Distributed Sensing and Data Collection Via Broken Ad Hoc Wireless Connected Networks of Mobile Robots

Alan FT Winfield

Intelligent Autonomous Systems (engineering) laboratory, University of the West of England Bristol, Coldharbour Lane, Frenchay, Bristol BS16 1QY, UK. Email: Alan.Winfield@uwe.ac.uk

Abstract. This paper reports on ongoing work to develop *ad hoc* wireless networking for application in distributed mobile robotics. The paper discusses a mission scenario in which a number of mobile robots are required to autonomously disperse into a physically bounded region, take sensor readings at predetermined time intervals, and then communicate the sense data back to a single collection point. There is assumed to be no overall command and control structure. The robots are assumed to be equipped with low-power short-range wireless network interfaces, which only allow direct communication with near neighbours, and overall an *ad hoc* (multi-hop) wireless networking scheme is employed. Furthermore, it is assumed that there are insufficient robots to allow full wireless connectivity through the region so that, at any instant, the robots actually appear as a set of disconnected sub-nets (i.e. a broken *ad hoc* network). The paper proposes a mechanism that exploits the mobility of the robots to overcome the lack of overall wireless network connectivity, resulting in a sub-optimal but robust scheme for garnering sense data. The paper presents simulation results that show how the arithmetic mean data collection delay varies with the number of robots deployed. The paper concludes that in mission scenarios that do not require data collection in real-time and where occasional data loss (erasure) is acceptable, then the proposed data collection mechanism is feasible, robust and economical in terms of the number of robots required.

Key words. Distributed sensing; *ad hoc* wireless networking.

1 Introduction

Consider a group of mobile robots that are required to autonomously disperse throughout a region, perform distributed sensing, monitoring or surveillance, and pass the sense data to a single collection point. Assume that the robots are equipped with low-power wireless transceivers whose range is too short to allow direct communication with the data collection point, but sufficient to allow robots to communicate with near neighbours. Under these circumstances the only practical approach to distributed command, control and sensing is to employ an *ad hoc* wireless networking scheme.

The networking literature defines an *ad hoc*, or *multi-hop mobile* wireless network, as a network with a set of geographically dispersed nodes in which each node is able to forward packets for other nodes that cannot communicate directly. Thus each node is able to act as a router, as well as a data source, or sink. A characteristic of *ad hoc* mobile wireless networks is that it is virtually impossible for each node to know the entire network topology at any given time, which means that optimal strategies for message routing in wide area networks cannot be used. Broadcast protocols that do not require knowledge of the entire network topology have been described [1], [2] and recently Basagni *et al* [3] introduced a mobility-transparent deterministic broadcast mechanism for such networks.

Since the advent of high performance wireless local area network technology at relatively low cost, its use for wireless control of mobile robots has become a practical proposition. Wireless local area networked miniature mobile robots, using TCP/IP (i.e. Internet Protocols), have proven highly successful as a platform for experiments in collective robotics [4], [5]. The issues of communication with large populations of mobile robots have received attention, for example in Gage [6]. Genovese *et al* [7] considered both swarm-like self-organisation and communication. Arai *et al* [8] proposed an information diffusion mechanism for distributing command data to multiple robots and analysed the overall time-delays incurred. However, *ad hoc* radio networks based on Internet Protocols have not yet, to the author's knowledge, been applied in the field of distributed mobile robotics.

This paper discusses the feasibility of applying *ad hoc* wireless networking to the problem of garnering sensory data from a spatially dispersed group of mobile robots. The paper looks at perhaps the simplest means of spatial deployment of a group of mobile robots into a physically bounded region, namely by random diffusion. Unlike the network communications literature this paper considers the likely scenario in which the *ad hoc* wireless network becomes disconnected and fragments into smaller sub-nets. The paper proposes a mechanism that exploits the mobility of the robots, and hence the highly dynamic network topology, in order to collect the distributed sense data despite an apparently 'broken' *ad hoc* network. Simulation results are presented, which show how the data collection time delay varies with the number of robots deployed and with each robot's communications range.

2 A model of robot connectivity for analysis and simulation

For the purpose of the simulation presented in this paper we shall assume an idealised wireless communications model for the mobile robot. Firstly, we shall assume that the wireless antenna fitted to each mobile robot is uniformly omni-directional in the horizontal plane. This is a sound assumption since the horizontal polar diagram of a simple vertical monopole antenna is a circle [9]. Secondly, we assume that the modulation scheme employed by the wireless transceivers experience a 'threshold effect'. In other words, as the distance between transmitter and receiver increases and the received signal strength reduces according to the inverse square law, the receiver's ability to demodulate the received signal will suddenly fail as the threshold is reached. Thus for any given transmitter power level (or receiver sensitivity) each mobile robot will be at the centre of a circle of radius *r* (i.e. the communications 'range'). Reliable communications between any two robots separated by less than *r* is thus guaranteed, whereas two robots more than *r* apart will be unable to establish a communications link. Again this is a reasonable assumption, since mobile robots are likely to employ Frequency Shift Keying (FSK), Phase Shift Keying (PSK) or one of its many variants, all of which experience the threshold effect [10]. A third assumption is that the wireless network will employ some form of time, frequency or code diversity in order that many robots (wireless nodes) may share the same RF spectrum and interoperate without compromising the first two assumptions. Again this is reasonable if, in practice, we employ either Frequency Hopping (FH) or Direct Sequence (DS) variants of spread spectrum modulation [11] together with some strategy for Carrier Sense Multiple Access (CSMA).

These assumptions provide us with an idealised model of wireless network communications for distributed mobile robots that lends itself to straightforward simulation but with reasonable confidence that the results will transfer successfully to physical implementation. Figure 1 illustrates this connectivity model for a set of 5 randomly dispersed mobile robots in a 2D space, numbered 0 to 4. Each mobile robot is shown with a circle that defines its radio range *r*. It is easy to see by inspection that robot 0 can communicate with robot 1. Robot 1 can communicate directly with robots 0, 2 and 3. Robots 0 and 3 can communicate either with 2 hops via robot 1 or with 3 hops via robots 1 and 2. Robot 4 on the other hand is currently out of range of any other robot but, depending on its heading, might move into range with robots 2 or 3.

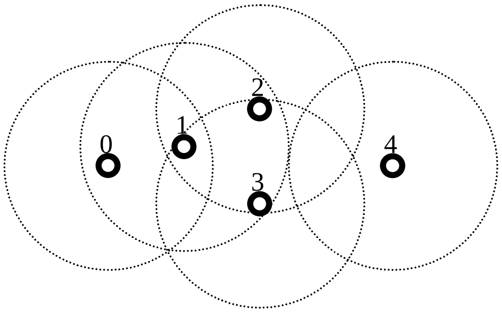

Fig. 1 Wireless range of randomly dispersed mobile robots in 2 dimensions

The networking communications literature conventionally models an *ad hoc* network as an undirected graph, $G=(V,E)$ in which $V=\{p_0,...,p_{n-1}\}$ is the set of wireless nodes and there is an edge $(p_i,p_j)\in E$ if and only if p_i and p_j are in range. In this case p_i and p_j are neighbours. The *distance* $d(p_i,p_j)$ between two nodes p_i and p_j, $0 \le i, j \le n-1$, is defined as the minimum number of hops between p_i and p_j. The maximum distance between any pair of nodes is called the *Diameter D* of the network. Figure 2 shows the same set of five robots in fig. 1 as an undirected graph, or to be precise two undirected graphs, one with four nodes and the other with one. The *Diameter* of the graph which comprises robots 0 to 3, $D = 2$. Additionally we define the *Connectivity C* of the overall network as the number of edges, and it is easy to see that for *n* robots $0 \le C \le \frac{1}{2}n(n-1)$.

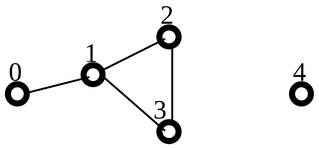

Fig. 2 Figure 1 represented as an undirected graph

3 Spatial organisation of ad hoc networked mobile robots

Consider the following mission scenario for distributed mobile robots. A group of mobile robots is required to perform some sensing or surveillance task that requires that the robots be spatially distributed. Once the robots have been deployed the sensor data from each robot is then moved across the group, via the wireless network, and garnered by a nominated robot (robot 0, say) on the edge of the group. Robot 0 is in turn connected to some base station for forward transmission of the collected sensor data. Consider now the possibility that directly deploying each of these robots into their final spatial distribution might be difficult or dangerous. This might be because they are to be deployed into a hostile environment or simply because there are too many robots.

In order to visualise this scenario consider, for example, that a large number of small submersible mobile robots are required to perform marine pollution monitoring. It would make sense to introduce the robots into the marine environment from the monitoring vessel *en-masse* and then rely on the robots to disperse themselves throughout a specified 3D volume of ocean. A number of small robotic deep-space probes might similarly be deployed from the hold of the *Space Shuttle* and then self-organise into a spatial formation before embarking on their mission. In all of these mission scenarios the ability of the robots to spatially self-organise, relative to each other, without the need for sophisticated means of absolute navigation might be a significant advantage.

Consider also a scenario in which the sensing robots act autonomously. In other words no command, control or time synchronisation messages whatsoever are required, from robot 0, to the rest of the group. In other words all motion control decisions are made locally, on the basis of local information (including, perhaps, information about local wireless connectivity). We shall assume that each robot has an accurate on-board clock, synchronised at the time the robots are deployed, and that data samples are undertaken at predetermined time intervals.

4 Spatial distribution within a physically bounded region

Consider now how spatial organisation can be achieved within a physically bounded region. For terrestrial mobile robots operating in, essentially, 2 dimensions the physical bounds might be natural terrain features such as hills or rivers; for laboratory robots the bounds are provided by the walls at the edge of an experimental arena. Submersible pollution monitoring robots operating within an inland lake would similarly be constrained by natural physical barriers. In such a bounded environment the simplest way of achieving spatial self-organisation would be by random diffusion. The robots are simply introduced *en-masse* at one point. Each then chooses a random heading and sets off along that heading. With an appropriate change of heading when a robot approaches a boundary, or when robots collide then, eventually, the robots will diffuse to fill the bounded region.

Fig 3. Initial robot deployment

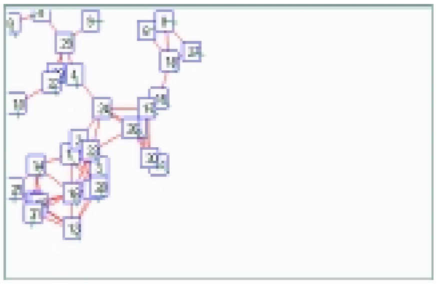

Fig 4. Robot distribution after 40 timesteps

To illustrate random diffusion consider the introduction of 32 robots into one corner of a rectangular arena, as shown in fig. 3. The robots have a wireless range $r = 50$, much smaller than the arena dimensions of 450 x 250. At time = 0 each robot chooses a random heading, from 0 to 359 degrees, and when in collision with the arena wall robots assume a new random heading anywhere within the 180 degree arc facing away from the wall. Initially we have a maximally connected network, $C = 496$, in which each robot is only one network hop from any other. Each robot moves a distance of 5 units along its heading in 1 time step, except robot 0 which remains anchored at its starting position (robot 0 is the fixed data collection point for the whole network). After 40 time steps the robots have formed a distinct wavefront and have filled about one quarter of the arena, as shown in fig. 4. At this point the connectivity has fallen to $C = 76$, but the network is still fully connected; in other words there is still a network path between robot 0 and every other robot. The lines shown joining individual robots in fig. 4 indicate neighbours that are in wireless range (or, edges, of the network graph). The short lines shown pointing from the centres of each robot indicate, compass fashion, that robot's current heading.

Figure 5 shows the situation after 80 time steps. The robots have now almost filled the entire arena, but the level of connectivity has fallen to just 39, less than 8% of the maximal connectivity. Most significant however, is that the *ad hoc* network has now broken into a number of fragments, or sub-nets, with a number of robots out of range of any neighbours at all. There is now no network path back to robot 0. After 120 time steps the robots have diffused throughout the region and any wavefront structure has now disappeared, as shown in fig. 6. The overall level of connectivity is very low at just 23 and there are no less than 7 sub-net fragments. It is also interesting to see how much the sub-net structures have changed between figs. 5 and 6.

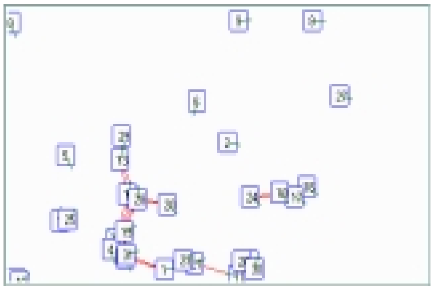

Fig. 5 Robot distribution after 80 timesteps

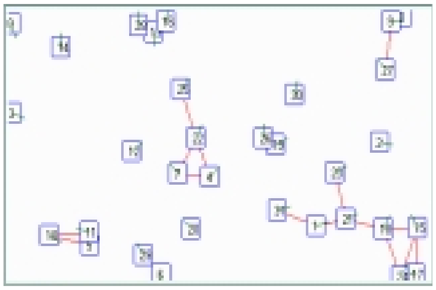

Fig 6. After 120 timesteps

Now it would be easy to conclude that the scenario illustrated in fig. 6 is not viable as a network for distributed sensing and that we need either more robots, or robots with a longer wireless range, in order to build a fully connected network. Such a highly fragmented (i.e. 'broken') *ad hoc* network would, at first sight, appear useless for collecting sensory data. However, this is not a static network but is implemented on mobile robots, whose most important attribute is of course their mobility. Figure 6 is just a freeze-frame of a highly dynamic process in which individual robots, and sub-nets, quickly move in and out of range of each other.

5 A mechanism for data collection via a broken ad hoc network

Let us propose a distributed data sensing and collection strategy for sparsely populated ad hoc connected networks of mobile robots, in bounded regions as follows. Assume that sensory data must be collected from every robot and that all data must be garnered at robot 0.

- At a predetermined time every robot simultaneously samples new sensory data (based upon their internal clocks, which are assumed to have been synchronised at time = 0);
- Every robot then immediately broadcasts its sensory data to every one of its neighbours;

- In turn each neighbour broadcasts any received data, together with its own sensory data, to its neighbours where it is stored in a data buffer. In this fashion sensory data is propagated through each of the sub-nets;
- Each time a robot detects that another robot has come within wireless range (and becomes a new neighbour) then all data within its buffer is broadcast to the new neighbour and to its neighbours, and vice-versa;
- Each robot's sensory data is marked with that robot's ID (IP address), and when robot 0 has received sensory data from every robot in the network then the data collection for this sample period is complete.

As a practical proposition this algorithm makes a number of assumptions. Firstly, that data collection is not required in real-time, and that the application can tolerate variable time delays between sensory sampling and completion of data-collection. Secondly, that the bandwidth of the wireless networking is sufficiently high and the volumes of message data sufficiently low, that all message buffers can be propagated between joining sub-nets before robot mobility re-fragments the joined sub-net. This is a reasonable assumption given typical wireless LAN bandwidths of 1-2 Mbytes/s. Finally, the algorithm assumes that each robot has sufficient local storage to be able to buffer messages from a single sample without overflow. This is again a reasonable assumption given current memory technology. The proposed algorithm may result in a robot receiving multiple copies of another robot's sensory data; it will however only need to store one copy in its data buffer, thus placing an upper bound on the size of buffer storage required per sample.

At first sight the use of a 'broadcast' mechanism to propagate sensory data across sub-nets may appear counter-intuitive given that the overall objective is to concentrate distributed sensory data to a single collection node. The rationale behind this approach is that by broadcasting the same sensory data to as many robots as are within sub-net connectivity at any given time, and therefore duplicating the same data widely across the network, we are minimising the time it takes for data to propagate back to the collection node. One way of visualising this approach is to think of the sensory data as 'infecting' the whole of a sub-net. As soon as any member of that sub-net contacts another sub-net, the sensory data will in 'virus fashion' infect the new sub-net, and so on. Thus in the same way that the robots physically diffuse across the bounded region, the sensory data diffuses, in reverse, back to robot 0. Note that we do not need to place a restriction on the sensory data sampling frequency. We do not, for instance, need to wait until data collection for one sensory sample is complete before triggering the next sample. Providing that there is sufficient buffer storage within each robot then we can exploit the highly parallel structure of the distributed robot network to allow multiple samples to be simultaneously propagating back to the collection node.

6 Simulation results

To assess the performance of the proposed algorithm the author has constructed a computer simulation in which the number of robots *n*, and their wireless range *r*, can be varied. The key performance indicator is the length of time for data collection to complete, in simulated time steps. Figure 7 shows the data collection time delay for range *r* = 100, with error bars. Figure 8 shows the arithmetic mean data collection delay graphically plotted against the number of robots, for a number of different values of wireless range *r*. All simulations were performed in an arena of 450 x 250 units, and each robot moved 5 units per time step. Each run, for a given number of robots and wireless range was undertaken 100 times, and the arithmetic mean plotted in fig. 7. To ensure that each run gave an independent value for the time to complete data collection, following completion the robots were moved through a further 100 time steps before initiating a subsequent sample and data collection.

Fig. 7 Data Collection Delay, *r* = 100

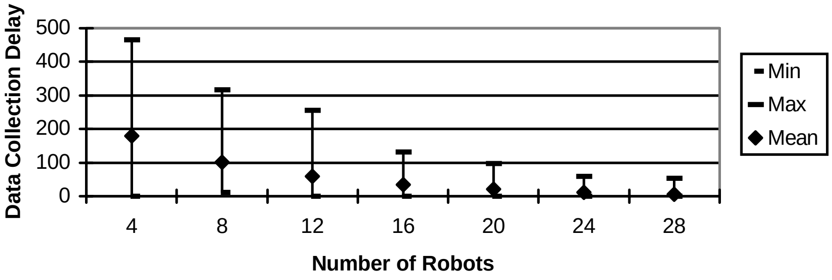

Fig. 8 Mean Data Collection Delay

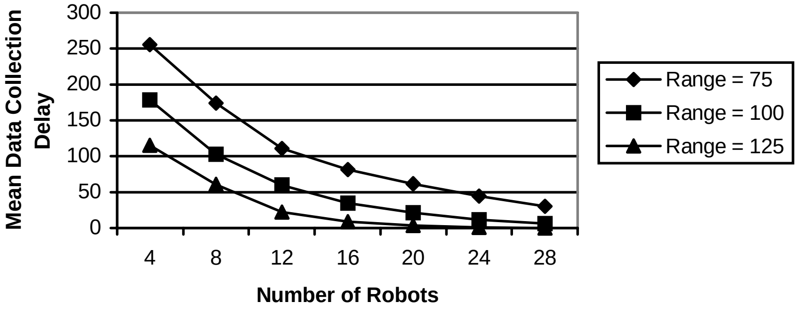

As would be expected the data collection delay asymptotically approaches zero as the population of robots in the bounded region increases. (Zero represents the time delay experienced by a fully connected network, i.e. one in which there is, at data sampling time, a multi-hop path from every robot back to robot 0). Also as expected, increasing the radio range *r* reduces the mean data collection delay for a given number of robots. These simulation results demonstrate that, at the cost of some delay time, it is feasible to propose distributed sensing and data collection across a ‘broken’ ad hoc wireless network by exploiting the mobility of the robots. Another way of thinking about the proposed mechanism is that we are trading wireless utilisation and buffer storage (which are high) for robot mobility which is low (compared with a scenario without wireless communications in which each robot would have to physically transport its collected data back to the collection point).

7 Discussion and Further Work

Consider now the buffer memory requirement of the proposed data collection mechanism. As already stated, for a single data sample, the buffer memory requirement is bounded by the number of robots. However in a practical distributed sensing scenario samples would be required periodically. This could potentially result in an unbounded buffer requirement because of the unbounded upper limit on the data collection delay time. In practice buffers need to have a fixed size limit and this will result in the possibility of erasures (data loss) because of buffer overflow. The probability of erasure can be minimised by careful design of the buffer size in relation to predicted data collection delays (from fig. 7). The provision of a circular buffer arrangement in which the oldest data is overwritten by the newest would provide a simple strategy for buffer management. It is, however, important to recognise that ‘erasure’ will very rarely mean the loss of an entire data sample. It is much more likely that the majority of the data from the oldest sample will already have propagated back to the collection node by the time its buffer needs to be overwritten by the latest sample. Thus the proposed data collection mechanism, via a broken *ad hoc* network, is likely to result in - at worst - partial data for occasional samples. In many remote sensing scenarios, for instance, pollution or environmental monitoring, this compromise would be perfectly acceptable.

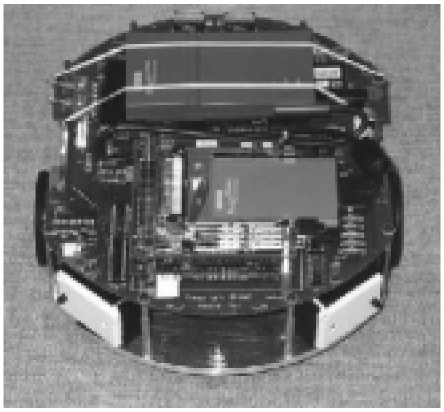

Fig. 9 The LinuxBot

As stated in the introduction, the paper reports on current work on the implementation of *ad hoc* wireless networking on distributed mobile robots. An important aspect of the work of the author’s laboratory is the transfer of theoretical concepts to engineering realisation. To this end we are currently implementing an *ad hoc* wireless network, using Internet Protocols, on a fleet of fifteen wheeled mobile robots (the *LinuxBots* [12]). Figure 9 shows a single LinuxBot. Thus, in the laboratory environment, we will be able to compare the real-life performance of the distributed data sensing and collection mechanism discussed in this paper, with simulation and evaluate the effects of noise and interference.

In parallel, we shall extend the strategies for robot deployment and data collection in physically bounded regions, as discussed in this paper, to unbounded regions. Recalling the principle that decisions regarding motion control must be taken only on the basis of local information and without the benefit of received command data, *is it possible to achieve spatial self-organisation in unbounded regions on the basis of network connectivity information alone?*

References

1. I Chlamtac and S Kutten. (1985) On broadcasting in radio networks – Problem analysis and protocol design, IEEE Trans. Commun., **COM-33**, 1240-1246.
2. I Chlamtac and S Kutten. (1987) Tree-based broadcasting in multihop radio networks, IEEE Trans. Comput., **C-36**, 1209-1223.
3. S Basagni, D Bruschi and I Chlamtac. (1999) A mobility-transparent deterministic broadcast mechanism for ad hoc networks, IEEE/ACM Trans. Networking and Computing, **7/6**, 799-807.
4. R Dawkins, O Holland, A Winfield, P Greenway and S Stephens. (1997) An interacting multi-robot system and smart environment for studying collective behaviours, Proceedings of the 8th International conference on Advanced Robotics, Monterey, CA.
5. AFT Winfield and OE Holland. (2000) The application of wireless local area network technology to the control of mobile robots, Microprocessors and Microsystems, **23/10**, 597-607.
6. DW Gage. (1993) How to communicate with zillions of robots, Proceedings of SPIE Mobile Robots VIII, Vol 2058, 250-257.
7. V Genovese, P Dario, R Magni and L Odetti. (1992) Self organising behaviour and swarm intelligence in a pack of miniature robots in search of pollutants, Proceedings of the 1992 IEEE/RSJ International Conference on Intelligent Robots and Systems, 1575-1582.
8. T Arai, E Yoshida and J Ota. (1993) Information diffusion by local communication of multiple mobile robots, Proceedings of the 1993 IEEE Conference on Systems, Man & Cybernetics, 535-540.
9. FR Connor. (1972) Antennas, Arnold, London.
10. BP Lathi. (1983) Modern Digital and Analog Communication Systems, Holt Rinehart and Winston.
11. P Wong and D Britland. (1995) Mobile Data Communications Systems, Artech House Publishers.
12. http://www.ias.uwe.ac.uk/linuxbot.htm.