The development of a robust, autonomous sensor network platform for environmental monitoring.

L. Sacks[1], M. Britton[1], I. Wokoma[1], A. Marbini[1], T. Adebutu[1]

I. Marshall[2], C. Roadknight[2], J. Tateson[2], D. Robinson[2], A. G. Velazquez[1/2]

[1] Department of E&EE, University College London, UK., [2] Btexect Technologies, UK.

Abstract:

This paper describes an approach to the approaches being explored for a Sensor Network platform being developed for the DTI/NextWave technologies programme. The approach being adopted is to develop the system as a community of devices which use self-organising techniques to provide key functions. The devices are, largely, based on commodity technologies, thus providing a low cost basis. We give an outline of the approach and project and illustrate the techniques being developed with specific functions for: control, management, data retrieval and data quality control. The target application is off shore sea shelf monitoring; but the techniques being developed may be applied to a range of problems.

1. Introduction

Developing complete sensor systems for environmental monitoring in harsh environments is currently a complex and expensive exercise. There is a trend towards leveraging commodity technologies to enable the construction of cost effective sensor platforms, enabled by recent growth in wireless communications and microprocessor controlled consumer devices. This leads to small, low cost, limited functionality devices working in a cooperative networked mode, to form a single system. This model brings cost advantages and extra science value. By constructing a *network* of sensors, a greater area may be covered allowing better spatial resolution. Further, redundancy in components may be introduced thus providing more robust and longer lived operations. These features are highly advantages in applications covering large, hostile environments such as: glaciers, volcanoes or off-shore sea beds.

This paper presents the development of such a sensor network (SN) platform for the monitoring of environmental impact on a coastal sea bed of a wind farm. Wind farms are seen as a key feature of the governments' sustainable energy policy. It is critical to ensure that their placement does not produce negative, environmental impacts. The complex interplay between the: oceans currents; wind; coast line and off shore features (e.g. sand banks) is poorly understood; but it is critical that natural costal defences such as sand banks are not disturbed or destroyed inadvertently. Thus there is a great need for continuous measurement of off shore features in such a context.

This approach and application requires trade-off between design and capability. Using low cost devices entails limited functionality and performance. Target environments considered here are, by their nature, highly turbulent which entails that the sensor platforms have to be able to adjust to the environment on a continuous, though hard to predict, basis. These problems have led us to explore a new approach to the system engineering. In this approach the system is composed of a *peer-to-peer* system, supporting intrinsically networked algorithms; which derive their full functionality from being symmetric, cooperative and self-organising. This contrasts to current approaches of confining the functionality to individual nodes and considering the sensor network as a slave device. To achieve this, our approach is to model the algorithms on behaviours found in natural biological systems. This paper presents this overall concept and some example details of the algorithm design.

2. Application Context & Motivation

The work reported has been carried out in preparation for implementation under the DTI NextWave [I] technologies program, for the project SECOS [II], part of the "Centre for Pervasive Computing in the Environment". The wind farm development located on the Scroby Sands [III] sand bank provides the key requirements capture for the SN development. The location of wind turbines on sand banks can

produce highly complex turbulent flows both locally, producing sand shift and scouring of the turbine bases; and on a wider scale having impacts on along-shore sand banks and sea barriers. Measuring the ocean in this environment is currently expensive. Typical landers will cost £100-200K to build including: sensors, acoustic releases etc. and have high deployment and retrieval costs (as the lander must be precision located). A single lander can only measure the environment at a single point; it is vulnerable to destruction by storms, burial by moving sand waves and accidents from trawlers – as such it is a 'one shot' proposal. In science terms, the key problem is that the current technologies do not give good coverage in either time or space. For most of each deployment little happens so power and data storage is wasted; operational cycles are always compromises between the need for high resolution measurements during the active periods (usually storms) and the uncertainties in the frequency, timing and severity of the active periods. Typically, most monitoring systems undergo fixed sleep-sample duty cycles which are good enough for monitoring regular features (e.g. tides) but cannot reactively monitor transit effects. Turbulence is far from regular through space and there is clear need for high special sampling to measure features such as sand: drift, scouring and build up.

The approach we are developing is to develop a SN platform consisting of a large number (30-50) of low cost (<£1000) nodes. Each node has basic functionality; a small μ-processor, low rate communications and sensing capabilities. Sensing capabilities will focus on features such as, optical backscatter – a measure of sediment load in the water – pressure, and temperature. In addition to the general sensor nodes, there will be a few master packages which can measure other parameters (e.g. surface pressure) and for communications with the science base-stations. For deployment the nodes could be scattered from a boat, if they are capable of determining their own location. Key amongst the technical requirements is the need for the nodes to be able to communicate between them selves on a nearest neighbour basis. This can be achieved using surface-floating tethered buoys supporting low cost, low power radio communications such as 802.15.4. operating in the IMS spectrum with low data rates (250kbs) and high node densities (~250) and low power demands (*e.g.* 30μW for a 1000/1 sleep/transmit cycle). For more general applications sonar communications are a viable alternative.

This approach brings its own set of functional requirements. The nodes must be able to configure them selves for monitoring frequency in space and time, to be able to condense data and communication it back to the base station, to be able to substitute for each other in case of failure, and be able to locate them selves. The nodes must respond to local changes in the environment (both the monitored environment and each other) as autonomously as possible. Thus there is a need for a general, localised control system with the whole system manageable by the scientist / operator. It is not possible for the operator to control the each node individually so the system must be manageable at a high level.

3.   A Kind of Operating System

From the above discussion it is clear that nodes loaded with classical micro kernel Operating Systems (OS) is inappropriate for this kind of application. Such OSs' take substantial resources and their functionality is generally confined to single nodes. In this context, we require a light-weight service kernel, supporting distributed algorithms. The best models we currently have for such systems can be found in the biological world in which systems at all scales (from inside the Cell, through to social systems) can be seen to derive their functionality as much from how components interact as from the intrinsic functionality of the components them selves; in a. fully dispersed architecture. The idea that the 'algorithms' which a system is performing is derived as much from the way things interact as from a specific algorithm is key behind the concept of Emergence. That concept of control occurring through dispersed cooperative action is the key notion behind evolving self-organising systems [IV]. Hence the OS framework should be based on concepts of emergent, self-organisation.

The engineering strengths of this approach are two fold. Firstly, by constructing the system out of a number of independent self-organising components in which much of the information and functionality is in the connectivity, we can build a system which is robust against failure or partial functioning of individual components. Secondly, the work performed by individual nodes is small, only being a partial part of the full algorithm, thus the processing & memory demands on individual nodes can be small. We refer to the operating system based on these attributes as the kOS.

We illustrate the functionality of the kOS through the development of some of the key functional

components. These components provide the following functionalities, typically found in operating systems in one form or another: Synchronisation between nodes for coordination of communications and work loading; User management through policies; Identity (by location); Persistent Data storage; Work scheduling & partitioning; Control Optimisation. In order to give a feel how the approaches describe here can be used to provide these functionalities; initial explorations of mechanisms are described below.

3.1. Synchronisation

On any system it is important that the individual components can synchronise their operation. In the SN platform the components not only need a degree of synchronisation but need to be able to adjust the sync-time base to enable energy conserving operation. The synchronisation mechanism has to increase frequency to able to adjust to transient changes in the system – for example when a user management command is entered – and return to the base state. To provide this functionality we have developed a system, loosely based on the concept of fire-fly synchronisation. Fireflies are known to emit

Figure 1. Synchronisation across a Transient event.

flashes at regular intervals when isolated but in a group they entrain the pulsing of their lights to converge upon the same rhythm as that of other fireflies in the group until synchronicity is reached. A non-deterministic distributed mechanism is used in the system where messages are passed between individual fireflies to achieve their coordination [V] which is based on a network of pulse-coupled integrate-and-fire elements [VI]. Each node adjusts its flash intervals through a concave control surface which enables synchronisation lock. In the SN platform, nodes exchange messages in line with the flash interval and 'lock' onto the message exchange phase in order to achieve synchronisation. An illustration of the synchronisation phase is shown in Figure 1 which shows the rate of message exchange before and after a transient event which has cause in frequency increase in the flash interval.

3.2. System Management

Any operating system has to allow the operator to control and configure the system. In this context, controlling an individual node is not meaningful as nodes will exchange roles for a range of operational reasons (power conservation, failure *etc.*). The management system we have explored is based on the concept of a Policy. Policies are high level operational rules which are meaningful to the operator and it is up to the individual components to adjust their behaviour to fulfil these policies. Policies in this context might determine such operational requirements and the length of the experiment or the

Figure 2. Policy Distribution

resolution of the data required by the scientists. This approach has been extensively explored in the world of network and service management [VII] including its use within self-organising systems [VIII]. In this context, the key problem is to ensure that all nodes in the system are aware of the introduction of new policies. The typical procedure for this would be to send the data to each node individually. This would require that the system support a complete routed network infrastructure such as those found on ad-hoc networks. Many ad hoc routing protocols have been devised. Some of the most widely known are DSDV[IX], TORA[X], DSR[XI] and AODV[XII]. A comparison of the performance of these protocols [XIII] has shown widely differing results in the size of routing overhead and can become unacceptable for a network size of 30 nodes. The main problems using these in a SN are the size of processor and memory required and the protocols are not energy usage aware. As an alternative, we have adopted the approach of a modified gossip protocol called Time-Stamped Anti-Entropy (TSAE) [XIV] which provides weak-consistency based information distribution [XV]. The

combination of TSAE and fire-fly protocols has been demonstrated in the context of policy based management for fixed networks[XVI]; the key point is that this never depended on specific routing or communications models, indeed, it performance on random networks was seen to be substantially better than on normally routed networks is illustrated in Figure 2.

3.3. Persistent storage / Data acquisition

The key use of persistent data storage in this context will be for data acquisition. In the context described here, each node will not be able to store much data locally; rather they will send the data back to base stations for archive and analysis [XVII, XVIII]. As discussed above, using classical ad-hoc networking technologies is extremely expensive and not necessary for control and management purposes. Thus the task is to devise a data retrieval system that does not depend on node identity. The approach taken here [XIX] is based on a least resistance flow concept and platforms on the basic synchronisation capabilities described above. Data is, effectively broadcast from each node to its nearest neighbours who, in turn, re-broadcast it based on a cost function. The cost function its self is developed through a form of back propagation from the base stations – who understand that they are the

Figure 3. Up and Down Stream Synchronisation

appropriate destination of the data. Figure 3 illustrates the forward and abackward develoment of the data flow synchronisation through the system as the flow paths develp thorugh time.

3.4. Data Quality selection.

With communications overheads being a major resource cost, it is essential that data is filtered at or close to the sensing nodes so that only changes in measurements are transmitted. The environment being measured is, however, turbulent, with high fractal components in the dynamics. Thus it is hard, if not impossible, to use classical control systems to track the sampling frequencies; a situation that is exacerbated by the fact that the nodes are likely to be 16bit, fixed point arithmetic devices which will struggle with floating point computations. For these

Figure 4. Adaptive Sensing Efficacy

reasons we have been exploring the use of adaptive control algorithms. The one described here is based on the evolution and exchange of short plasmid strings representing the key sensing operations – more or less like a distributed genetic algorithm. Each node continuously evolves its notion of sampling and exchanges its best guess representation with neighbouring nodes. The discreet operations are: Sense, Forward, Delete, Compressing and Idle - each operation being performed for a variable amount of time. The ratio of these operations is determined both through the evolutionary process and locally under Policy control (in line with the nodes local objectives). These rules are applied for each measured quantity individually or in combination (as part of a local data fusion approach) with differential priorities being given to each measurement – under policy control. This approach has been modelled for a scenario with 3 kinds of measurement and a network of 30 nodes. Figure 4 illustrates how each measurement can be differentiated under varying (policy) requirements on retrieved data.

3.5. Location Based identity.

In the applications considered here, the important form of reference is functional in the sense that the key issues are the science values at spatial locations; independently of which node is actually sampling that location. Indeed, it is possible that through the lifetime of an experiment the node(s) responsible for sampling some part of the environment will change due to failure or movement. In our architecture, a node is principally identified functionally by its position and thus it is necessary that nodes can

determine where they are. It will be possible to determine some nodes location absolutely by equipping them with GPS systems or attaching them to wind turbine masts. For the majority of nodes it will be necessary that they can locate them selves relatively to other nodes through techniques such as sonar ranging or – for floating buoys – radio ranging. Nodes can then exchange measurements of relative position with nearest neighbours to locate them selves in absolute term. Some approaches to this have been explored else where [XX,XXI]. Current developments [XXII] focuses on robust operations.

3.6. *Work partitioning*

Most operating systems will have some means of multi-tasking, allowing the system to perform a range of tasks simultaneously. In this context individual nodes will have a single tasking kernel and the multi-tasking aspect of the work will be achieved through partitioning of work between the available nodes. To do this in a self-organising manor, it is necessary that the nodes agree locally, in groups, which work will be undertaken by which nodes. To develop this approach, we have explored the use of Quorum Sensing [XXIII]. Quorum Sensing is a form of intercellular signalling that regulates gene expression in response to cell-population density. In this process, each cell is able to individually sense when a quorum of bacteria, a minimum population, has been achieved and results in a behavioural change in the bacteria. An example of this is bioluminescence which occurs in the organs of deep sea fish due to the bacterium, Vibrio fischeri. Having established a core quorum forming system, the sensor nodes will be able to from them selves into groups dedica:ted to sampling at required rates; performing data fusion, hibernating and so on.

3.7. *Practical Implementation*

The implementation of the kOS in software is driven by a requirement for only minimal OS-like functionality. As the nodes self-synchronise their functional operation, we have no requirement for concurrent task operation and we do not therefore require a multi-tasking system. We can then have run-to-completion tasks that are queued in a simple FIFO manner-greatly simplifying system operation and minimising the code footprint. This also results in predictable task execution and removes any need for inter-task communication. One non-standard feature required for the kOS is in-application reprogrammability, where new sensing instructions are installed in non-volatile memory areas of nodes from remote workstations.

The choice of microprocessor is an important design decision. In most modern microprocessors, processing speed and programmable non-volatile memory is cheap in terms of cost and power consumption, and is generally more than adequate for our purposes. Also, many microcontrollers have features we require such as analogue-to-digital converters (ADCs) and several timers or counters. Driving requirements are therefore for low power consumption devices (including enough low-power modes of operation) and suitable communications and peripheral interfaces, as both networking and sensing are core features of the system. As we have only a simple memory organisation and speed is not a primary concern we also find that a Von Neumann (Princeton) architecture is adequate.

4. Related Work

Researchers at the University of California Berkeley have worked for several years on the design and implementation of a self-contained, millimeter-scale sensing and communication platform in the Smart Dust project [XXIV]. The aim is to design a sensor device that will have some processing ability, wireless communications and a battery supply, while being inexpensive enough to deploy in large numbers. A number of such "motes" have been implemented, and a minimal component-based operating system TinyOS [XXV] has been written specifically for the platform. Other related projects include Manatee [XXVI], involving Bluetooth-based communications and environmental sensing. The University of California LA's Centre for Embedded Networked Sensing [XXVII] conducts a considerable amount of research in this area.

5. Conclusion

This paper has described our approach to the developments of a sensor network platform based on, biologically inspired, self-organising algorithms. At this stage we have established the viability of implementing key functions using this approach. Critical work is now underway to implement the system in suitable packaging and with appropriate technologies for the target applications. This overall approached hold out the possibility of low cost sensor applications for a range of science tasks; and with low management overheads and hardware costs.

6.   References

[I] http://www.nextwave.org.uk/

[II] SECOS: "Self-Organising Collegiate Sensor Networks" http://www.adastral.ucl.ac.uk/sensornets/secos/

[III] http://www.offshorewindfarms.co.uk/sites/scroby-sands.html

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

[V] U. Wilensky, K. Reisman, "Learning Biology through Constructing and Testing Computational Theories – an Embodied Modelling Approach".

[VI] X. Guardiola, A. Diaz-Guilera, M. Lias, C. J. Perez, "Synchronisation, diversity, and topology of networks of integrate- and-fire oscillators", Physical Review E, October 2000.

[VII] M. Sloman, "Policy Driven Management for Distributed Systems", Journal of Network and Systems Management, 1994

[VIII] L. Sacks, O. Prnjat, I. Liabotis, T. Olukemi, A. Ching, M. Fisher, P. Mckee, N. Georgalas, H. Yoshii, "Active Robust Resource Management in Cluster Computing Using Policies", Journal of Network and Systems Management, Special Issue on Policy Based Management of Networks and Services, 2003.

[IX] C Perkins and P Bhagwat, Highly Dynamic Destination-Sequenced Distance-Vector Routing (DSDV) for mobile computers, Proceedings of the SIGCOMM '94

[X]VD Park and MS Corson, A Highly Adaptive Distributed routing Algorithm for Mobile Wireless Networks, Proceedings of INFOCOM '97, pages 1405-1413, April 1997

[XI] DB Johnson, Routing in Ad Hoc Networks of Mobile Hosts, Proceedings of the IEEE Workshop on Mobile Computing Systems and Applications, pages 158-163, December 1994

[XII] C Perkins, Ad Hoc On Demand Distance Vector (AODV) Routing, Internet-Draft, draft-ietf-manet-aodv-04.txt, October 1999

[XIII] J Broch, DA. Maltz, DB. Johnson, Y-C Hu, J Jetcheva, A Performance Comparison of Multi-Hop Wireless Ad Hoc Routing Protocols, Proceedings of Mobicom '98

[XIV] C. L. van Eijl, C. Mallia, "The time-stamped anti-entropy protocol- a weak consistency protocol for replicated data".

[XV] R. Golding, D. D. E. Long, "Modelling replica divergence in a weak-consistency protocol for global-scale distributed data bases", http://www.cse.ucsc.edu/~golding/

[XVI] I. Wokoma, I. Liabotis, O. Prnjat, L. Sacks, I. Marshall, "A Weakly Coupled Adaptive Gossip Protocol for Application Level Active Networks"; IEEE/Policy 2002, Monterey, CA, USA, June 2002.

[XVII] Tian H., Stankovic J., Chenyang J., "SPEED: A Stateless Protocol for Real-Time Communication in sensor networks", International Conference on Distributed Computing Systems (ICDCS 2003).

[XVIII]Deepak G, Krishnamachari B, Woo A, Culler D, Estrin D and Wicker S, "Complex Behavior at Scale: An Experimental Study of Low-Power Wireless Sensor Networks", UCLA/CSD-TR 02-0013.

[XIX]A. E. Gonzalez-Velazquez, L.E. Sacks and I. W. Marshall, "Simple spontaneous mechanism for flexible data communication in wireless ad hoc sensor networks", The London Communications Symposium 2002

[XX] Savvides, A. Park, H. Srivastava, M. 'The bits and flops of the N-hop multilateration primitive for node localization problems' In Proceedings of 1st ACM International Workshop on Wireless Sensor Networks and Applications, Sept, 2002

[XXI] Savarese, C. Rabaey, J. Langendoen, K. 'Robust Positioning Algorithms for Distributed Ad-hoc Wireless Sensor Networks' USENIX Technical annual Conference, pgs 317-28, Monteray, CA, June 2001.

[XXII]D. P. Robinson and I. W. Marshall "An Iterative Approach to Locating Simple Devices in an Ad-Hoc Network" The London Communications Symposium 2002.

[XXIII] "Quorum Sensing in Bacteria", Melissa B. Miller and Bonnie L. Bassler, Annual Review in Microbiology, pp.165-99, 2001.

[XXIV] K. S. J. Pister, J. M. Kahn and B. E. Boser, Smart dust: Wireless networks of millimeter-scale sensor nodes, 1999, Highlight Article in 1999 Electronics Research Laboratory Research Summary.

[XXV] J. Hill, R. Szewczyk, A. Woo, S. Hollar, D. Cullar and K. Pister, System Architecture Directions for Networked Sensors, ASPLOS 2000, Cambridge, November 2000. http://webs.cs.berkeley.edu/tos/papers/tos.pdf

[XXVI] Manatee project web page. http://www.distlab.dk/manatee/

[XXVII] The Centre for Embedded Networked Sensing web page. http://cens.ucla.edu/