Adaptive fault tolerance in wireless sensor networks

David Irons[1], Fabrice Saffre[2] and Chris Cannings[1]

[1]University of Sheffield
[2]British Telecom
d.irons@sheffield.ac.uk

Systems including wireless sensor networks and amorphous computing involve scattering a large number of nodes (sensors / processors) over a geographical area and then creating local connections to carry out some global objective. For example, monitoring, positioning and reporting back on changes in environmental features; such as heat, light, sound and / or motion. Geographical, cost or size constraints may also imply that the nodes need to be dispersed semi-randomly, implying a great deal of uncertainty about their relative positions and the exact size of the area being covered.

Various research challenges exist in these systems, such as creating strategies for self organisation, energy management, fault tolerance, maintaining security and adapting to cope with geographical features (e.g. walls, hills, buildings). Moreover, these challenges become increasingly important as technological advances lead to smaller and smaller nodes, making them (potentially) more error prone and difficult to position accurately. In this study, we introduce a fault tolerance / energy management strategy that allows the system to robustly self re-organise, when nodes break or run out of power. We then test this strategy on different geographical landscapes.

In order to model these sensors networks, we use random geometric graphs. Here, n nodes are randomly placed on a surface and edges can form between nodes if they are within distance R of one another. Additionally, we assume that each node has a randomly assigned lifespan corresponding to time taken for the node to run out of power / break down.

We then test the following strategy. Suppose each node can take one of two states, 'active' or 'hibernated', with the hibernated state consuming less power. Additionally, suppose any active node is capable of sending all other nodes within distance r ¡ R into hibernation. Now, beginning with every node in hibernation, the system evolves by individual hibernated nodes (asynchronously) trying to become active (i.e. turning active if there are no existing active nodes within radius r). Whenever the lifespan of an active node is exceeded, it disappears from the system, allowing the opportunity for hibernated nodes to take its place.

In order to score the integrity of the system, we compare the evolution of such networks to an equivalent system where all nodes are active. In particular, we compare various statistical properties of the network at regular time intervals; including (1) the proportion of active nodes in the largest component, (2) diameter and (3) area covered by the largest component. Such performance measures were chosen since these systems may need to adequately monitor a geographical area (over a period of time), with network connectivity being essential for communication of data / results. We find that, for a large range of parameters, the new strategy presented here is more robust, in that the integrity of the system is maintained over a longer period of time. In particular, robustness increases as the initial number of nodes (n) increases. We believe that as nodes become cheaper to mass produce, having a large number is a feasible strategy for these systems.

As well as testing our strategy on flat two dimensional surfaces, we also extend our approach to more complicated three dimensional geographical landscapes with the additional constraint that nodes can only communicate if they are in line of sight of one another.