Self-organising Sensor networks

Ian W Marshall[1], Christopher Roadknight[2], Ibiso Wokoma[1] and Lionel Sacks[1]

1. EE dept. University college London, 2. BTexact Technologies Adastral Park, mailto: imarshal@ee.ucl.ac.uk

Sensor Networks [6, 5] consist of a large number of low-cost low-power devices, each with sufficient hardware to monitor one or more variables and send and receive the readings for these variables to other devices. 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 typically make use of ad-hoc networking, but normally lack the processing power to utilize the full richness of many proposed ad-hoc network protocols.

In the DTI funded project SECOAS [14] we are investigating how sensor networks can be used to collect rich datasets in an offshore environment, thereby enabling more accurate models of poorly understood interactions between coastal sedimentation processes. This will involve a new way of thinking for coastal oceanographers, marine scientists, managers and engineers. It is vital therefore not to further overload these users with the need to learn how to configure complex networks. We are therefore aiming to build a self-organising, collegiate system. 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. In addition each sensing scenario will have it’s own very specific characteristics depending on mobility of devices, time span of the sensing task, inhospitability of the environment, etc. For example fish move very quickly, glaciers very slowly so the optimal algorithms and application of sensor networks for each task will be different. Every sensing task must therefore be approached in a requirements centered way. Our hope is that our autonomy based approach will enable this specialization to be evolved in-situ with minimal manual intervention.

We have previously proposed and simulated evolutionary algorithms for use in an active network [10, 9]. We are now applying these algorithms to sensor networks. In our current model there are up to 100 sensing devices. Each sensing device can make 3 different measurements (pressure, temperature and turbidity), forward measurements to neighbours, delete measurements from its stack, combine (average) measurements, or idle. Each device is expected to minimize its battery usage (extend life) whilst collecting as many measurements as possible and ensuring they get relayed to one of 3 base stations (sink). To enable efficient routing to the sinks, nodes carry out an assessment (using a gossip based mechanism) of their nearest neighbours and discover a hierarchical level for themselves based on the number of hops to the sink. To minimize collisions neighbours at a similar level form groups (a quorum algorithm), which only send data once. Individual nodes also evolve internal rules for probability of executing each behaviour, and successful nodes share these with their neighbours using “plasmids” i.e. headers attached to data packets. Users control activity using high level policies – for example in the case we present here the user has declared that turbidity measurements have high priority and pressure measurements have low priority. This allows the network to optimize its use of limited resources as shown in Fig 1

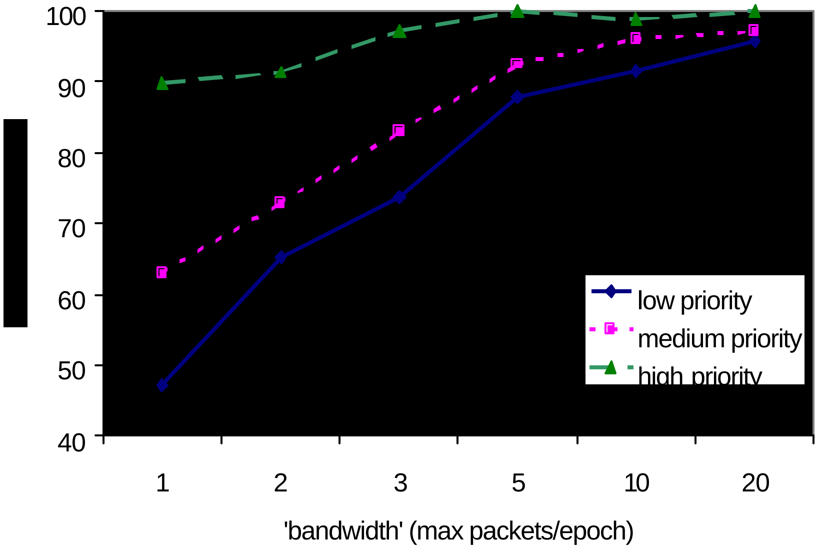

Figure 1. Number of packets sent and percentage dropped as ‘bandwidth’ increases

Hopefully it is clear from the figure that when bandwidth is constrained the network automatically drops a higher percentage of the low priority traffic, and delivers most of the high priority traffic. Taking this research further will involve:
Fine tuning the learning algorithm for different scenarios.
Carrying out further investigations of how the different learning approaches (ie. local rule based, genetic) interact when they are carried out in parallel.
Making routing decisions adaptive and subject to the same evolutionary pressures as other decisions.
Adding 'listening' to the list of possible states a node can be in. This means nodes can only accept data when they are in the listening state.
Automating the reward and penalty functions, so they too and configured in a hands off way.
Implementing the decision making solutions onto real sensor network devices

This will all be carried out as part of the SECOAS project [14]

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/secos/