Annealing Sensor Networks

Andrew Jennings and Daud Channa

Electrical & Computer Engineering School, RMIT University, Melbourne, Australia
ajennings@rmit.edu.au, S3047387@student.rmit.edu.au

Abstract. With a continuing improvement in the capabilities of intelligence per unit of energy, we should reconsider the organisation of sensor networks. We contend that solutions should be model-free, locally based and need to be highly dynamic in nature. Here we propose an approach inspired by simulated annealing. In the context of several application scenarios we explore the potential for adding intelligence to sensor networks.

1 Introduction

Sensor networks [1] are envisioned to become an integral part of our lives. These networks are being applied to provide various tasks such as surveillance and monitoring systems for commercial and military applications. Applications are being developed to gather process and utilize the information from the surrounding environment as required. These requirements have kept challenging researchers in the design of better architectures and protocols for sensor networks. We now have some early deployment of sensor networks, showing that we have successfully established the basic protocols. These follow two main directions: clustering of nodes [2], and synchronised sleep cycle networks with a flatter structure [3]. Now the main challenge is to establish a wide range of application systems. To deploy applications we need methods of coordination that are efficient both in delivering services and conserving energy.

The growth and advancements in technologies and the constant reduction in the size and cost of Micro Electro Mechanical Systems (MEMS) have given rise to a whole new dimension of networking which involves sensors and actuators that are quickly deployable and self organizing. They have resulted in a new dimension in network computing, namely pervasive sensing and control. The ongoing rapid advancements, developments, and research in the sensor and actuator networks only leaves one to foresee that they will soon intervene and associate into all living habitats of humans and their surrounding environment.

In most scenarios the network must be functional over long periods of time, it is crucial for the operation, management and continued lifespan of the network to control the behaviour/reaction of the sensors and actuators to the different occurrences of events in the network. The sensors in these networks are limited in energy, memory and computational facilities, while generally the actuators have an ample supply of resources as their mobility can enable them to recharge thus utilizing their resources to the fullest without energy constraint.

The deployment and maintenance of the nodes must be cost effective, because it will be unfeasible to configure these large networks of small devices. The sensor nodes along with the actuators must be self organizing and provide a means of programming and managing the network as a whole, rather than administering individual nodes and actuators.

R. Khosla et al. (Eds.): KES 2005, LNAI 3684, pp. 581–586, 2005.

2 Status and Scenarios

Although it is not yet clear which applications are viable for sensor networks, we have selected three scenarios to motivate our work. They serve to highlight the issues that we now face. We aim to create solutions for these situations.

2.1 Pedestrian Crossing Guard

We aim to improve the safety of pedestrian crossings using sensor networks. How might we prevent the running down of pedestrians?

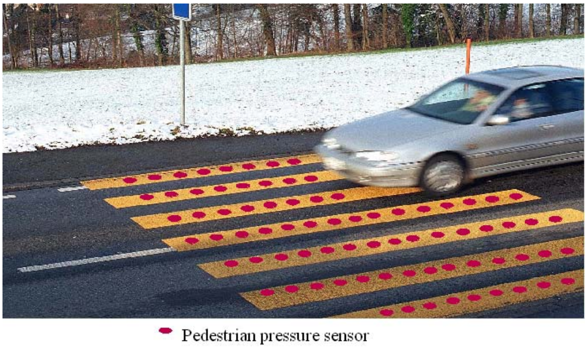

Fig. 1. A proposal for an instrumented pedestrian crossing

One possible solution is to use an array of pressure sensors on the roadway surface to detect pedestrians walking. This would be in addition to proximity sensors and visual surveillance [5], as illustrated in Figure 1. Through combination of sensor readings, we can improve the accuracy of recognition. If we want to make use of the sensor readings then we need confidence that there is a low probability of false positives - otherwise car drivers will not accept the system.

With a reliable method of pedestrian detection, we can perhaps move to the next level with these systems. If we can reliably detect a car traveling at a speed likely to result in a collision, then the crossing system can intervene and communicate with the car - potentially it could also override car braking systems. This would bring the car to a halt. We have the possibility of eliminating the possibility of cars contacting pedestrians at crossings. There may also be a role for coordinating with robot teams to enable more active monitoring of pedestrians – here we need development of team behaviour [6]

2.2 Animal Counting

Environmental monitoring was one of the earliest motivations for exploring sensor networks. A typical task is the estimation of animal populations. We would like to know how many of a particular type of animal are within a geographic area. In contrast to urban applications, this setting is very demanding in terms of energy management. Note that the estimation of populations is more difficult than simply tracking animals - we need some confidence of the identity of animals. Is the set of readings for a single animal, or two that are within the same area?

2.3 Perimeter Surveillance

This is a classical application of sensory technology. We have a perimeter that we wish to patrol, with video and movement surveillance. To augment this, we would like to deploy proximity sensors to improve accuracy. These scenarios can give us a framework to consider sensor network application approaches. They provide challenges and a range of difficulties. All are real applications that may have some prospect of widespread deployment. At this stage of development of the field, it is important to focus on feasible applications to focus the research.

3 Energy and Intelligence

One of the central tenets of sensor networks has been the need to keep nodes simple and careful in their use of energy. We could not, for example, implement the full TCP/IP stack on sensor nodes. This would be a waste of energy, since the nature of the communication is quite different.

Progress in battery technology is painful and very slow. But when we consider intelligence per unit of energy, then progress is quite dramatic. So we should be more open to incorporating intelligence in the sensor node, as long as that results in significant energy conservation.

3.1 Local Resolution

One of the original proposals for sensor network protocols was "directed diffusion" (DD) [4]. It is a robust protocol that can work in very tough environments. Even with extensive network breakages, it will continue to operate. As Figure 2 illustrates, "interests" are propagated to areas of the network, and "gradients" are used to reinforce the successful delivery of packets across the network.

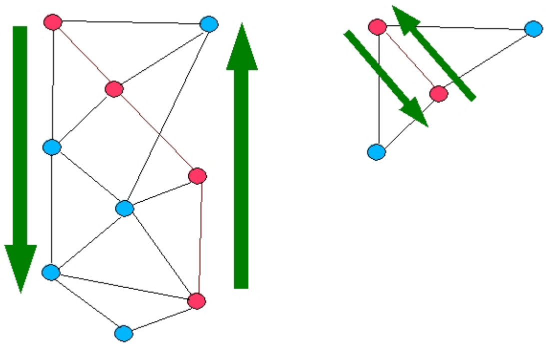

Fig. 2. Directed diffusion (network wide) versus local resolution

Given the constraints placed on directed diffusion, it is an appealing solution. But how might it change if we allowed local processing, rather than propagating results across the broader sensor network? DD assumes that we have to send this outside the sensor network, but with local processing we can avoid the energy consumption of network wide reporting.

This creates both a need for local algorithms that can actually resolve sensor data, and also a means of coordination. There are clear energy advantages in local resolution.

Similar difficulties arise in the location of mobile sensor nodes. The author in [7] has presented the use of simulated annealing for this problem. Here we are concerned with organisation and messaging for fixed nodes.

3.2 Model Based, or Model Free?

If we want to improve monitoring, then perhaps a more accurate model of the context will help? In the case of the pedestrian crossing, we could develop a tracking model. For example, Kalman filtering could aid following an individual through the network. But if we incorporate this model, what will happen in non-modelled cases? How will our model-based approach react when a person falls over, and lies stationary on the crossing? In the worst case, we might decide that the crossing is clear, and let the cars proceed.

Similarly, in surveillance of a perimeter, we might improve accuracy by statistical training to detect people walking across the field. But will we detect somebody crawling across the space?

Once we fix a model for the sensing environment, defining the range of possible targets, then the task of constructing the sensor network is reduced to optimisation. There is no need for further intelligence. So the real challenge for sensor networks is how to deal with unusual data. Consider surveillance when a bunch of leaves falls to the ground. Do we raise an alarm or simply log the event for further processing? If we log it, what priority do we give to the event?

This is a familiar problem in AI, bringing us to the very familiar challenges of semantics. How do we deal with images that do not have familiar content? How can we go beyond simple statistical pattern matching? These are very difficult, but also very important problems.

3.3 Dynamics

Consider the problem of tracking (and identifying) an animal that gives us unusual readings. Perhaps it is of a size that we have not encountered, or it genuinely is a new entry to the region. Clearly this is important, and we would like to track its trajectory. But in order to do this, we need to estimate velocities and alert the relevant part of the sensor network. Once we have lost contact, it will be difficult to sustain the identity. Remember, we are interested in counting animals, so identity is important. Clearly we need application protocols that can deal effectively with highly dynamic situations.

4 Annealing Sensor Networks

The simulated annealing algorithm is a successful method of searching for optimal solutions in complex spaces. Most importantly, it is *model free*: any problem can be formulated as an annealing process. In analogy with the process of annealing crystals, it has an associated temperature. At high temperature, large parameter changes are possible, but as the process cools only smaller changes are possible.

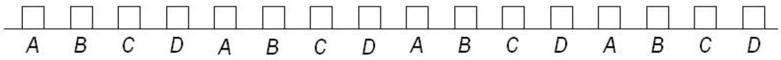

Fig. 3. Activation cycles (sleep cycles) for a single node

We propose an approach to organising sensor networks that is inspired by simulated annealing. Regions have a temperature, which indicates the intensity of sensing. Figure 3 illustrates the sleep cycle of a single node. At a low temperature, the nodes cycle only at A, but as the temperature increases we also cycle at B,C,D progressively. Since these cycles are divisions of the fundamental cycle (the A cycle), these schedules do not conflict.

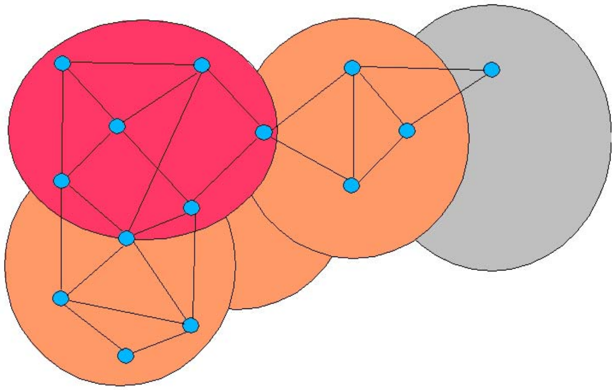

Fig. 4. Network temperature heating process

In the event that a node encounters an unexpected stimulus, it can cause a local rise in temperature. If several nodes in a region send this message, then a rapid rise in local temperature can take place. We can allow this temperature to spread rapidly in space if we desire, or rapidly decay. In accordance with the physical analogy, local heating cannot spread vast distances without decay. Figure 4 illustrates the process.

To effectively make use of annealing sensor networks, we need local resolution of sensor information. For example, in the case of animal tracking, a local decision is needed on the temperature response. Note that an identification is not needed, but only local decision making. We are investigating how to provide this on the typical processors used for sensor nodes. It certainly seems possible to accomplish this computation on the nodes. Of course over time, we can expect the intelligence/energy quotient to continue to grow.

Annealing sensor networks (ASN) are a development of directed diffusion networks. There are some important differences. The adaptive sleep cycle provides for rapid response. Local resolution of control is an important difference. Where directed diffusion incorporates routing, ASN's only advise routing.

5 Discussion

We have proposed an approach to model-free local behaviour for sensor networks. Given that we have no local model of expected behaviour, how can we achieve local resolution? Each node can keep a statistical database of patterns it has encountered. When patterns within a statistical tolerance appear, this can trigger the appropriate behaviour.

Fully distributed control of sensor networks in this manner raises some important new issues. How do we maintain the currency of statistical data? How can we make changes to behaviour whilst ensuring network stability?

It is interesting that classical problems of semantics come to the forefront when we want to further explore sensor networks. Here the resources we have to bring to the problem are limited. We have an unlimited source of data, through lifelong observation of the world through the network sensors. Potentially we can bring vast computation to the task, through recording and processing offline. But we are limited in human intervention. This leads us to explore computationally intensive approaches. Perhaps we should consider the task as "data mining for sensor organisation methods".

References

1. Deborah Estrin, Ramesh Govindan, John Heidemann and Satish Kumar "Next Century Challenges: Scalable Coordination in Sensor Networks" *In Proceedings of the Fifth Annual International Conference on Mobile Computing and Networks (MobiCOM '99), August 1999, Seattle, Washington.*
2. Wendi Heinzelman, Anantha Chandrakasan, and Hari Balakrishnan "Energy-Efficient Communication Protocols for Wireless Microsensor Networks", Proc. Hawaiian Int'l Conf. on Systems Science, January 2000.
3. Wei Ye and John Heidemann and Deborah Estrin "An Energy-Efficient MAC Protocol for Wireless Sensor Networks" *Proceedings 21st International Annual Joint Conference of the IEEE Computer and Communications Societies, 2002*
4. Chalermek Intanagonwiwat, Ramesh Govindan and Deborah Estrin "Directed Diffusion: A Scalable and Robust Communication Paradigm for Sensor Networks" In Proceedings of the Sixth Annual International Conference on Mobile Computing and Networks (MobiCOM 2000), August 2000, Boston, Massachusetts.
5. Christian Wohler "Autonomous in situ training of classification modules in real-time vision systems and its application to pedestrian recognition" Pattern Recognition Letters Volume 23, Issue 11 Pages: 1263 - 1270 (September 2002)
6. H. Van Dyke Parunak, Sven. A. Bruekner & James Odell "Swarming pattern detection in Sensor and Robot Networks" Altarium Institute, Working Document
7. Martinson E. and Payton D. "Lattice Formation in Mobile Autonomous Sensor Arrays", HRL Laboratories, LLC. 2003-2004