Sensor Networks of Intelligent Devices

C M Roadknight and I W Marshall

***Abstract*—Making wireless sensor networks work efficiently and effectively is a key technology challenge for the 21st century. We show that novel decentralized evolutionary algorithms, that we proposed for cluster management, are a potential means of automating the management of sensor networks. Specifically we show the algorithm can enable preferential and reliable delivery of the most important data using only high level user priorities as inputs.**

I. Introduction

The trend for decreasing size and cost of networking devices has created opportunities for networked systems in many new areas. In the field of sensing it has created a new class of networked systems called Sensor Networks [6, 5]. These consist of a large number of battery-powered devices, each with sufficient hardware to monitor one or more variables and send and receive the readings for these variables to other devices. This basic hardware outline gives scope for very complex systems of interacting devices that can carry out sophisticated sensing tasks in a much more robust, economic and effective manner than conventional systems. Given continued miniaturisation and cost reduction, it seems certain that the field of sensor networks will become more accessible and more prevalent as an area of research and application. There are therefore strong reasons for research into future applications, most specifically into how large sets of devices are managed, optimised and deployed. We think some of the key issues include battery efficiency, routing and how Artificial Intelligence can be used to facilitate device autonomy. An alternative motivation for studying sensor networks is that they provide a simplified research environment in which to explore critical topics in the more wide reaching field of pervasive computing. Many of the more complex issues that hold back the widespread adoption and deployment of pervasive computing (e.g. security, charging, interoperability) can be minimised by assuming a single owner of all the devices. This allows us to focus more precisely on key research issues concerning automated configuration and maintenance. This paper provides early solutions to how sensing device can make local decisions on their sensing behaviour, how the devices can make decisions on what they should be doing in a scaleable, hands-off way. Flexible solutions to these problems will enable any sensor network task to be provisioned efficiently given a set of user defined constraints.

Manuscript received October 15, 2003.

C. M. Roadknight is with the BTexact, MLB pp13f, Adastral Park, Martlesham Heath, Suffolk IP5 3RE UK (phone: +44 (0)1473 606523; e-mail: christopher.roadknight@bt.com).

I. W. Marshall., is with the BTexact Adastral Park, B54, Adastral Park, Martlesham Heath, Suffolk IP5 3RE UK (e-mail: i.w.marshall@bt.com).

II. Sensor Networks

Wireless sensor networks are becoming a powerful tool for monitoring a range of diverse situations [3, 8]. While the devices themselves are mostly still in the prototype stage [7] the theory surrounding these devices is a fast moving area of research. Ad-Hoc networks are a collection of mobile devices with wireless networking capability that may form a temporary peer to peer, multi-hop network without the aid of any established infrastructure or centralised administration. Sensor networks have much in common with this network paradigm, but have some unique properties, including;

1. Measurement. The primary purpose is to make and deliver measurements (eg. motes [6] recording temperature levels). The field of active networks [12] has much to offer in this area, methods for moving processing into a fixed network have been available for some years and the argument for a similar approach for lower bandwidth devices is stronger.

2 Limited power supply: Many approaches are being proposed that optimise the use of limited resources [e.g. 11].

3. Lack of persistence: Devices in a wireless sensor networks are untethered and have a degree of unreachability. This will mean that protocols and algorithms developed and optimised for fixed networks will not be optimum [1].

4. Remote management.

5. Local Interactions: Provision of localised algorithms [4], local code acting to achieve a global aim, appears to be one solution to decentralised wireless network management.

6. Reprogramability: Devices have bi-directional communication to other devices, this is a means to reprogram and update device software locally.

Matching user requirements and budget to capabilities is an important area of action. We propose that making realistic simulation software that provide a virtual experimental testbed is invaluable to end users in scaling the proportions of their experiment, in giving realistic options to what can and cannot be done on chosen budgets. The requirements of the users will have much more flexibility than with fixed networks. They will want options to make real time changes to measurement regimes, to modify granularity over important periods of time as they arise, to move devices around to monitor important regions more closely. Most sensor network research uses offline analysis of data [2] this can often mean that a whole

year passes before modifications, improvements and fault rectification is made.

III. The Self-Organising Collegiate Sensor (SECOAS) Network Project

SECAOS [14] involves a new way of thinking for coastal oceanographers, marine scientists, managers and engineers; a change in direction, moving away from large expensive sensor packages to small, self-organising, collegiate systems. The advantages of this are numerous: large packages are expensive to build, maintain and deploy; they need to be protected against trawlers, they must be recovered (usually essential, to retrieve the data). Rarely are more than two or three such systems available for a study (usually only one!), so site selection can be problematic. They have many expensive sensors, high precision and accuracy, low temporal drift and compensated for temperature and pressure effects. Ironically, due to temporal and spatial variability in natural coastal systems, high precision is not necessary for many parameters. For example Vincent et al, [13] examined the uncertainty in measurement of suspended sediment concentration by an optical backscatter sensor (OBS) resulting from the effects of time-varying sediment size and concluded than ±10% was the best that could be achieved. Currently oceanographers don't understand many aspects of sandbank dynamics. An alternative is to use a network of sensors to measure the spatio-temporal landscape. This system is robust even when nodes are destroyed or the network topology changes. Furthermore, nodes can be easily added and reconfigured. While this approach has always been desirable, the availability of low-cost microprocessors and radio devices have made this approach more feasible. The measurement packages designed for the SECAOS project will be small, cheap, simple sea-bed Packages (level 1) that are scattered over an area of oceanographic interest; typically 30-50 Packages, each with the ability to communicate with each other via links to floating buoys. A smaller number (3-5) of more complex surface buoys (Level 2) would communicate, control, monitor and organise the Level 1 Packages, interact with the other Level 2's (radio) and with the outside world. Sensors should be relatively cheap (so we should begin with a basic suite consisting of pressure sensor, optical backscatter sensor (OBS) and a thermistor,) and require low power.

IV. Evolving Decision Making Functions For Device Autonomy

We believe it is desirable that sensor network devices have as much autonomy as possible. Given the mobility of devices and increased likelihood of failure, devices that can learn, adapt and make sensible decisions for themselves will be far more robust and their resulting measurements should be more reliable. As the number of devices increase, as envisaged in "Smart dust" [6] type research, the idea that each devices behaviour can be remotely managed on an individual device level become untenable. We have previously proposed and simulated evolutionary algorithms software deployment on an active network [10, 9]. In this paper we show that a similar approach could also be used effectively on a sensor network. The model devised is that of a simple Ad-Hoc sensor network, a network of devices with the task of gathering data from a site while also optimising their battery usage. Each device within the network is given the capability to move around geographically following a bounded random walk. Each device could be active or inactive during each time window and each device has a battery that was used and monitored, and can be trickle recharged with periods of inactivity. Data collected by the sensors had to be routed to some central data 'sinks'. To enable efficient routing to the sink, nodes carry out an assessment of their nearest neighbours and discover a hierarchical level for themselves based on the number of hops to the sink. Firstly every node will send out a message looking for an acknowledgement from a sink, this message has a maximum range. Every node that is within range of a sink then becomes a level 2 node (sinks are level 1). Every node that is not level 1 or 2 then sends out a message asking for replies from level 2 nodes, if they get one they become a level 3 node. The remaining nodes then request and acknowledgment from a level 3 node, if they get one, they become a level 4 node, and so on until the maximum hop number is reached. Nodes then forward to the NEAREST node that is at a lower level than itself. There are three qualities of data, these could be 3 types of data (e.g. humidity, light levels, temperature) that the user had decided were 3 different levels of importance. Each device puts every item of data sensed or received via forwarding into a 'First in-First out' queue. This data is then acted upon (deleted, combined or forwarded). The queue length is initially set at 50, if sensing puts the length at above 50 then data is dropped. Nodes are able to carry out one role per epoch. Sensing, Forwarding, Deleting, Compressing or Inactivity. They decide on what state to be in during each epoch based on a set of values, these are initially random but are modified. E.g. A node may have the behaviour values:

P(Sense) = 20%, P(Forward) = 50%, P(Delete) = 2%, P(Compressing) = 3%, P(Inactive) = 25%

Therefore it will be Sensing 20% of the time, Forwarding 50% of the time and so on. These values are modified in two ways, local rules and evolutionary, fitness based rules. Local rules act on these values based on the internal values such as battery level and queue length. Eg.

If battery < 100 then P(forward) = P(forward) * 0.95 and P(sense) = P(sense) *0.95

The rules can be as complex as needed, but the important part is that the most suitable values for variables within these rules can be evolved, taught or learnt. Fitness-based rules use a fitness indicator to decide if the P value for the node should be changed, either randomly or by copying the values from a neighbouring node. This provides a longer-term selection process for the best combinations of P values. Nodes are given fitness rewards for sensing data and forwarding data, depending on the quality of that data, they are also given

penalties for deleting data or dropping it due to full queues. The initial settings for the node are randomly decided. Figure 1 shows a snapshot of the network. The different shades of the nodes represent the 5 behaviours and the values beneath the nodes represent node routing level and fitness.

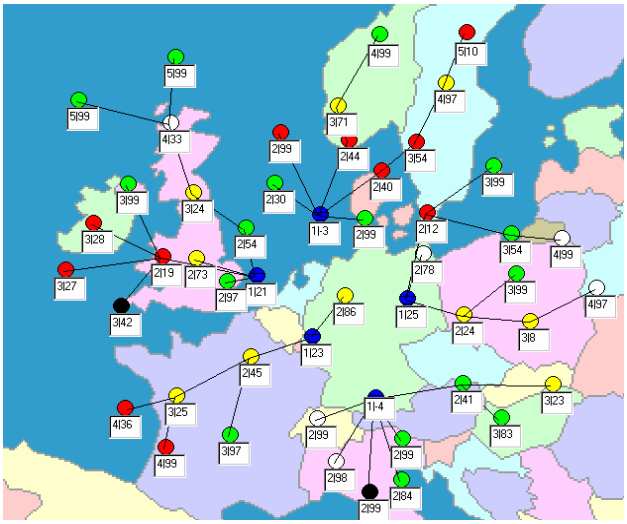

Figure 1. Snapshot of Ad-Hoc network. Node Routing Level | Fitness.

Sensing = Circles, horizontal line. Resting = Circles, no lines. Sink = Square, horizontal line. Compressing = Black square. Routing = Square, vertical line. Deleting = Circle, vertical line

The importance of the local rules is shown in figure 2. When the local rules are switched off, the behaviour of the network nodes is significantly altered. Less queue management is carried out and much more relaying. Without the local rule though, less data is received at the sink and there is no difference in the amount of measurements for the three different measurements. A set of experiments were carried out to show how the behavior of this sensor network could be beneficial and suitable to a sensor network user. For instance, the quality of any service provided must be assessed.

Figure 3 shows how a decrease in the rate at which devices can transfer data effects the success rate of the three different data types. A decrease in maximum transfer rate could occur in several ways, changes in environmental conditions or falling battery power being the most likely. Decrease in performance seems to be dependent on the importance of the three data types. High priority data decreasing from 100% to 90%, medium priority data decreasing from 97% to 63% and low priority data decreasing from 95% to 46%.

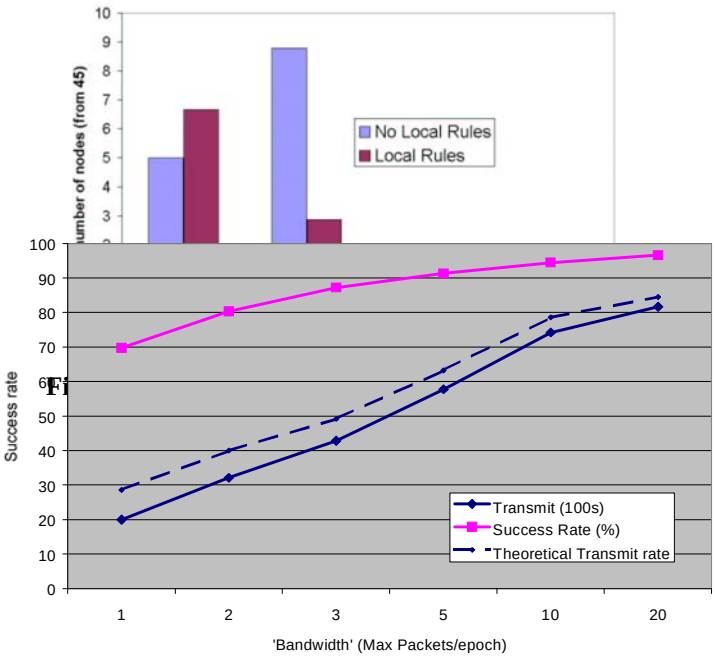

Figure 3. Number of packets sent and percentage dropped as ‘bandwidth’ increases

This would be a desirable feature given that the less important data is dropped preferentially when the network is more ‘stressed’. This is achieved entirely by the delete function within the node. When the node carries out the delete function it looks at the ‘importance’ of the next reading in the queue and decides if to delete it or not. It is programmed to be more ruthless to less important readings, thereby freeing up places in the queue for more important entries

We were interested in node behaviour, particularly how much time each node spent sensing and relaying. Figure 4 shows how, when the number of nodes that that target node acts as a conduit for increases, the number of sensing epochs decreases, with the exception of when the nodes involved are adjacent to the sink, when the number of sensing epochs increases slightly. This is shown for 5 runs with different random number seeds. In other words:

A. The less nodes that I am a conduit for the more sensing I do.

B. The higher the % nodes that are adjacent to a sink the more sensing I do

(0%, unless otherwise shown).

Some nodes are obviously sensing more than others, creating a system where some points are being more regularly monitored than others. Also nodes adjacent to sinks have the benefit of a constantly ON receiving node. Sinks do not suffer from battery depletion like other nodes so are always available as receivers of data.

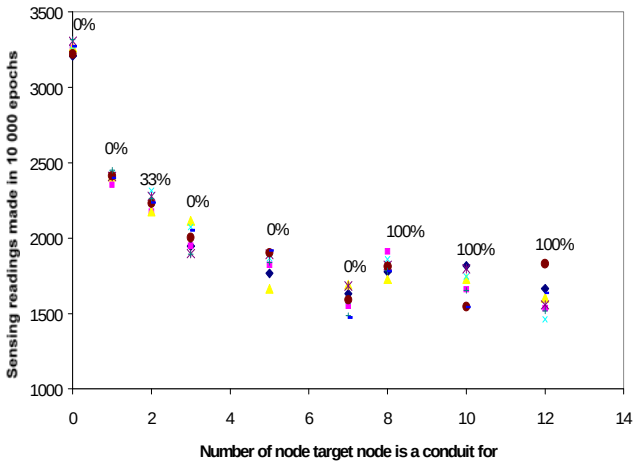

Figure 4. Effect of position in network on node behaviour

The variations in sensing can be explained by the fact that nodes that act as conduits need to spend more time in 'relay' mode to cope with the increased packet rate. While 'hub' nodes sense less the fall in sensing is not as severe as to make the nodes useless as sensing devices, regardless of how loaded they are. This graceful degradation in sensing performance would be key in any real world implementation. While the amount of sensing decreases when nodes are moving, the characteristics stay the same. The decrease in sensing can be explained by a lack of reliable connectivity. The static network is designed so that every node can reach every other node, this is not the case when nodes are moving. The final experiment shown here demonstrates the complex nature of the fitness function when coupled with 'local rules'. Nodes are rewarded or penalised when they carry out one of the network functions. For instance, every time they sense they are given a reward, which will influence the long term survival of their simple genome.

V. Conclusions

Every device in a sensor network needs some kind of intelligence. This may be a simple set of rules about when they sample, what to sample or it maybe be something more sophisticated and complex that takes into account internal and external conditions to make a decision about it's actions. In this paper we introduce some initial results for device intelligence that is constructed out of simple modifiable rules coupled with fitness function based adaptation. This research was carried out to demonstrate possible solutions to providing device autonomy. Firstly showing that embedding code within each node to carry out some of the decision making usually associated with the human user of the network Sense Fitness

Secondly, the behaviour of the network as a whole is shown to disply attractive features like load balancing and quality of service when topological effects are investigated. It is encouraging that while the individual nodes are acting in a self-optimising manner, the network as a whole is displaying characteristics that are robust and scalable. The importance of the learning techniques is then demonstrated. A degree of self-regulation is shown in using multiple adaptive techniques, where inefficiency or naive settings in one aspect of the learning algorithm can be regulated by a different aspect of the algorithm.

Taking this research further will involve fine tuning the learning algorithm for different scenarios and carrying out further investigations of how the different learning approaches interact when they are carried out in parallel. Automating the reward and penalty functions, so they too and configured in a hands off way will also be essential. Implementing the decision making solutions onto real sensor network devices, this will be carried out as part of the SECAOS project [14]

Each sensing scenario will have it's own very specific characteristics. Mobility of devices, time span of the sensing task, inhospitability of the environment. For example fish move very quickly, glaciers very slowly so the optimal algorithms and application of sensor networks for each task will be different. Nodes must adapt, without user intervention, to carry out the task efficiently and effectively.

References

[1] Balakrishnan H., V. N. Padmanabhan, S. Seshan, and R. H. Katz, "A comparison of mechanisms for improving TCP performance over wireless links," IEEE/ACM Transactions on Networking, 1997

[2] Benton J., J Fuhrer J. Gimeno, B.S., Skärby, L., D. Palmer-Brown, G. Ball, C Roadknight, & G Mills. An international cooperative programme indicates the widespread occurrence of ozone injury on crops. Agriculture, Ecosystems & Environment 78, 19-30, 2000.

[3] Brooks T. Wireless Technology for Industrial Sensor and Control Networks. Sicon'01. P73-77

[4] Estrin R. G. D. and J. Heidemann. "Scalable coordination in sensor network." In Proc. ACM/IEEE MobiCom, 1999

[5] Hill J., R. Szewczyk, A. Woo, S. Hollar, D. Culler, and K. Pister. System architecture directions for networked sensors. In Proceedings of the 9th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, pages 93- 104, Cambridge, Massachusetts, Nov. 2000

[6] Kahn J. M., Randy H. Katz, and K. S. J. Pister. Next century challenges: Mobile networking for smart dust. In Proc. ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom 99), Seattle, WA, August 1999

[7] Levis P. and D. Culler. Mate: A tiny virtual machine for sensor networks. In International Conference on Architectural Support for Programming Languages and Operating Systems, San Jose, CA, USA, Oct. 2002.

[8] Mainwaring A., J. Polastre, R. Szewczyk, D. Culler, and J. Anderson, Wireless sensor networks for habitat monitoring., ACM International Workshop on Wireless Sensor Networks and Applications, WSNA '02 (Atlanta, GA), September 2002.

[9] Marshall I.W.and C.Roadknight "Provision of quality of service for active services" Computer Networks, April 2001

[10] Marshall I.W.and C.Roadknight, "Adaptive management of an active services network", British Telecom. Technol. J., 18, 4, pp78-84 Oct 2000

[11] Schurgers C, V Tsiatsis, S Ganeriwal and M Srivastava. Optimizing Sensor Networks in the Energy-Latency-Density Design Space. 70 IEEE Transactions on mobile computing, vol. 1, no. 1, january-march 2002.

[12] Tennenhouse D L., J M. Smith, W D. Sincoskie, D J. Wetherall, G J. Minden. A Survey of Active Network Research. IEEE Communications Magazine 1997 vol 35. No 1. P80-86.

[13] Vincent C.E, S.J. Bass and J.M. Rees (2003) Uncertainties in the estimation of suspended sediment concentration due to variations in sediment size. Proceedings of Coastal Sediments '03, Clearwater, Florida. (on CD, pdf, 11pp)

[14] http://www.adastral.ucl.ac.uk/sensornets/secoas/