Phase transitions in multi-robot interactions

Colin G. Johnson* and Steven J. Brodie
Computing Laboratory
University of Kent
Canterbury, Kent, CT2 7NX
England C.G.Johnson@ukc.ac.uk

5th April 2001

Abstract

Phase transitions are where a small change in a parameter causes a qualitative change in the behaviour or state of a system. We can apply this concept to determining the connectivity of a population of mobile robots. Consider a population of mobile robots which can pass messages to other nearby robots. When the number of robots is small, or their working environment is large, robots will only be able to pass a message to a small number of others. However there is a point where if we increase the population by a small number, increase the communication range slightly, or reduce the working area by a small amount, the connectivity of the population increases suddenly to a point where almost all robots are communicating with each other. The consequences of this for multi-robot learning are discussed.

1 Introduction

Many mobile robotics systems consist of teams of robots which communicate with one another, e.g. sending each other information about the environment that they share, or about successful strategies for achieving certain tasks. It has been shown [7, 8] that such communication can improve the overall performance of such teams on complex tasks.

If we are to make effective use of such teams of robots, we need to understand what conditions will facilitate the maintenance of an effective level of communication. Given a particular environment, we would like to know how many robots need be deployed to cover the space effectively, and how good the communication need be between pairs of robots in order for the team of robots to be connected with respect to communication.

In this paper we show how *phase transitions* can occur in such systems. This is where small changes to the system make a sudden qualitative change in the global structure of the system. This type of behaviour gets its name from physics, where changes between phases (e.g. liquid to solid, liquid to gas, or *vice versa*) occur with small changes of temperature or pressure.

2 Random graphs and small worlds.

As an underlying model of the communication between agents, we make use of graphs which represent snapshots of the situation at any particular time, where the nodes represent the robots, and the edges represent the ability of a pair of robots to communicate. The graph-theoretic property needed for the network of robots to be able to intercommunicate is that the graph must be connected, as then a message can pass along a path from one robot to another. In practice we mean that a large number of nodes much be connected for a large proportion of the time.

One type of graph which we will see arising below when we apply this model is the *random graph* model [1]. A random graph $\mathcal{G}(n, p)$ is a graph produced by taking n nodes, then taking each pair of nodes

Proc. TIMR 01 - Towards Intelligent Mobile Robots, Manchester 2001.
Technical Report Series, Department of Computer Science, Manchester University, ISSN 1361 - 6161.
Report number UMCS-01-4-1. http://www.cs.man.ac.uk/csonly/cstechrep/titles01.html

*Author for correspondence

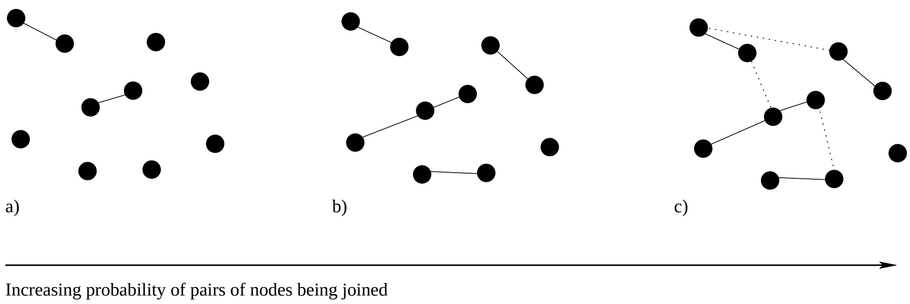

Figure 1: The formation of a giant component.

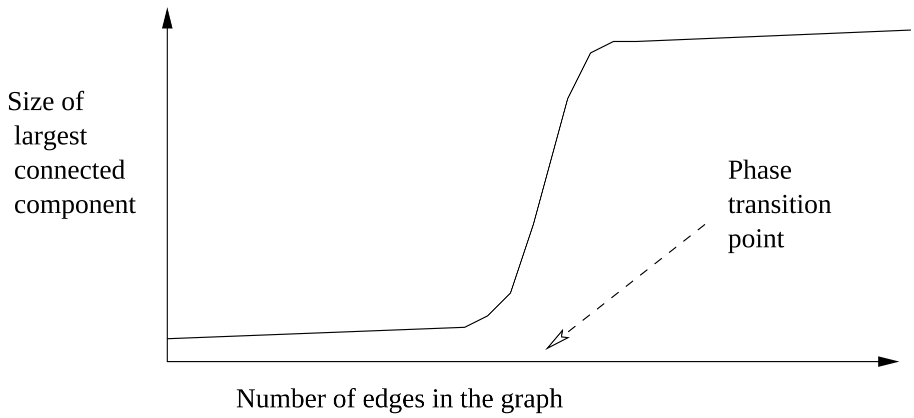

Figure 2: Caricature of the typical behaviour of a random graph process as we increase the number of edges.

in turn and joining them together with probability p. One of the most interesting aspects of random graphs is that as we increase the probability of two edges being joined, the clustering behaviour in the graph exhibits a non-linear behaviour. In this context a cluster is a set of nodes any pair of which can be joined by a set of edges in the graph, i.e. what is commonly called a *connected component* [11].

If the probability of joining each pair of nodes is small, i.e. the number of edges is small, then a number of small clusters (typically of size two or three) will form (figure 1a). However as we increase the number of edges we reach a point at which new edges have a high probability of linking two clusters (figure 1b). As a few more edges are added the clusters merge together increasingly, leading rapidly to the formation of a *giant component* (figure 1c) where most of the nodes in the graph are connected together.

This transition from lots of small components to a giant component happens with only a small change in the number of edges added to the graph. This is a process similar to a phase transition in a material, where the transition from (say) liquid to solid happens with a small change of temperature or pressure. For this reason we can refer to the critical value around which this phase transition happens as the *phase transition point*. Figure 2 shows a caricature of typical behaviour around this point.

The occurrence of these phenomena in graphs have been investigated by Erdős and Renyi [3]. They showed that for a random graph with n edges, the expected number of edges required for the giant component to emerge is $\frac{n}{2}\ln(n)$.

These ideas have been extended to other kinds of random graphs other than those where there is a uniform probability of each pair of nodes being formed. These ideas have been used to explain a number of effect. One example is the effect whereby every person appears to be connected to every other person in the world by a small string of acquaintances (the *small worlds* effect). Another example is showing the importance of casual acquaintances in transmitting infectious disease outside family groups. These ideas are explored further in [1, 10].

3 Applying these ideas to mobile robotics.

One way in which robots can work well together is for them to be able to communicate information which has led to their success at particular tasks, either by sharing information about their common environment or by communicating strategies which have led to success at some task [7, 8].

In order for this to be successful the population robots need to be in communication with each other for much of the time: every pair should be connected either by a direct communication link or via an intermediate chain of one or more robots.

If a phase transition point occurs in their ability to communicate then we need to know where that point might be, otherwise we run into difficulties and problems. An example of a situation in which such a problem could arise would be taking a functioning population of robots and placing them in a slightly larger environment, or removing a robot from the population. It would only take a small increase in the size of the environment in order to render the whole population useless.

3.1 A simple system for simulation.

As a test example for the simulation program described below we use the following model. We have a number of robots, represented by circles moving in a two-dimensional rectangular arena. As the robots move around this arena they communicate messages to other robots in the population by sending a message to all other robots which are within a certain radius of the centre of the current robot.

3.2 Predicting the phase-transition point.

In this system, where does this phase transition point occur? Take the number of robots (r), the communication radius (c), and the size of the arena (assume a rectangular arena of dimensions $h \times w$). Begin by considering the relationship between the number (n) of nodes in a random graph and the number (e) of edges. In order for the giant component to emerge we need

$$\frac{n}{2}\ln(n) < e \tag{1}$$

Consider the graph which underlies a particular configuration of robots in the arena, where we assign a node to each robot and an edge between every pair of robots which can intercommunicate. Clearly the number r of robots is equal to n in the equation above.

We can work out the number of edges thus. Firstly consider the area of the arena in which the robot can communicate (figure 3). The area of a given robot's communications radius is πc^2 (ignoring edge effects, so this will only be valid for communication radii which are small with respect to the area of the arena), and the area of the arena is width $\times$ height. So the proportion of the arena which the robot can communicate to is

$$\frac{\pi c^2}{wh} \tag{2}$$

Now there are $r-1$ other robots in the arena. If we assume that at a typical point in the run the robots are well distributed around the arena, then there will be

$$\frac{(r-1)\pi c^2}{wh} \tag{3}$$

robots in the communication radius of a typical robot. Finally we sum over all of the robots, and divide by two because the process above counts each robot twice (for each pair the first is counted when it is within the second one's communication radius, and *vice versa*).

Therefore for a given set of parameters we would expect the phase transition to occur when

$$\frac{r}{2}\ln(r) < \frac{r(r-1)\pi c^2}{2wh} \tag{4}$$

As an example of the use of this let us calculate how big a communication radius we need for a phase transition for a 100×100 arena containing 20 robots. Firstly we rearrange equation 4 so as to isolate c

$$\sqrt{\frac{2whr\ln(r)}{2\pi r(r-1)}} < c \tag{5}$$

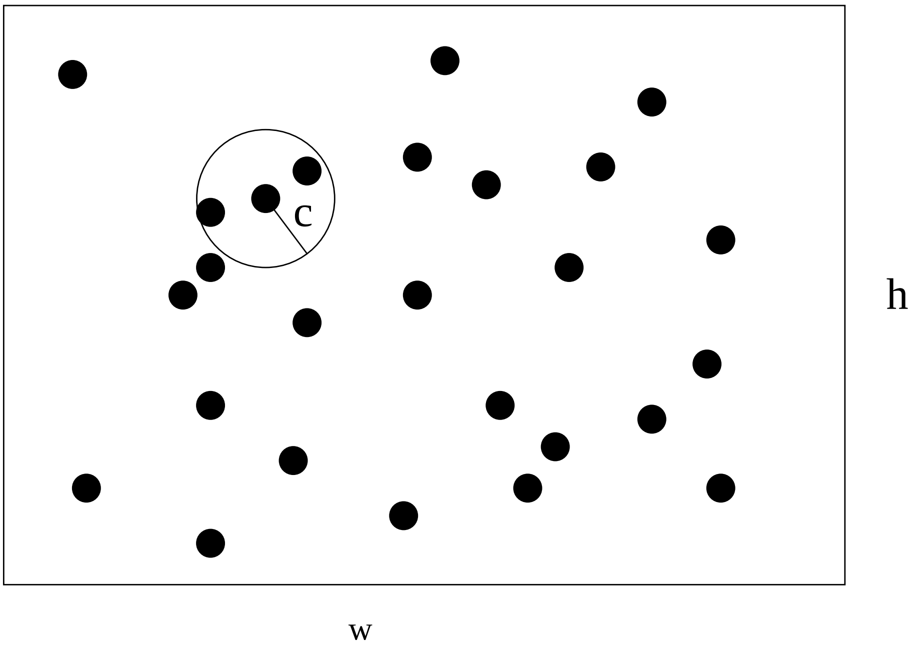

Figure 3: Counting robots in the area occupied by a robot's communication radius.

and we assume that the positive square root is taken as the negative result is physically meaningless. If we calculate this for our figures above we get that $c > 22.4$ as a criterion for the phase transition to occur. To investigate this further we shall do some simulations, results are given below.

We could extend this model to the more realistic case where there is a probability that the message-passing fails. To do this we multiply the probability of two robots being linked by this additional probability.

4 Experiments and results.

We have written a simulation program with which to investigate these phenomena. This simulates a number of mobile robots moving in a rectangular arena, and can be run either as an interactive graphical simulation (figure 4) or as a batch process for doing statistical analysis.

We are currently using this program to investigate the phase transition phenomena described above. For example we can create an experiment whereby the robots wander randomly around the arena, and record the average cluster size. In this case a cluster is the number of other robots which can be communicated with, either directly by being in the communication radius of the current robot or via passing a message through other robots. We then average this over 2000 timesteps of motion.

We would expect there to be a phase transition point at the location described above, i.e. somewhere around a communication radius of 22.4. The results are given in figure 5. This plots the maximum cluster size for a given communication radius. We can clearly see that the cluster size increases rapidly once the phase transition point has been passed.

Clearly this will have an impact on the ability of the robots to learn from one another, whether that is learning specific features of the environment or learning general strategies about how to cope with the environment.

We are currently carrying out work which looks into the impact that this behaviour has on the learning behaviour of robots. We are using a number of tasks, for example a foraging task and a rubbish-collection task, and investigating whether these phase transitions in communication affect the ability of the robots to learn from one another, including learning information about the environment and learning strategies about how to act in that environment. A preliminary summary of these investigations suggests that this affects the communication when the robots are carrying out tasks where the information is constantly

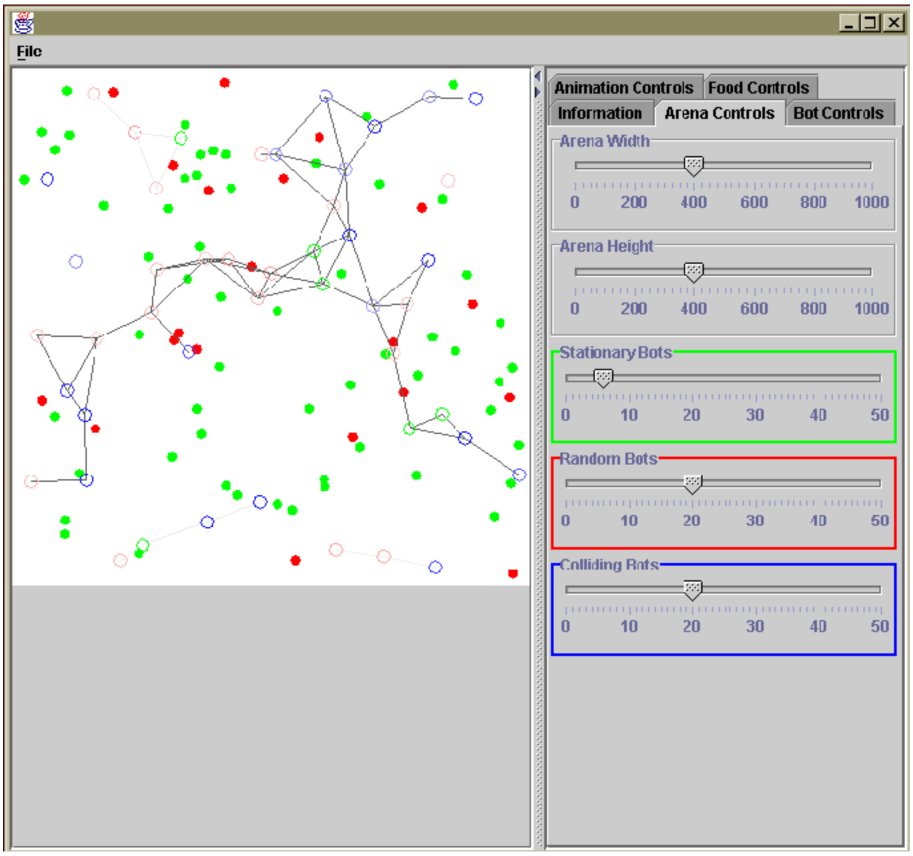

Figure 4: The simulation program in action.

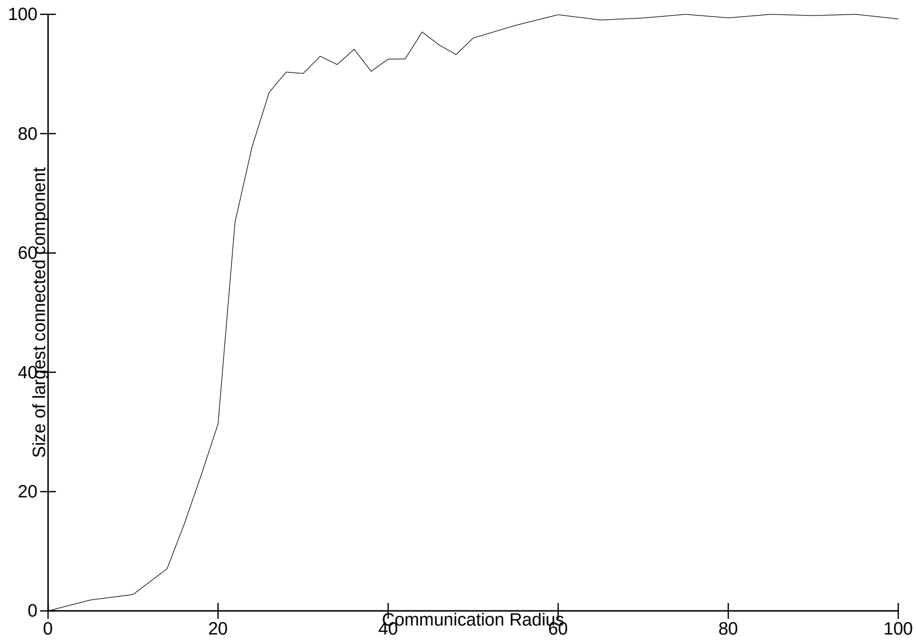

Figure 5: Phase transition in the maximum cluster size as we increase the size of the communication radius.

changing, e.g. a foraging task where fecund areas of the environment are regularly changing and the robots are communicating conjectured positions of such regions. Where the task is such that the information, once learned, continues to be of use (e.g. a good heuristic strategy), the phase transition has little effect, because robots are gradually moving from small cluster to small cluster even when the communication radius is small.

5 Discussion.

It seems likely that these arguments can be extended to more complex behaviours. One example is where robots have to explore a large region and then come back together to pool their results. Using the above ideas we can decide how close the robots need to come to get a collective knowledge of this information.

Another way in which this could be used would be in creating good configurations for populations of mobile robots rather than analysing existing situations. One example would make use of another special kind of random graph called a *small world network* [10]. Such networks model the communicative behaviour of social networks in populations of people, where large numbers of people are connected by small chains of acquaintances. This is popularly known as the “six degrees of separation” effect (the idea that every person in the world can be connected to every other through a chain of an most six acquaintances) [6], which has been experimentally investigated [9].

Such networks are characterized by having small groups (e.g. families) in which lots of information is shared, connected by networks of acquaintances which span otherwise disconnected parts of the network. Such networks have high connectivity for small numbers of links. This might provide a good model for exploratory robotics using many small robots [4]; small clusters of robots explore particular areas of interest, whilst communicating results via a small number of long-distance links. Such systems could be evolved in a bottom-up manner by including relevant connectivity measures in a fitness function for an evolutionary robotics system.

Another issue is that in many situations the robots will not be evenly distributed around their environment, either due to the behaviour being carried out, or due to the nature of the environment. In these cases we need more complex measures of connectivity. Examples of this are given in the books by Bollobas [1], and in particular in the book by Watts [10]. Watts defines two quantities, the *characteristic path length* and the *clustering coefficient*, which are shown to characterise the connectivity of graphs. Other ideas are suggested in the popular-science book by Gladwell [5]; formalizing these ideas mathematically and experimentally would be another interesting direction.

There are other areas in which similar questions can be usefully asked. As an example consider the following proposal for bringing telephone technology to parts of the world where installing a cable or base-station infrastructure is excessively costly or problematic. The idea is that each telephone will act both as a normal phone and also as a low-powered base-station, having enough power to transmit to another phone in the vicinity. This would then make a connection to another phone, and so on, until a set of links was established. If we were to install such a system we would need to know in advance how many units would need to be installed, how powerful they would need to be, and how the network might be affected by the removal of a small number of units.

Further details can be found in the second author's masters thesis [2].

References

[1] B. Bollobas. *Random Graphs.* Oxford University Press, 1985.

[2] Steven J. Brodie. Phase transition phenomena in robotic communications. Master's thesis, University of Kent, 2000.

[3] P. Erdős and A. Rényi. On the evolution of random graphs. *Publications of the Mathematical Institute of the Hungarian Academy of Sciences*, 5:17–61, 1960.

[4] A.M. Flynn and R.A. Brooks. Fast, cheap and out of control : a robot invasion of the solar system. *Journal of the British Interplanetary Society*, 42:478–485, 1989.

[5] Malcolm Gladwell. *The Tipping Point.* Little, Brown, London, 2000.

[6] John Guare. *Six Degrees of Separation: A Play.* Vintage, 1990.

[7] David A. Keating and Ian D. Kelly. Increasing mobile robot learning rates through sharing of experiences. In *Proceedings of the UKACC International Conference on Control*, 1998.

[8] Ian D. Kelly, David A. Keating, and Kevin Warwick. Mutual learning by autonomous mobile robots. In *Proceedings of the First Workshop on Teleoperation and Robotics*, Linz, Austria, 1997.

[9] Stanley Milgram. The small world problem. *Psychology Today*, 2:60–67, 1967.

[10] Duncan J. Watts. *Small Worlds: The Dynamics of Networks between Order and Randomness.* Princeton University Press, 1999.

[11] Robin Wilson. *Introduction to Graph Theory.* Addison-Wesley-Longman, 1996. Fourth Edition.