1). These robots are roughly 12" in diameter. Each robot is equipped with 12 Polaroid sonar sensors (although these were not used in the experiments report here), and communicates to off board processing through serial links. Agamemnon[1], the robot on the right, is equipped with a video camera for sensing and utilizes a wireless serial link for offboard communication. Rosie[2], the robot on the left, has been augmented with a textured pattern region with which it communicates with Agamemnon. Rosie uses a tethered serial link for communication to offboard computation.

4.2 Sensing and Communication

Experiments have been conducted using several communication strategies for the Convoy task. For the three strategies described here, the follower and leader communicate with each other only through follower's passive visual sensing of the leader's pose in the follower's own coordinate system. The first

---

[1] Agamemnon normally resides at the Centre for Intelligent Machines, McGill University.

[2] Named for the robotic maid in the Jetson's, Rosie normally resides at the Vision, Graphics and Robotics Laboratory, York University.

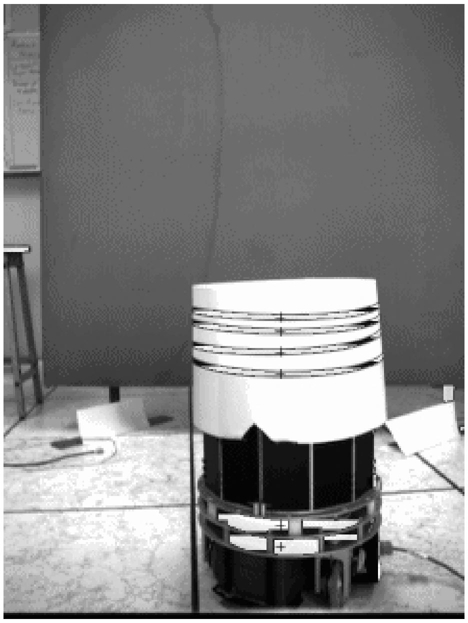

Figure 3: Agamemnon's view of Rosie. This image shows the output of the image processing step. Horizontal lines have been coloured white, and small crosses have been placed in candidate horizontal stripes. Given the height of the top stripe, and the relative position of the third stripe, Agamemnon can compute Rosie's relative pose.

set of experiments involved a convoy strategy without hints, as described above.

Communication via hints was achieved by encoding the leader robot with a geometric pattern that allowed it not only to be robustly detected (assuming a relatively simple environment), but which also allowed its relative pose (distance and orientation) to be rapidly computed. The pattern is a cylindrical set of almost-horizontal lines on the leader's turret (the robot Rosie, see Figures 1 and 3) that encodes pose as an essentially analogue value.

An alternative approach that has been used to simply encode orientation of an object being imaged uses a binary pattern that is wrapped around a target object [12]. Such an approach, however, depends critically on an image with sufficient resolution to determine the individual bits of the binary pattern encoding orientation values.

In the mobile robotics context, it is important that the pattern (to be applied to the robot) be of small size and yet encode pose accurately. In contrast to the binary encoding alternative, the encoding pattern we have used describes distance as well as orientation and degrades gracefully as resolution decreases (for example as the distance from the camera to the pattern increases).

This pattern is comprised of a set of black lines on a white background. Three of these are parallel (and roughly horizontal) and are used to encode distance as well as to provide robustness. By locating the four black stripes on Rosie (as shown in Agamemnon's view of Rosie in Figure 3), Agamemnon can compute the distance between the two robots by locating the pattern in the image and computing the distance between the furthest pair of lines in the image. The additional horizontal line and its position is used to verify that the correct pattern has been located (as opposed to background noise). The 3rd stripe (counting from the top) in the pattern wraps helically down the robot (diagonally when the pattern is laid flat). As a consequence, its vertical position with respect to the other lines directly encodes the orientation of the robot.

Very simple image processing steps are applied (by the follower robot) to compute the pose in Agamemnon's view of Rosie. The estimated orientation is accurate to roughly five degrees at one meter distance while the distance estimate has an accuracy on the order of centimeters.

4.3 Herding

In addition to the techniques described in the previous section, a simple herding behaviour was also examined whereby the leader robot manipulated the position of the follower by exploiting constraint (2): essentially "pushing" the "follower" robot away by moving towards it (without any physical contact).

For larger collections of robots (examined in simulation only) this generalizes to a behaviour where the lead robot adjusts the position(s) of one or more followers by essentially pushing them (analogous to the manner in which dogs can be used to herd sheep or cattle). In our real experiments, we have implemented these behaviours with two robots (as above). In the herding experiment, Rosie moves autonomously under external control, while Agamemnon moves to center Rosie horizontally and maintain Rosie at a particular distance. This behavior was implemented on Agamemnon by defining an energy function that has a minimum when Rosie is centred and at an appropriate distance. This function allowed the robots to repel one another when close together and to attract when widely separated. The specific form of this artificial force function is

$$f_R(r) = \begin{cases} -K_1/d^2(r-d)^2 & 0 \leq r < d \\ K_2/d^2(r-d)^2 & d \leq r < 2d \\ mr + [K_2 - m(2d)] & r \geq 2d \end{cases} \qquad (4)$$

where r is the distance between two robots, K_1, K_2, d and m are constants.

Agamemnon computes the distance and defines centrally through its passive visual sensor as described above. Once the current energy value is available, Agamemnon moves to reduce the energy term.

In the herding experiments we found that Rosie could successfully herd Agamemnon around our lab environment, provided that Rosie made sufficiently small steps so that it did not move out of Agamemnon's field of view. A panoramic sensor such as [14] would be a useful sensor for these types of situations.

4.4 Convoy

In a first implementation, the Herding algorithm described above was used. Rosie moved and Agamemnon followed, trying to maintain a preferred distance away from Rosie. This technique was successful, but as with the Herding experiment, small movements of Rosie were required in order to keep Rosie within Agamemnon's field of view. The herding-like convoy algorithm operates using completely implicit communication. The leader completely ignores the follower and it is up to the follower to correctly track the leaders position.

In the more complete implementation of the convoy task, the follower robot (Agamemnon) followed the lead robot (Rosie) while the leader moved about the environment in a manually selected trajectory. Using the approach without "hints" the rate of progress was severely constrained not only by the speed threshold defined in the previous section, but also by the need to avoid occlusion by obstacles in the environment.

In a second experiment, we implemented explicit one-way communication in which the leader (Rosie) "telegraphed" its movements to Agamemnon using the visual pattern encoding. This telegraphing was accomplished by having the lead robot rotate to the direction in which it was going to move and then to pause for a moment so that the follower could observe the intended heading (as well as distance). The leader robot would then begin moving to the new position while the follower robot would begin moving to the leader's former position. This rotation, pause and translation step signaled to the follower the bearing at which the leader's new position would lie. When the follower arrived at the leader's old position it then rotates to look in the direction that the leader had moved in order to re-acquire it.

Using this explicit one-way communication technique, the two robots could move much farther and faster using the simpler methods described above. Empirical estimates put the difference in net speed at roughly a factor of four, but this is highly dependent of environmental factors. This mechanism did require considerably more cooperation between the two robots but it did provide a considerable improvement in performance over simpler convoy algorithms.

5 Discussion

This paper presents a number approaches to the use of multiple mobile robots to effect group behaviour, in particular coordinated "convoy" motion. Without advance notice of what the leader is going to do, the convoy behaviour becomes extremely inefficient. It is interesting to note that this is essentially consistent with human convoy behaviour: consider how difficult it is to efficiently follow another driver who does not use his or her signal lights or brake lights.

In terms of the collective taxonomy defined in prior work [5], this approach describes a system that is described as

Size: Pair

Range: COM-NEAR: local communication

Topology: TOP-BROAD: broadcast messaging

Bandwidth: BAND-MOTION: expensive communications with low bandwidth

Reconfigurability: ARR-STATIC: static hierarchy of robots.

A helical stripe pattern was developed to encode relative pose and efficient leader-follower behaviour was achieved. This led to a substantial improvement in the net rate of motion that a group or robots could achieve. These approaches were examined experimentally using two RWI-B12 mobile robots. The results verified the desirability of visually encoding additional information.

Acknowledgements The authors gratefully acknowledge the financial support of the Natural Sciences and Engineering Research Council.

References

[1] R. C. Arkin and J. D. Hobbs. Dimensions of communication and social organization in multi-agent robotics systems. In SAB 92, 1992.

[2] Tucker Balch and Ronald C. Arkin. Communication in reactive multiagent robotic systems. Autonomous Robots, 1(1):27–52, 1994.

[3] R. A. Brooks. Intelligence without reason. Technical Report AI Memo No. 1293, MIT, 1991.

[4] E. Dickmanns and A. Zapp. Autonomous high speed road vehicle guidance by computer vision. In Proc. 10th World Congress on Automatic Control, volume 4, Munich, Germany, 1987.

[5] G. Dudek, M. Jenkin, E. Milios, and D. Wilkes. A taxonomy for swarm robotics. In IEEE/RSJ IROS, pages 441–447, Yokohama, Japan, 1993.

[6] D. W. Gage. Command control for many-robot systems. In Proc. AUVS-92, Technical symposium of the association for unmanned vehicle systems, pages 22–24, Huntsville, AL, 1992.

[7] S. Hackwood and G. Beni. Self-organization of sensors for swarm intelligence. In Proc. Int. Conf. on Robotics and Automation, pages 819–829, 1992.

[8] M. Mataric. Minimizing complexity in controlling a mobile robot population. In Proc. IEEE Int. Conf. on Robotics and Automation, pages 830–835, 1992.

[9] E. Nitz and R. Arkin. Communication of behavioral state in multi-agent retrieval tasks. In Int. Conf. Cons. R. and A., Hawaii, 1993. (to appear).

[10] L. Parker. Alliance: An architecture for fault tolerant cooperative control of heterogeneous mobile robots. In Proc. IROS 1994, Munich, Germany, Sept. 1994.

[11] R. Skuthankar. Raccoon: A real-time autonomous car chaser operating optimally at hight. In Intelligent Vehicles '93, pages 37–42, Tokyo, Japan, 1993.

[12] R. Szeliski. Uncertainty models for 2-1/2d and 3-d surfaces. In L. Harris M. Jenkin, editor, Spatial Vision in Humans and Robots, pages 229–247. Cambridge University Press, New York, NY, 1993.

[13] T. Ueyama, T. Fukuda, and F. Arai. Configuration of communication structure for distributed intelligent robot system. In Proc. IEEE Int. Conf. on Robotics and Automation, pages 807–812, 1992.

[14] Y. Yagi, S. Kawato, and S. Tsuji. Collision avoidance using omni-directional image sensor (COPIS). In IEEE Robotics and Automation, pages 910–915, 1991.