

“Small-World” Networks of Mobile Robots

Stergios I. Roumeliotis and Maja J. Mataric

stergios|maja@robotics.usc.edu <http://www-robotics.usc.edu/~stergios/aaai2000.html>

Department of Computer Science
University of Southern California
Los Angeles, CA 90089-0781

In order for a group of robots to coordinate collaboration during multi-robot tasks (Fontan & Mataric 1998), they need to communicate in an intelligent, purposeful way (Gerkey & Mataric 2000). If small groups of robots are involved, global communication (broadcasting, or one-to-all) is usually sufficient. The main advantage is that all the acquired information is available to all the members of the group. However when the number of robots increases so does the amount of data to handle. Information overflow can affect the performance of the separate teams working on different sub-tasks while limitations of the communication channel can cause interference and reduce the overall performance. A potential alternative to this would be to support local (one-to-a few) communication amongst the robots. The connectivity of such networks (communication topology) is usually assumed to be either completely regular (each robot communicates with its immediate neighbors forming a communication lattice), or completely random (each member of the group communicates with some random other members). However many biological, technological, and social networks lie somewhere between these two extremes (regular lattices vs. random graphs). Recently, a new form of coupled systems called “small-world” networks (Watts & Strogatz 1998), (Collins & Chow 1998), have been used to successfully describe the interactions of systems that can be highly clustered, like regular lattices, yet have small characteristic path lengths like random graphs.

Regular networks are also known as “large-world” networks. They are highly clustered while the characteristic path length is large, scaling with the typical dimension n of the network. High clustering is appropriate for certain types of robotic tasks when information produced by individual robots is more likely to be used by neighboring robots. For example, a large number of robots distributed amongst a few teams each performing a different task that requires only local collaboration (e.g. one team is drilling and analyzing soil samples from a certain area while another team is assembling a solar panel) would benefit from such an arrangement of the communication sub-networks between the robots. On the other end, if these teams perform tasks that call for global collaboration (e.g. each team maps neighboring areas) then the size of the characteristic path length results in increased delays for the information to flow from one robot to all the others in the colony. This type of communication topology would reduce the efficiency of the collaboration amongst the sub-groups. In this latter case a random-graph type of connectivity would be ideal. Small characteristic path lengths would ensure that the information gathered by each individ-

ual team member would be diffused during a few only communication cycles (re-broadcasting of the data to robots connected on the same Ethernets) to all the robots in the colony. In many cases the members of a robot colony are required to switch from tasks that require primarily local collaboration to tasks that depend on global communication. Instead of redesigning the communication topology each time there is a new task from a regular to a random graph and vice-versa, “small-world” networks can be used. Starting from a ring lattice and rewiring a few edges at random with some probability p we can “tune” the graph between regularity ($p = 0$) and disorder ($p = 1$). Graphs belonging in this intermediate region $0 < p < 1$, shift gradually from regular network to a random network. Both the characteristic path length $L(p)$ and the clustering $C(p)$ of the network can be described as functions of the amount of randomness p .

The implementation of a “small-world” network topology when designing the communication graph of a large colony of mobile robots has the following advantages: 1. The communication overload that each of the robots would experience if all of them were connected on the same network (1 Ethernet, one-to-all communication) is obviated, 2. The amount of clustering for local teams of robots remains almost the same and thus relevant information produced within this team is quickly shared amongst its members, and 3. The characteristic path length is significantly reduced compared to the case of a regular (lattice-like) network, therefore facilitating the fast diffusion of information across the colony when this is necessary for a global task. The critical parameters of such a network are the total number of robots in the colony n , the connectivity dimension k and the degree of randomness p . In most cases n is pre-specified while k and p are the design parameters to be determined in order to bring an initially regular network to its “small-world” form.

References

- Collins, J., and Chow, C. 1998. It’s a small world. *Nature* 393(6684):409–410.
- Fontan, M., and Mataric, M. 1998. Territorial multi-robot task division. *IEEE Transactions on Robotics and Automation* 14(5):815–822.
- Gerkey, B., and Mataric, M. 2000. Murdoch: Publish/subscribe task allocation for heterogeneous agents. In *Fourth International Conference on Autonomous Agents*. (to appear).
- Watts, D., and Strogatz, S. 1998. Collective dynamics of ‘small-world’ networks. *Nature* 393(6684):440–442.