Engineering for real – The SECOAS project

I.W.Marshall (University of Kent), **A.E.Gonzalez** (University College London), **I.D.Henning** (University of Essex), **N.Boyd** (Salamander Ltd), **C.M.Roadknight** (BT plc), **J.Tateson** (BT plc), **L.Sacks** (University College London)

Introduction

SECOAS was funded by the DTI as part of the Envisense centre (pervasive computing for natural environments) within the Next Wave Technologies and Markets initiative. The objective was to deploy a sensor network to monitor sedimentation processes at small scales in the area of Scroby Sands just off the coast at Great Yarmouth, Norfolk. Scroby Sands is the site of a wind-farm, and the DTI business case was based on improving the monitoring and impact assessment of offshore infrastructure. The research motivation from the academic side was to demonstrate the potential of a range of collegiate AI ideas that the group had previously simulated [1,2,3].

Successful deployment of a sensor network in a natural environment requires the devices to survive for long periods, without intervention, despite the fact that the conditions encountered are very likely to be antithetical to electronic devices. This means that any deployed devices must be both robust and able to deal with partial failures gracefully, without requiring large amounts of power. The project team believe that embedded AI is a good solution to enabling adaptation to failure and power management, since many AI algorithms are known to be tolerant to partial inputs and noise, and our own simulations had shown this to be true of the particular algorithms we were aiming to test. In other words the approach was likely to lead to reasonably robust software. On the other hand it is not possible to prove properties of this type of algorithm, and simulations can never capture the full complexity of reality (and are thus only indicative). It is therefore necessary to test ”in situ”, by undertaking a real deployment. SECOAS was designed to combine the software expertise of Kent, UCL and BT with the hardware expertise of Plextek, Essex and Salamander, and provide robust hardware and a challenging scenario that would represent a good test of the AI based approach.

What Happened

During the life of the project there were 3 full trials of sensing nodes, an initial deployment of one node (measuring Pressure, Temperature, Turbidity and Conductivity) for one week, an early deployment of 5 nodes for 2.5 weeks, and a final deployment of 10 nodes for 2 months. In all cases the AI algorithms performed well, and further tests are certainly justified. However none of the tests allowed an exhaustive characterisation of the software performance since the rate of failure of the nodes was significantly higher than expected. During the initial deployment (intended as a technology trial) no problems were observed. During the 5 node deployment one node failed completely and one node failed after 24 hours. Both failures were due to water ingress at an unplanned cable joint, introduced during deployment for operational reasons. No other hardware failures were observed. These trials gave the team confidence that the hardware was reasonably robust and had a good chance of surviving the planned final deployment of 2 months, providing the equipment was deployed as intended. In the final trial however 4 nodes were destroyed through external intervention, 2 nodes failed as a result of water ingress down a weak antenna cable (that was not intended to be submerged), and the remaining 4 nodes lasted for only 4 weeks. As a result the statistics generated by the recovered data samples are not sufficient to make conclusive claims about the software performance (remember the measured parameters exhibit long range dependency as a result of the turbulent flows, so very large sample sizes are needed). However the project did return some interesting oceanographic data, and some useful lessons for pervasive system engineering. These lessons are briefly outlined in the next section.

Lessons

A major difficulty faced by the project was aligning the language and methodologies of the software team and the hardware team. There does not appear to be any established literature on hardware /software co-design of extended systems of this nature, and the team had to create its own ad-hoc solutions (we spent a lot of time engaged in cross-disciplinary training). The difficulty was most clearly expressed close to deadlines when the software and hardware engineers had no mechanisms for dealing with the instability of each others outputs.

A second related difficulty was enabling full understanding of the limitations and failure mechanisms of the hardware and software across the whole team. This was most clearly expressed at module interfaces, where the software engineers tended to assume module clocks were as accurate as required, and the hardware engineers assumed the clocks would operate in spec. It turned out during the first multimode trial that neither group was correct, and more sophisticated interface protocols were used in the final trial.

A third key problem was the need to consider unintended interactions with non-project participants (such as local fishermen). The team did not start with sufficient expertise in this area.

Clearly for the future it is necessary to develop a methodology that systematizes approaches and solutions to these and similar problems.

We acknowledge the Support of the DTI and our industry sponsors (BT, Salamander and Plextek) and the contributions of many other team members during the life of the project. We particularly note the assistance of UEA with the oceanographic aspects of the work.

[1] A Novel Mechanism for Routing in Highly Mobile Ad-hoc Sensor Networks. J Tateson and I W Marshall. In Willig Karl and Wolisz, editors, *Wireless Sensor Networks*, number 2920 in LNCS, 2004.
[2] A Weekly Coupled Adaptive Gossip Protocol for Application Level Active Networks. I Wokoma, I Liabotis, O Prnjat, L Sacks, and I Marshall. In *3rd International Workshop on Policies for Distributed Systems and Networks (POLICY 02),* Monterey, California, 2002.
[3] Emergent organization in colonies of simple automata. IW Marshall and CM Roadknight. In J Kelemen and P Sosik, editors, *Advances in Artificial Life*, number 2159 in Lecture Notes in Artificial Intelligence, pages 349-356. Springer Verlag, 2001.