A Scalable, On-Board Localisation and Communication System for Indoor Multi-Robot Experiments

Ian Kelly[1*] and Alcherio Martinoli[2*†]

[1]Robotics Research Laboratory, University of Southern California, Los Angeles, CA 90089, USA

[2]Swarm-Intelligent Systems Research Group, Nonlinear Systems Laboratory, EPFL, CH-1015 Lausanne, Switzerland

Keywords: Autonomous Robots, Smart Sensors, Local Positioning, Local Communication, Flocking.

Abstract

We describe the first prototypes of an inter-robot infrared localisation and communication system that exploits standard of-the-shelf RF components and uses time division multiplexing of a single carrier frequency. The system detects the relative positions (both range and bearing) of autonomous mobile robots with an update rate of up to 20Hz, a range of up to three meters and an accuracy of 40cm for range and 45 degrees for bearing. In addition, each robot can send at least one byte of data to all the other robots within range per update cycle. Flocking on a group of eight robots is used as a non-trivial real-world test of this system. We conclude the paper by discussing advantages, limitations, and future improvements of the system.

1 Introduction

For many real-world tasks the lack of reliability of any single robot is unacceptable [Unsal94], therefore recently considerable interest has emerged in systems with multiple co-operating autonomous mobile robots. Typical tasks that have been proposed in the literature for such systems are:

Cooperative transportation of heavy objects [Bay95, Kube00],
Distributed search [Genovese92, Hayes02a],
Distributed coverage, mapping, and exploration [Levy92, Howard02].

The application fields vary quite a lot. They include, but are not limited to: extraterrestrial exploration, lawn mowing, humanitarian demining, building inspection in catastrophic scenarios and military operations, wireless network deployment and maintenance in catastrophic and battlefield scenarios, and localisation of pollutant sources.

In order for multiple robots to co-operate with each other, often a method of locating teammates and some form of communication is required. Communication can be achieved through the

---

[*] This work has been carried out while both authors were at the California Institute of Technology, Pasadena, CA 91125, U.S.A. It has been completely supported by the Caltech Center for Neuromorphic Systems Engineering under the American NSF Cooperative Agreement EEC-9402726.

[†] Correspondence author. A. M. is currently sponsored by a Swiss NSF professorship.

environment [Bonabeau99] or directly by a dedicated system on the robots. Here we use the direct approach and have designed and constructed an on-board infrared communication system, which in addition gives the relative ranges and bearings to other local robots.

As a first non-trivial test-bed experiment for our system, we have implemented a simple flocking algorithm using a group of eight robots. Flocking is one of a number of useful basic behaviors that can be utilised in combination with other behaviours for solving many higher level tasks such as those previously mentioned (for instance, for moving a fleet of robots from A to B when there is not specific interest in searching the zone between A and B). Furthermore, it can be considered a primitive form of formation movement in which the overall cohesiveness is more important than the relative geometric placement of the units [Fredslund02]. For these reasons, flocking has recently stimulated considerable engineering development with artificial systems, both in simulation [Reynolds87], [Ota93], [Hodgins94], [Hayes02] and real robots in two-dimensions [Mataric95], [Kelly96], [Hayes02], and more recently in three-dimensions [Welsby01].

2 Localisation and Communication System

This inter-robot localisation and communication system uses standard of-the-shelf RF components and design principles to reduce both cost and development time - with the exception that infrared light is used as the carrier instead of a radio signal. Infrared light was chosen for many reasons: small highly directional receivers are readily available, low power requirements, and inexpensive of-the-shelf transducers.

2.1 System Overview

This system is controlled by its own dedicated on-board PIC microcontroller, thus freeing the robot's main processor for higher-level tasks. Each robot is equipped with a ring of twelve Infra-Red (IR) Light Emitting Diodes (LEDs) that emit a carrier (signal) at a set frequency. The LEDs are split into four independently controllable regions (front, back, left, and right) of approximately 90 degrees each, and can give a full 360 degrees of coverage. The transmission power can be set to one of three different levels allowing for different ranges of up to 3.1m. Each robot also has four photodiode receivers spaced at 90 degrees from each other that are connected to a tunable superheteroding FM narrowband receiver via a multiplexer. The signal strengths received at the four receivers allow the relative distance and angle to other robots to be calculated.

Data are transceived by frequency modulating the carriers. The system is designed to operate on a single carrier frequency using time-division multiplexing with dynamic selection of transmission time-slots. However, since both the receivers and transmitters are re-tunable, in software, groups of robots could operate at different frequencies, allowing experiments to be conducted using two or more unique groups of robots within the same area.

Figure 1 shows a block diagram of the system and Figure 2 shows one of our robots (Moorebot, see also [Winfield00] for more details on the basic robotic platform) with the localisation and communication board mounted on top.

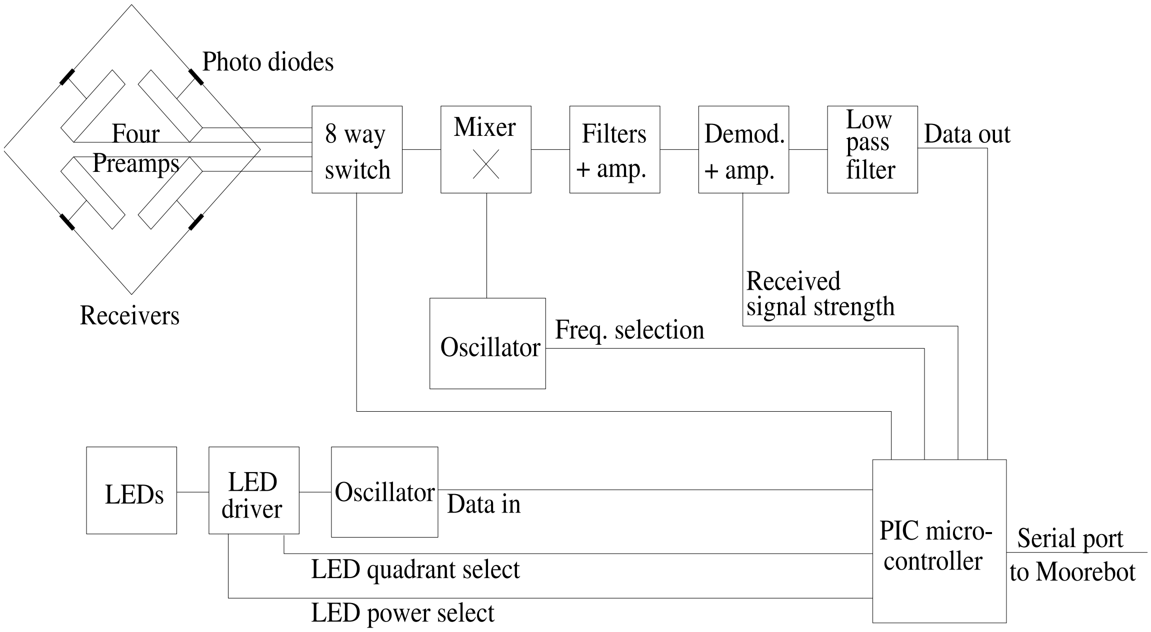

Figure1: Block diagram of the communication and localisation system.

Figure 2: Close-up of a Mooreboot with the localisation and communication board (board on top of the PC104 stack on the back of the robot).

2.2 The Hardware

To reduce development time we have based our system on one previously developed by [Kelly96] and have considered the systems requirements from work carried out by [Hayes02] using an indoor localisation system based on a combination of an overhead camera and a global radio LAN. The hardware system used a standard "of-the-shelf" narrow-band FM demodulator. This system was originally designed to use frequency division multiplexing and so used heterodyning. In heterodyning, the input signal is mixed (multiplied) with a local oscillator. This local oscillator works at a set frequency (e.g. 455kHz), higher or lower than the desired frequency. This mixing produces images of the desired frequency at both the sum and difference of the local oscillator and desired frequency. A fixed high quality filter is used to select either the sum or difference, thus reducing the size of the receiver. Tuning the receiver to another frequency is accomplished by simply changing the frequency of the local oscillator. Since the receiver can operate at different frequencies two or more separate groups of robots could operate within the area, with each transmitting at a different frequency.

It has already been shown [Kelly96] that at low frequencies there are high levels of noise due to the high efficiency fluorescent tubes used in many laboratories. Therefore the communication system needs to operate at higher frequencies than the majority of this noise, to obtain a wider bandwidth for channel modulation and potentially for a multi-channel system. Unfortunately, inexpensive infrared LEDs work most efficiently at frequencies of less than 10kHz, which is within the band of high ambient noise and would offer a low communication bandwidth. Communication grade LEDs (such as Sweetspot^{TM}) and laser type devices will both operate at much higher frequencies than inexpensive infrared LEDs (about 60MHz for communication grade LEDs and 50MHz plus for laser devices). However, lasers are undesirable for safety reasons, and both Sweetspot^{TM} ($25 each) and laser ($100 plus optics) devices are very expensive compared to normal LEDs ($0.50). Therefore inexpensive LEDs were chosen, using frequencies starting from 220kHz, which have relatively low back ground noise levels. To overcome the loss of power output from the LEDs at these higher frequencies we use a ring of twelve LEDs (each having a half power angle of 60 degrees) arranged 30 degrees apart in one plane, such that the light from adjacent LEDs overlaps. Thus the light intensity stays approximately constant with respect to angle. To allow greater flexibility the LEDs are arranged into four independently controllable banks: front-left, front-right, back-left, and back-right. Each bank can be switched on or off as required, in software. The overall transmission power can be controlled by selecting, again in software, which source resistor is used to limit the current flow though the LEDs, thus effectively controlling the transmission range.

2.3 Sharing the Available Bandwidth

Transmitting information from many robots onto a single carrier frequency is achieved by time division multiplexing. An alternative method would be to use different frequencies for each robot, as used in [Kelly96]. However, this method, in addition to not be scalable with increasing numbers of robots for a finite bandwidth, generates “beating” problems between the fundamental frequencies and harmonics of the signals and fluorescent tubes (in other words the different signals would mix/multiply with each other producing interfering mirror frequencies). This effectively creates high levels of noise and thus rules out the use of many channels. With increasing numbers of robots the problem worsens. Moreover, since inexpensive infrared LEDs output power starts falling of at about 10kHz, robots transmitting at different frequencies also emit at different light amplitude levels, decreasing the accuracy of range measurements.

By using a single frequency and time division multiplexing these problems are overcome. First, since all the robots transmit at the same frequency they transmit the same output power. Second, with only one robot transmitting at any given time, there are no problems with “beating” which leads to more accurate positional fixes. Third, there is no need to add additional channels for extra robots which makes the system scalable to increasing numbers of robots, assuming that the communication range is not global. However, increasing the number of robots within communications range of each other increases the chance of packet collisions and it has been shown that only about 30% of the theoretical data capacity of computer networks, such as Ethernet, can actually be utilised under high demand [Guy, 1992]. With a non-global system the number of robots that can locate and communicate with each other simultaneously is set by the range, the communication bandwidth, the positioning update rate of the system, and by the information processing capacity of the individual robots. The overall performance in term of processing speed depends on the density of robots and on the slowest component in the sensing-communicating-computing chain.

Furthermore, in our system robots do not have unique identities, unless an identity field is added to the transmitted data packet. Transmitting identity information would make the system less scalable, and further limit the number of robots that can share the channel within a local region, since each one needs more bandwidth (to transmit its identity). Without identity transmissions a method of eliminating robots positions being included more that once during any update-period is required (see subsection 2.6 for more details).

For all the above reasons, especially scalability, we have chosen to use time division multiplexing, and in particular our system uses carrier sense multiple access (CSMA) with a random back-off time if the carrier is busy when a transmission is attempted.

2.4 Range and Bearing

The range to another robot is determined from the strength of its received signal. This Received Signal Strength Indication (RSSI) is directly obtainable from the demodulator IC. By placing four photo-detectors arranged 90 degrees apart in one plane, as shown in Figure 3, the received signal strengths of all four receivers can be compared to give the direction and range of the other robots. In order for this direction detection scheme to work, the photo-detectors have to have receptive half-angles of greater than 90 degrees. This arrangement of LEDs and photodiodes also allows transmissions to be received regardless of the angles of rotation of the robots. LEDs and photo-diodes (with visible light filters) to the above specifications are readily available for 950nm infrared light.

The system detects the relative positions (both range and bearing) of autonomous mobile robots with an update rate of up to 20Hz, a range of three meters and an accuracy of 40cm and 45 degrees for bearing (see section 3.1) using an 8-bit analogue-to-digital converter. The accuracy of the system in range is limited by three factors: first, the nonlinear response of the RSSI vs. distance (while a change in the LSB at 30 cm translates in a mean change of 3.25 cm in distance, 1 LSB represents on average 15 cm at 3 m); second, small heterogeneities among optoelectronic components (and therefore among different sectors and robots) prevent the achievement of increased accuracy of the measurement of the relative range; third, possible misalignments between receptor and emitter may change the intensity of the received signal for the same position of the transmitting robot. The resolution of bearing measurements is limited by the number of sectors (or photodiodes) used to cover a 360 degrees horizon and again small differences among optoelectronic components which prevent a reliable interpolation of the signal between two neighboring sectors.

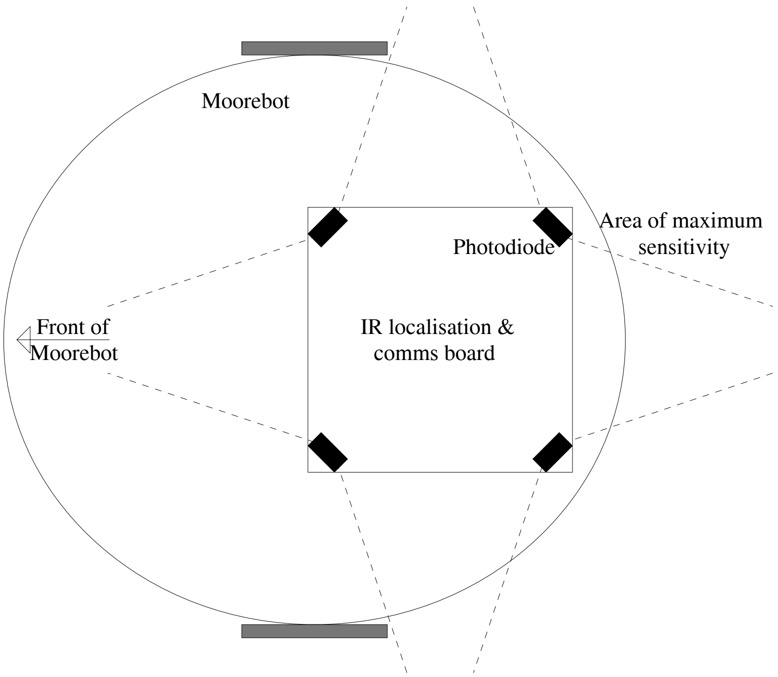

Figure 3: Photo detector arrangement.

2.5 Control Firmware

The firmware is written in C running under interrupt on a PIC microcontroller. The main program loop is shown in Figure 4. The theoretical maximum update rate without the system suffering from degradation through packet collisions is given by:

$$\text{Max. Update rate (Hz)} \approx \frac{0.3}{\text{Transmission Time} * \text{Number of neighboring Robots}}$$

The "0.3" factor represents the achievable 30% of the theoretical data capacity that comes from data collisions [Guy92] - this allows us to avoid a high number of communication collisions by keeping the transmission time sparse. Furthermore, in our case the transmission time is 17ms (for localisation and transmission of data). The 17 ms length of each robot's transmission is the results of an engineering design trade-off required to obtain reliable positioning information and data transmission as quickly as possible.

Finally, the theoretical maximal update rate decreases as a function of the number of neighboring robots, a number which is usually less than the total number of robots used in the experiment and is strongly influenced by experimental constraints (collective pattern of motion required by a specific task, range of detection/communication, etc.). For instance, if in an experiment it is known that at most one robot will be within the detection/communication range, our theoretical maximal update rate will be 17.65 Hz. We could therefore choose an update frequency not too far away from this value (for instance, 20 Hz) in order to avoid that the system performance is too much degraded due to signal collision. If instead, in our experiment we can have situations where a single robot is surrounded by nine other teammates and we would like that it can communicate and locate all of them all the time, we would select an update rate of 2Hz since the theoretical maximal update rate is about 1.96 Hz.

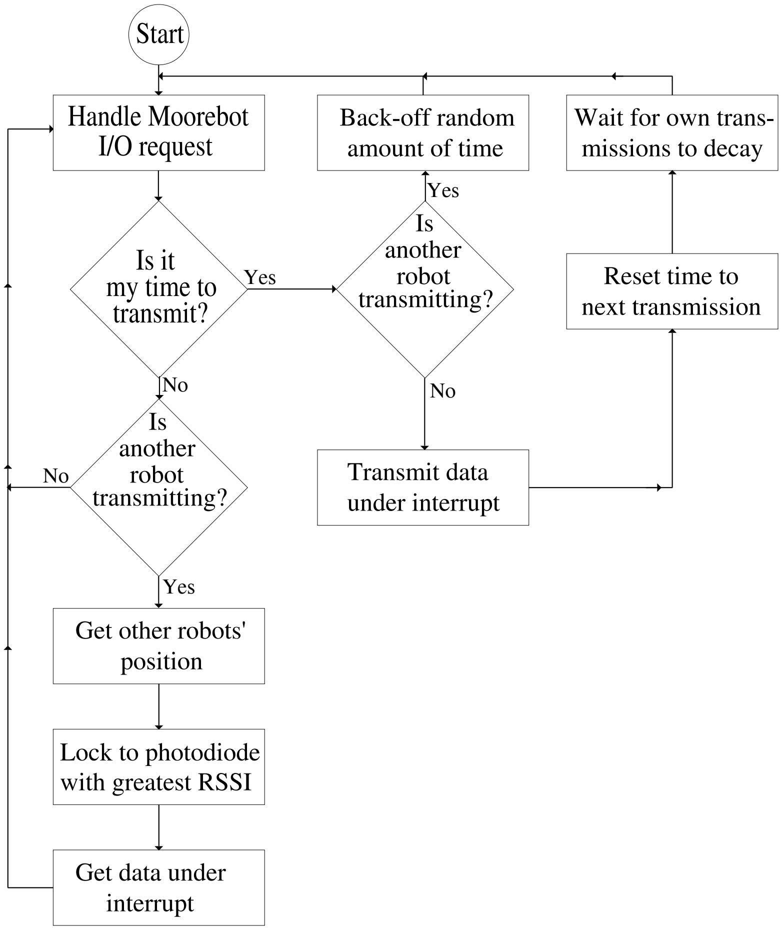

Figure 4: Flow chart of main program loop.

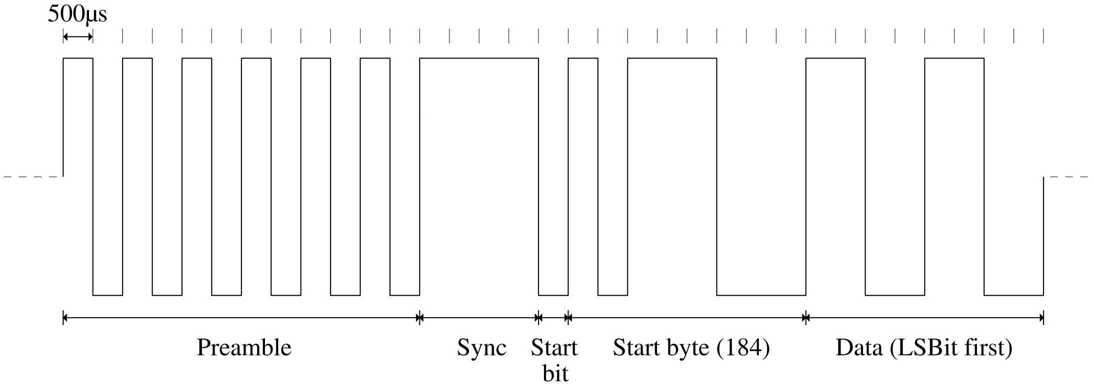

Figure 5: Transmission waveform when sending the byte 51.

An example of the waveform used in our system is shown in Figure 5. The preamble serves two purposes: first, to setup the FM demodulator's voltage biases on the receiving robots, and second, to allow time for the receiving robots to work out and lock their receiver to the photodiode which points most directly towards the transmitting robot. It is therefore important that the length of the transmission is long enough to lock robots' receivers but also short enough so that the robots cannot move too far away within a waveform period. When reading the signal strengths at the four photodiodes the system takes an average of eight consecutive readings of each photodiode to reduce the effects of noise. With this averaging process it takes 4ms to read all four photodiodes, and tune the receiver to the correct photodiode. To this 4ms we have to add a further 10.5ms to read the start bit/byte and the data. During the total transmission time of 17ms (which includes extra time for the receiving robots to detect the transmission) a Moorebot can move a distance of 1.7cm assuming that they travel at a maximum speed of $1ms^{-1}$. This additional source of inaccuracy due to movement within a measurement period is below the 1-bit resolution obtainable in the best case with the range system (3.25cm) and is therefore negligible.

Data transfer is achieved under interrupt by frequency modulation at a rate of 2000 baud per channel. The error-free range of the communication system for data transfer is 6m under the worst possible alignment conditions (see section 3.2). This distance is much larger than 3.07m (3.05m maximum robot detection range plus 1.7cm possible distance of movement within transmission time) thus allowing for a longer data packet if required, but of course at expense to the maximum update rate.

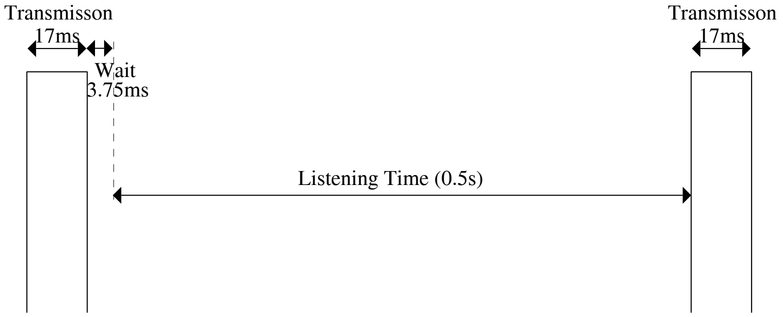

Figure 6: Overall timing.

The maximum data transmission rate is limited by the pass-width of the IF filters used. After a robot completes its transmission it waits for 3.75ms before searching for transmissions from other robots. This is to allow the output from its own transmitter to fully decay, so that it does not try to receive its own transmissions as another robot. The presence of other robots transmitting is detected by comparing the received signal strength (RSSI) at the four photodiodes to pre-established, hand-coded thresholds. The 6ms of preamble is, on its own, enough transmission time for the other robots to obtain the range and bearing. Therefore, if data is not required it would be possible to increase the update rate of localisation by almost three times by not transmitting the data byte, the start byte and the start bit. Given the local nature of this system this method of comparing the RSSI to a threshold can lead to two robots transmitting simultaneously to a robot that is in-between them. However, since we lock onto the photodiode with the strongest signal strength, one of the two transmitting robots will be correctly detected and the other will get ignored. Figure 6 shows the overall transmission and reception timings.

2.6 Identifying the Neighboring Robots

Each transmitting robot does not have a unique identity and is only distinguished by the fact it transmitted. A non-computationally intensive method is therefore required in order to differentiate transmitting teammates based on the corresponding time period the packet was received. Whenever another robot's transmissions are received, its position (and the byte of data received) are stored sequentially in five different arrays, one for each photodiode sector plus one for the data. These arrays have a size limited to the maximum number of robots used in the experiment plus one. This is a conservative approach although in reality we might know in advance that the maximal number of neighboring robots (overall or showing up in a specific sector) is much smaller than the total number of robots used in the experiment. Once the limits of these arrays are reached further data are stored starting from the beginning of the arrays thus overwriting older robot locations. The last updated positions in each of the arrays are known by both the IR systems firmware and by the Moorebot.

Using arrays of fixed length and overwriting them whenever they are full would not be enough to discriminate among different neighboring robots. We could certainly use a priori information such as robot kinematic/dynamic properties in order to correlate successive relative positioning data and therefore uniquely identify their position, a process that is often implemented in position tracking algorithms using vision (e.g., an approach based on Kalman filters). However, a much simpler and less computationally expensive method exists if we exploit the fact that each robot is also part of the local network: it is actively transmitting, has to negotiate its time slot at the beginning and synchronize with the neighboring robots, and knows its own transmission time stamps. Indeed, robots store a unique label when they transmit by writing the number “255” at the next available position in the arrays. With this information the Moorebot can easily figure out how many other robots are currently within range, their positions and the byte they transmitted as follows. It suffices to identify two “255” markers defining the local network “period” and assign incrementally, starting from the first marker, the received position data to different robots up to the second marker.

If the number of network nodes is static for a while this simple algorithm will keep track of the different robots in the neighborhood. If one of the robots leaves the local network or a new robot joins it, there will be first again a renegotiation of times slots in order to avoid collisions and again a certain rhythm and series of time stamps corresponding to different robots. Notice that, if no markers or only one marker is found, then all of the positions in the array represent robot locations and all the robots are within reception range of each other. Of course this assumes that all the robots are set to transmit at the same rate and that one or more robots are not forever forced to “back-off” their transmission time into the next “transmission-period”.

3 System Testing

In this section, we present the results of some preliminary tests we have carried out using our ten prototypes mounted on ten different robots.

3.1 Range and Bearing

Tests on the accuracy and repeatability of the range and bearing system are presented in Figures 7 to 9. All the tests were carried out with the transmitting robot moving towards the stationary receiving robot at a speed of 6.6cm/s, starting from a distance of 3.5m; for every 3.1cm of movement the signal strength at the receiving robot was stored. In all the trials the transmitting robot transmitted at full power. The back-left photodiode was used for reception in all the tests.

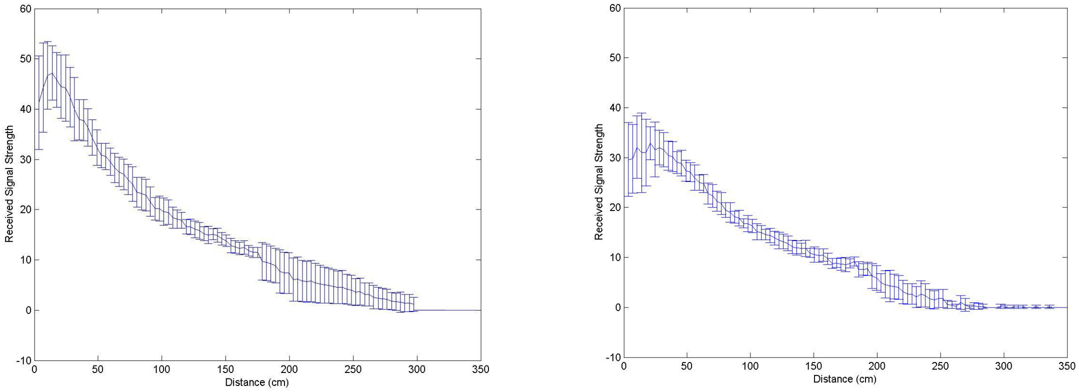

Figure 7: Average response and standard deviations (over 10 trials) of the received signal strength vs. the real distance for an arbitrary pair of robots. *Left*: under the best conditions (receiver and transmitter aligned). *Right*: under the worst conditions (receiver and transmitter misaligned).

Figures 7 shows the differences in received signal strength over distance, measured from the transmitter to the receiver, with one robot receiving a second robot's transmission. In both these tests the same two robots were used. In Figure 7, *left* the tests were done under the "best" possible conditions, where the signal strength is the greatest – that is with a transmitting LED facing straight towards the receiving photodiode. In Figure 7, *right* the same tests were carried out but this time under the "worst" possible conditions, where the signal strength is the lowest – that is with the transmitting robot facing the receiving robot at an angle half way between the nearest two LEDs, likewise with the receiving robot facing the transmitting robot at an angle half way between the nearest two photodiodes. Both the curves in these graphs are nonlinear, with a decrease in received signal strength at distances less than 10 cm due to the lense optics on the LEDs.

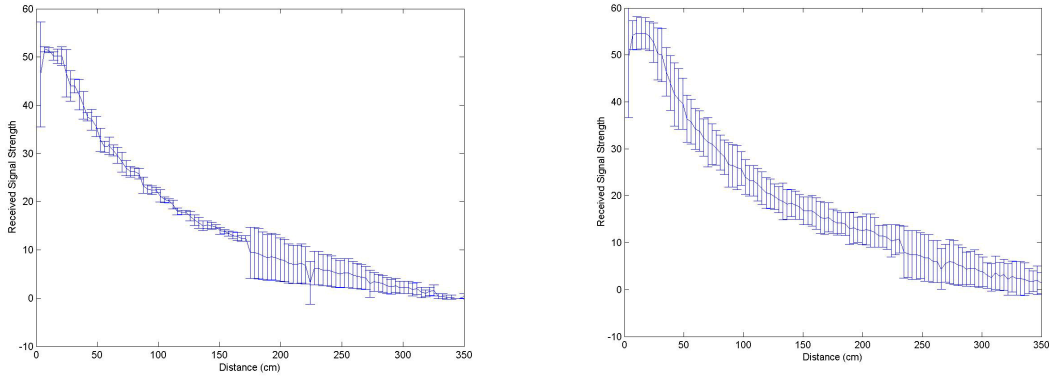

Figure 8: Average response and standard deviations (over 10 trials) of the received signal strength vs. the real distance for an arbitrary pair of robots under the best conditions (receiver and transmitter aligned). *Left*: using the same receiving robot and different transmitting robots. *Right*: using the same transmitting robot and different receiving robots.

Figures 8 shows repeats of the experiments depicted in Figure 7 but with different transmitting robots (*left*) and different receiving robots (*right*). As expected these are all very similar to Figure 7, *left*. The standard deviations in the above graphs show that there are relatively small differences between the received signal strengths of different photodiodes (on the same robot or on different robots). They also show that different robots transmit slightly different transmission powers.

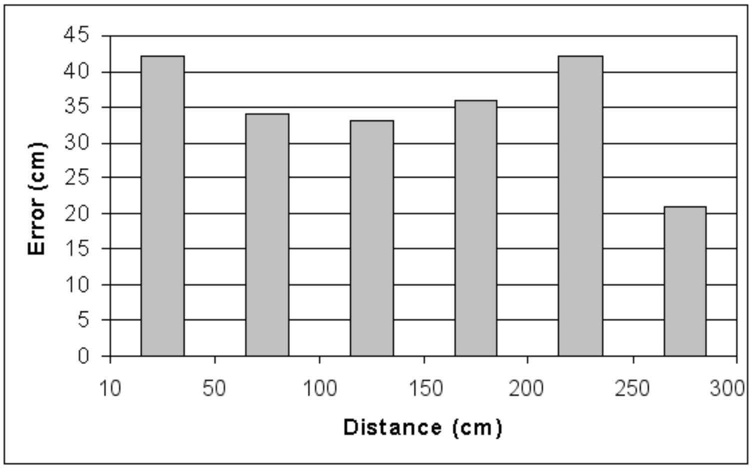

Figure 9: Maximal error among different robots (over 10 trails in both sets) vs. distance under the best and worst alignment conditions.

In Figure 9 the absolute maximal difference between the maximum/minimum and the average of the received signal strengths are shown as a measure of the distance they would represent, thus showing the distance error vs. range (we consider here data collected in both optimal and worst alignment conditions, i.e. Figure 7 in its integrity). In this histogram the first bin ignores the drop that occurs at distances less than 10 cm. In Figure 9 it can be seen that the combination of intrinsic power emission nonlinearities of the optoelectronic components and the nonlinear decrease of the amplitude as a function of the distance makes hard to predict precisely the evolution of the error exclusively based on theoretical predictions. We can however see that for any distance in the range of view of our relative localisation system the measured distance error is maximally about 40 cm. However, the resolution of the range measurement for a given sector is the distance represented by a change in the least significant bit, and is on average 3.25cm close up and 15 cm at 3m.

The accuracy of 45 degrees for the bearing was obtained by measuring the angle of rotation required to reliably change the received signal strength by more than the two least significant bits over 10 trails.

3.2 Communication

The range of the communication system was measured with one robot receiving a known data stream from another robot. This was done under the worst possible conditions, which is when the transmitting robot’s signal strength is the lowest possible (the angle directly in-between two LEDs) and the receiving robot’s signal strength is also the weakest (the angle directly in-between two photodiodes). The transmitting robot was moved from a range of 8m in a straight line towards the receiving robot in increments of 1m.

As shown in Figure 10 the error-free range of the communication system for data transfer is 6m under the worst possible conditions of transducer alignment. This is twice as much the range within which reliable distance assessment is obtained. The reliable data transmission range is greater than that of reliable distance measurement, because the RSSIs for range have to be significantly greater than the background noise levels, whereas data can be correctly received with signal strengths much closer to the background noise levels.

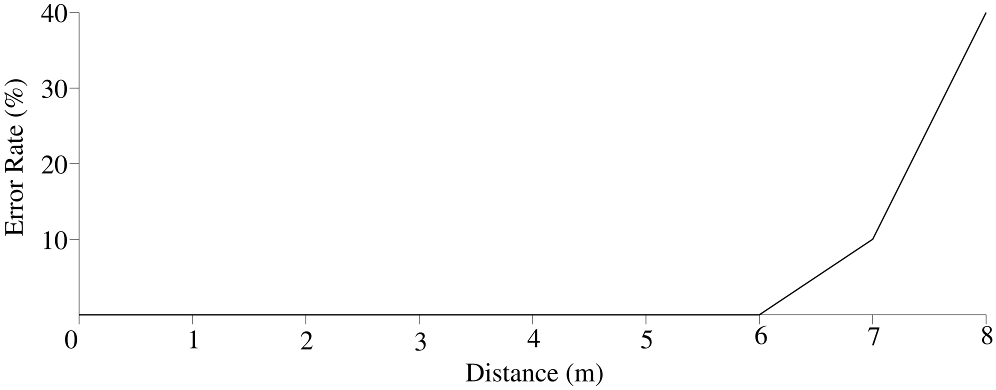

Figure 10: Error rate of the communication system under the *worst* conditions.

3.3 Flocking

As a further test of the system, real-time flocking was implemented on a group of eight robots using the same algorithm presented in [Kelly96/7]. This flocking algorithm presents many tests to the communication and localisation system: the robots are moving – often in close proximity, there is no set order to the robots, robots go out of and come back into range, and the algorithm relies on data being transceived between robots. Since in irregular formations, such as those generated by our simple flocking algorithms, one robot can be surrounded by an undetermined number of teammates, we conservatively set the maximum update frequency to be 2 Hz so that, in the worst case, the whole eight robots could form a local communicating network without incurring in an unreasonable number of message collisions.

The flocking algorithm is reactive, using a subsumption-like architecture [Brooks86] and characterized by a simple dynamic form of leadership. The four basic rules of the flocking algorithm are:

1. avoid objects (most basic behaviour with highest priority);
2. if no other robots are visible, become a leader and wander;
3. if in a flock, try to maintain position;
4. if a flock can be seen in the distance, speed up and head towards it, with more priority being given to following the closest visible leader.

Avoiding objects and wandering is achieved using a belt of eight proximity sensors, with the position and distance to other robots being determined from the infrared communication and localisation system. The localisation system is used to form an attractive force which brings the robots together, whilst the proximity sensors are used to act as a repulsive force to prevent collisions. Although the range measurement performed by our system is characterized by noise and it is fairly inaccurate (worst case 13% of error), it is good enough for maintaining flock cohesion: high resolution in range estimation is necessary only in the proximal region around a robot in order to achieve smooth flocking. Reynolds [Reynolds 87] came to the same conclusion in a noise-free environment. To help prevent head-on collisions only the rearward LEDs are switched on, thus the robots are only attracted towards the rear of each other. When a very close object is detected the avoidance behaviour forces the robots to slowly reverse, thus helping to avoid deadlocks and maintain the minimum separation distance between robots. The communication system is used by each robot to inform the other robots whether it is a leader or a follower.

The selection of the leader is dynamic, because like Reynolds's boids [Reynolds87] (and real flocks of animals) the flock should be able to split up to go around obstacles, and rejoin once past the obstacle. If a leader is pre-defined this is not possible. Also, since our robots operate in a finite bounded environment, there would be problems when the flock meets a boundary of the environment. In this case the pre-defined leader would have to fight its way through the other robots. Finally, if the pre-defined leader should stop working (i.e. die or is killed) then the whole of the flock would also fail.

Under a system where any robot can become a leader and can relinquish leadership when required, one or more leaders can co-exist. In this system the flock can split up into two smaller flocks to go around both sides of an obstacle and then rejoin once past the obstacle. If the leader should get trapped between other robots, then by definition it is now in the flock and therefore simply gives up leadership. One of the robots on the outside of the flock will take over the leadership and the rest will follow it. To ensure that this new leader does not simply turn around and rejoin the main body of the flock there is a short period of time for which it is not allowed to relinquish leadership to any robots.

From the point of view of any single robot, when no other robots are visible in front, it can become a leader. As leaders, robots wander around by moving forward in a straight line until an object is encountered, upon which they turn away from it. As stated above, to inhibit a leader from rejoining the flock immediately after leaving it, leaders do not start looking for other robots until after a few seconds of taking leadership. After this delay, when one or more other robots are visible ahead, leadership is relinquished. Robots that are followers head towards the greatest density of visible robots, with a higher priority being assigned to following the leader. This attracts robots towards each other, thus forming flocks. If as a follower, any robot sees a flock in the distance, it will catch up with them by increasing its speed. Whilst any robot is in a flock it tries to match the speed and direction of its nearest neighbours.

As shown in Figure 11 (which shows snapshots taken every two seconds from a group of eight robots flocking, see also http://www.coro.caltech.edu/Projects/Flocking/ab_flocking.htm for the full movie) the group of robots maintain a relatively close formation flock. The maximal distance between the robots is on the order of 5-6 robot lengths even when robots change positions within the flock.

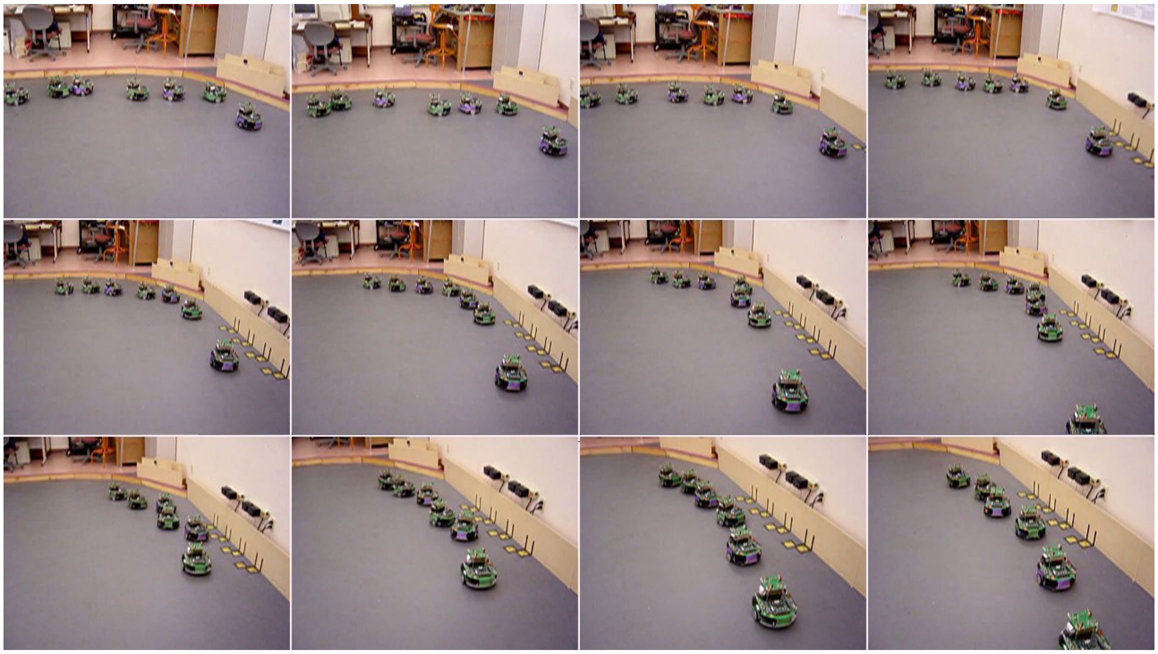

Figure 11: Flocking snapshots.

4 Discussion: System Limits and Possible Improvements

Like many other systems this infrared local positioning/communication system is a balance of design requirements and what is actually possible. This section defines some of the key limits of the system and how they could be improved. Some of those are due to hardware bottlenecks intrinsic to the positioning/communication approach used while other can certainly be relaxed by a further optimization of the system. For instance, among the intrinsic limitations of the system there are those mentioned in subsection 2.2 and 2.4: the modulation frequency is limited by the reduced bandwidth of LEDs and the distal range accuracy is limited by the measurement principle based on the amplitude of the received IR signal (decreasing with inverse of the distance at power two). Among the limitations that could be reduced or overcome after an optimization process, we see a clear potential for improvement on the following issues:

1. The current system assumes that the total number of robots used in the experiment is a priori known. This allows the engineer to establish the size of the arrays storing time stamps (see Subsection 2.6). If the maximum number of robots is a priori unknown it would still be possible to use this same system but the storage array would have to be set to the biggest size the PIC microcontroller has enough memory for (which off course sets the limit to the maximum number of robots that it could receive during any single time step). Under these conditions it would be impractical to send the whole array back to the Moorebot each time it requests the positions of the other robots, but instead the PIC itself would have to figure out the number of unique transmissions received and transmit just these to the Moorebot.
2. At the moment the system uses a fixed transmission interval between transmissions. The current way of deducing the number of other robots currently within range requires that this is kept constant (see section 2.5) and lies on the best estimate of the user of how many robots will get close together during the experiments. Using a dynamic transmission interval which would be a function of the number of robots actually participating in a given local network, would allow the system to increase the communication throughput as well as relative positioning rate as a function of number of robots participating to the network. The lower the number of robots and the higher the positioning/communication rate. However, the initial protocol for establishing the local network not only should include the randomisation of transmitting time slots but also a negotiation of the update frequency for the positioning/communication system.
3. Currently, the area of emission (sector and range) is static and pre-established at the beginning of the experiment. Similarly to point 2, adapting the area of emission as a function of the requirement of the local network would allow the system to reduce interferences with neighboring local networks and therefore in turn increase communication throughput and rate of the local positioning system.
4. The maximum bandwidth of the system is set by the two 455kHz intermediate filters. The original IR system on the Dwarfs [Kelly96, 97] that this system is based on was designed to be multi-channel and therefore, in order to increase the number of concurrent channels available, each of them was given a low bandwidth. To save development time the same filters were used as on the original system. Since only one channel is now used it could be made to have a much wider bandwidth.
5. As stated in section 2.5, the 6ms of preamble is, on its own, enough transmission time for the other robots to obtain the range and bearing. Therefore, if data is not required it would be possible to increase the update rate of localisation up to 50Hz within range of each other by not transmitting the data byte, the start byte and the start bit.
6. The accuracy of the obtained bearing could be improved by increasing the number of photodiodes (e.g., from four to eight per robot) and by a dedicated interpolation algorithm. Interpolating the data of several sectors receiving a signal from the same source may allow not only for an increased bearing accuracy but also for an improved range precision. In particular, by using a mathematical model of the reception profiles of the two photodiodes receiving the strongest signal strength, it would be possible to increase the accuracy of range measurement with respect to the incident angle of reception.

7. A more systematic calibration procedure for individual robots and individual sectors could lead to a more homogeneous system and therefore an increased overall accuracy. For instance, we could tune the resistors controlling the emission power for each sector and the gains of the photodiode amplifiers, so that the differences in emission and reception due to heterogeneities in the manufacturing of the optoelectronic components or inaccurate sensor positioning could be minimised.
8. Finally, it is worth noticing that although the range estimation principle based on the amplitude of the incoming IR signal presents intrinsic limitations, in particular in the distal accuracy, an increased update rate in position obtained by any of the point 2 to 4 above would allow the system to average over multiple samples within the same time window, and therefore statistically improve the overall signal-to-noise ratio in range and bearing.

5 Conclusion

In this paper we have presented the design of a prototype integrating local positioning and communication system. Although the prototypes have not been systematically characterized yet, we have shown a few preliminary but promising results, including a non-trivial flocking experiment using eight real robots, which have provided valuable insight on the performance of the system. As outlined in Section 4, the system can be further optimised on several points, both at firmware and hardware level. The electronic design we proposed in this paper, a combination of RF transceivers and IR front-ends, exploits technologies based on off-the-shelf components. Indeed, the solution we adopted takes advantage of the fact that electromagnetic waves under light form can be easily controlled both in range and cone of emission by using tiny and inexpensive optoelectronic components available on the market. However, we believe that in the future hybrid design such ours will be no longer competitive with integrated multi-antenna RF transceivers, electronic modules capable of not only high bandwidth communication but also of controlling range and direction of emission/reception (see for instance [Guan01]).

6 References

Bay, J. S. (1995) *Design of the 'Army-Ant' co-operative lifting robot*. IEEE Robotics & Automation Magazine, March, pp. 36-43.

Bonabeau E., Dorigo M. & Theraulaz G. (1999) *Swarm Intelligence: From Natural to Artificial Systems*, SFI Studies in the Science of Complexity, Oxford University Press, New York, NY, 1999.

Brooks, R.A. (1986) *A robust layered control system for a mobile robot*. IEEE Journal of Robotics and Automation. Vol. RA-2, No.1, pp. 14-23.

Fredslund, J. & Matarić, M. J. (2002) *General Algorithm for Robot Formations Using Local Sensing and Minimal Communication*. Special Issue on Advances in Multi-Robot Systems, Arai T., Pagello E., and Parker L. E., editors, *IEEE Trans. on Robotics and Automation*, Vol. 18, No. 5, pp. 837-846.

Genovese, V. et al. (1992) *Self organising behaviour and swarm intelligence in a pack of mobile miniature robots in search of pollutants*. International Conference on Intelligent Robots and Systems.

Guan, X. & Hashime, H. (2001) *Si-based Fully Integrated 24GHz Transceiver.* http://www.chic.caltech.edu/SiTranceiver.pdf

Guy, C.G. (1992) "Data communications for engineers." Macmillan. pp. 133-134.

Hayes, A. T., Martinoli, A. & Goodman, R. M. (2002a) *Distributed Odor Source Localisation*. Special Issue on Artificial Olfaction, Nagle H. T., Gardner J. W., and Persaud K., editors, IEEE Sensors, Vol. 2, No. 3, pp. 260-271.

Hayes, A. & Dormiani-Tabatabaei, P. (2002b) *Self-Organized Flocking with Agent Failure: Off-Line Optimization and Demonstration with Real Robots*. IEEE Int. Conf. on Robotics and Automation IROS02, May, Washington DC, USA, pp. 3900-3905.

Hodgins, J. & Brogan, D. (1994) *Robot herds : Group behaviours for systems with significant dynamics.* Artificial Life 4, Eds. R.A.Brooks & P.Maes, MIT Press, pp. 319-324.

Howard, A., Matarić, M. J. & Sukhatme, G. S. (2002) *An Incremental Self-Deployment Algorithm for Mobile Sensor Networks*. Autonomous Robots, Special Issue on Intelligent Embedded Systems, Vol. 13, No. 2, pp. 113—126.

Kelly, I. D. & Keating, D. A. (1996) *Flocking by the Fusion of Sonar and Active Infrared Sensors on Physical Autonomous Mobile Robots*. Proc. of the Third Conf. on Mechatronics and Machine Vision in Practice, Gaimardes, Portugal, 1996, Vol. 1, pp. 1-4.

Kelly, I.D. (1997) *The Development of Shared Experience Learning in a Group of Mobile Robots*. Ph.D. Thesis, Department of Cybernetics, University of Reading, UK, pp 2/4-2/12.

Kube, C. R. & Bonabeau, E. (2000) *Cooperative Transport by Ants and Robots*. Robotics and Autonomous Systems, Vol. 30, No. 1-2, pp. 85-101.

Levy, S. (1992) *Artificial Life : The quest for a new creation*. Penguin Books. pp. 300-301.

Matarić, M. J. (1995) *Designing and understanding adaptive group behavior*. Adaptive Behavior, Vol. 4, No.1, pp. 51-80.

Ota, J., Arai, T. & Kurabayashi, D. (1993) *Dynamic grouping in multiple mobile robots system*. Proceedings of the 1993 International Conference on Intelligent Robots and Systems, pp. 1963-1970.

Reynolds, C.W. (1987) *Flocks, herds and schools : A distributed behavioural model*, Computer Graphics : Proceedings of SIGGRAPH '87, Vol. 21, No. 4.

Unsal, C. & Bay, J. S. (1994) *Spatial self-organisation in large populations of mobile robots*, IEEE International Symposium on Intelligent Control, Columbus, Ohio.

Welsby, J. & Melhuish, C. (2001) *Autonomous Minimalist Following In Three Dimensions: A Study with Small-Scale Dirigibles*. Towards Intelligent Mobile Robots. Manchester, UK. Computer Science Tech Report UMCS-01-4-1.

Winfield, A.F.T. & Holland, O.E. (2000) *The application of wireless local area network technology to the control of mobile robots*. Microprocessors and Microsystems, Vol. 23, pp. 597-607.