

# Deploying a Sensor Network in an Extreme Environment

K.Martinez, P.Padhy, A.Elsaify, G.Zou, A.Riddoch, J.K. Hart\*  
School of Electronics and Computer Science  
School of Geography\*  
University of Southampton  
United Kingdom  
+44 (0)2380 594491  
{km, pp04r, ae, gz, ajr}@ecs.soton.ac.uk  
jhart@soton.ac.uk\*

H.L.R. Ong  
Department of Engineering  
University of Leicester  
United Kingdom  
+44 (0) 116 252 5683  
h1r01@leicester.ac.uk

## Abstract

*A wireless sensor network has been designed and deployed to gather data from nodes deployed inside glaciers. This paper describes the solutions to power management, radio communications and other challenges faced in the system together with a discussion of the performance of the final system. 18 months of data have now been received, which provide an insight not only into the glacier's behaviour but also into the design decisions. The system uses custom PIC-based sensor nodes and an ARM-based base station which controls weather and differential GPS. Different versions have been installed in Norway from 2003-5 and this paper describes the lessons learnt from coping with the extreme conditions that of glaciers.*

## 1. Introduction

Sensor networks which are designed to be deployed in the natural environment present a unique set of challenges for system designers. The growth of sensor network research has enabled an increasing number of real deployments to be made. These deployments provide an ultimate test for the growing body of theory and algorithms. This paper highlights the design issues and solutions to problems unearthed from practical experience of installing three generations of systems in a glacial environment.

The numbers, types and complexity of successful environmental sensor networks are rapidly growing [9]. Berkeley's habitat modelling at Great Duck Island, which is analysing bird nesting habits, uses motes in the birds' burrows and is probably the largest deployment of its kind. It is one of the eight different projects undertaken by the Centre for Embedded Network Sensing at UCLA [14]. NASA's JPL laboratory has five deployments [4]. Recent developments have enabled mobile nodes to cover larger areas, using aer-

ial tramways [12], robotic vehicles (including small sub-kilogram intelligent tele-robots) [1], aeroplanes [10] and micro-submarines [13] for example.

The SEACOS project uses sea-bed sensors connected to buoys with long range wireless communication. It records wave properties around an off shore wind farm and uses intelligent sensors capable of dynamic self-configuration [2]. There are a growing number of sensor networks being developed in the oceans for hazard warnings around the world [7].

The FloodNet [5] project links a sensor network to a model and control system running on a GRID. This is an increasing trend as GRID technology can provide not only powerful distributed processing but many services, such as storage, for sensor networks [6]. This type of system has the potential to provide massive amounts of data which has many uses. The NWIS system in the USA provides water data from 1.5M fixed nodes for example [15].

The Volcn Tungurahua project [16] used motes to monitor volcanic activity using microphones to monitor sounds emanating from the volcanic vent during eruptions. Advances are still needed to make battery-powered event-monitoring systems such as this because most current nodes require too much power for continuous operation.

## 2. GlacsWeb and Objectives

The study of glaciers is an extremely important contribution to our understanding of climate change as well as geophysical processes [3]. Glacier behaviour depends on local conditions, and in particular the nature of the bed it rests on. Currently 90% of the discharge in Antarctica is from the ice streams, a behaviour that is controlled by sub-glacial processes and not necessarily by climate. Therefore, it is important to understand how glaciers contribute by releasing fresh water into the sea because this activity could cause

rising sea levels and great disturbances to the thermohaline circulation of sea water, which brings warmth to areas such as north-west Europe.

Previous projects that have attempted to measure properties of the sub-glacial bed used custom instruments including long wires to link to a logger on the glacier surface [3]. The loggers were intermittently read manually and hence there was no way of knowing if there were problems such as power failures. We also felt that the wires may have played a part in interfering with the measurements themselves.

GlacsWeb aims to measure the behaviour of the glacier bed using probes (sensor nodes) of size similar to that of palm-sized stones. By measuring their movements and other properties, a different view of what happens under glaciers is provided. The probes have to be wireless and need to last for many years so that they become completely embedded within the glacier and reflect the seasonal changes. Briksdalsbreen in southern Norway was chosen because of its high flow, bed characteristics and accessibility to region [8].

When we first approached this problem in 2001 there were many unanswered questions such as what radio frequency to use, how strong the probe cases would have to be and what kind of sensors to install. We decided to produce a completely custom solution so that the power management, software and hardware were all under our control. This led us to a greater understanding of the relationship between the designs and the real-world performance of the system, something which has proved important in other pervasive computing projects.

### 3. System Design

Our primary concern was to choose an appropriate frequency for radio communications to avoid any impediments caused due to liquid water present in bubbles and layers of the ice. Using the dielectric constant for ice found in the literature (3.17), measured at a frequency of 1.8 GHz, we computed that the losses over 100m would only be around -25 dB. This, however, was clearly over optimistic as this value was calculated for fresh water ice which did not contain the complex mixtures present in a real glacier. The power used was another factor as conventional transceivers typically have a legal limitation. Version 1 probes used an add link 868MHz, 10mW transceiver but this was found to have an extremely short range in glacier ice. Version 2 probes were installed with a Xemics 433MHz transceiver together with a power amplifier to boost signals to around 100mW and were tested to have a range of over 100m. Tests were usually carried out by lowering a transceiver, driven by a PIC, into a bore-hole and then walking around the glacier with a laptop and receiver testing where the signals reached. Although we recognised that a custom-made antenna design is probably needed we decided to use off-the-

shelf helical antennas. These are much larger than dielectric antennas for example but more predictable in the presence of PCBs and metal. Dielectric antennas were initially chosen because of their omni-directional nature. The weak performance on-axis of the helical antennas, however, was not found to be a problem in the later systems since the probes move around anyway, and there is usually a time when the angle is acceptable. A minor change to the antennas from compressed helical to slightly longer helical antennas improved signal strength so much that the probe cases were modified in 2005 to cope with the extra length. Another critical affect on range was the size of the ground plane for the antenna. Version 3 probes will use a 173MHz transceiver in order to reduce losses further, however this will require custom or larger antennas.

Having established the low level radio hardware a communications protocol had to be chosen next. As the hardware would inevitably change over time and not be related to a standard radio network a completely custom solution was chosen. In order to reduce complexity and improve our diagnostic capabilities a simple master-slave protocol was used rather than a complex ad-hoc network typically found in sensor networks. The probes remain passive and are interrogated by the base station using a fixed command set. By doing this we avoided the issue of collisions and MAC protocols. How to get the data out from the glacier area

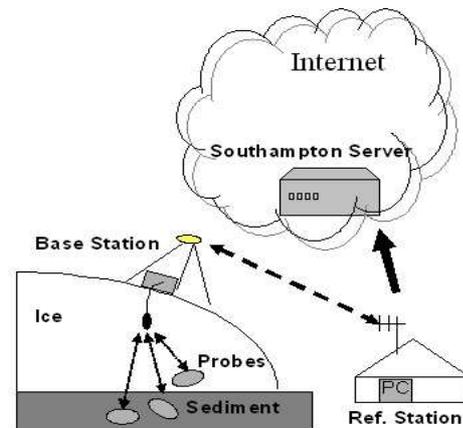


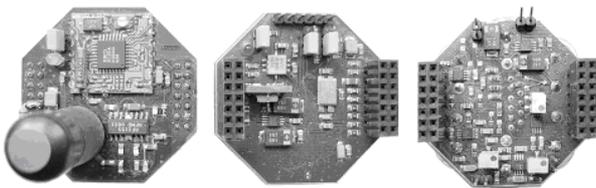
Figure 1. System Overview of Glacsweb

back to the UK was another problem which was eventually solved by using 500mW 466MHz radio modems over the 2.5km distance between the base station and the PC inside a cafe down in the valley. Although there was a GSM phone signal present on the glacier this would have proved costly and possibly unreliable in winter. Version 1 had a custom packet protocol for this long range link, which required more software and debugging. This was replaced by the point-to-point protocol (PPP) link reducing the need for software on the Linux PC in the cafe. PPP retries and error

checking cope with the small error rates normally encountered on this link. However it is still not 100% reliable. Files can be transferred daily to a database server in Southampton using the ISDN dial-up service available in the cafe.

This led to an overall system design illustrated in figure 1. The probes are lowered down holes drilled with hot water and the base station stands on top of the deployment area. Practical experience showed that radio losses in the first 10m of ice were far too high, especially in the summer, so the base station's radio link is lowered 20-30m into a hole on an RS-232 cable. Once we had a chance to test the radio in the cold winter months we found a surface radio could communicate with probes as well. This means a winter-only backup radio is possible.

The initial probe design contained one PCB with most of the functionality, in addition to a one off-board transceiver PCB. To make better use of the circular space inside the probe case a cylindrical modular design was made as shown in figure 3. While this required more connectors it allowed each module to be tested separately and problematic modules easily replaced. The transceiver module (figure 2, left



**Figure 2. Transceiver, Digital and Analogue boards**

PCB) can be used separately, with the serial port adapter and software running on a laptop. In the figure the Xemics module can be seen in the top of the transceiver, with the power amplifier below it. The component values and layout between the RF amplifier and antenna were revised a few times before coupling losses were negligible. This module - powered from the cells located on the analogue module (discussed below) - regulates and decouples the power supply for stable operation and to minimise noise that could otherwise affect the sensitive analogue module.

The digital module (centre) contains the PIC micro-controller (underside), a real-time clock (RTC), a digital temperature sensor and two dual-axis MEMS accelerometers to measure tilt. The PIC is in-circuit programmed via the headers, and communicates with other boards via the connectors (left and right of the circuit board). Like the transceiver, the digital module self-regulates and decouples the power.

The analogue module (right circuit board) contains two single-ended and two differential low-noise inputs. The gains of the differential inputs are adjustable via surface mount potentiometers. This arrangement allows fine tun-

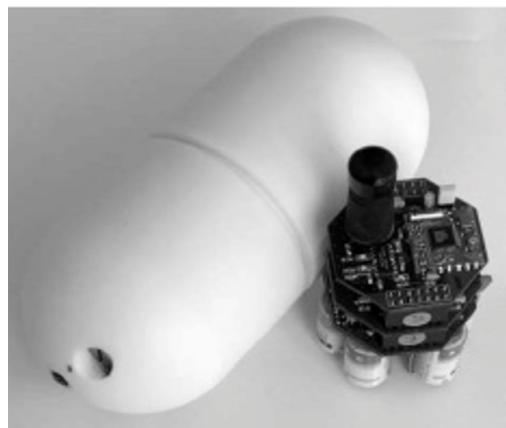
ing of the sensitive sensors typically connected to differential inputs. On the other hand, precision resistors determine the gain of the single-ended inputs as they are typically connected to less sensitive sensors. This reduces the component footprint. The analogue board regulates and decouples the power supply to minimize noise. It also generates a precision 2.5V reference voltage that is used by the analogue-to-digital converter of the PIC. The six lithium thionyl chloride cells and the connector for the sensors are located on the underside of the circuit board.

### 3.1. Physical Issues

The probe enclosures were made out of solid rods of plastic (polyester) cut in a CNC milling machine. Although these have been tested under high pressure by a test system in the School of Oceanography, Southampton University, it is difficult to know if their seals fail under the glacier. Potting compound is typically used in these situations to completely fill the container but we found the radio electronics was too sensitive to do this and therefore the probes were left with air inside.

The base station needed to control the probes, weather station, GPS, radio modem and other instruments while using little power. The first design used three (one master and two slaves) PIC processors but it could not be remotely configured so a low power ARM-based single board computer (BitsyX from Applied Data) was used. The base station now runs embedded Linux and is in standby mode most of the time, where it only consumes 120mW per hour. This is significantly higher than the PIC-based solution, but it is still comfortably within our battery capacity and considerably outweighs the PIC solution due to its re-programmability.

A custom version of cron (a unix-based scheduler) was



**Figure 3. A probe with its casing showing the pressure gauge and resistivity bolts at one end**

written which wakes up the system and runs jobs. This provided a way of running separate GPS jobs for example and a familiar interface for configuring jobs. A custom interface card was made that can power on/off all the peripherals and share serial ports among them. When powered on, the BitsyX and the interface card use an acceptable 1.45W of power. A parallel system of lead-acid batteries, solar panels and wind generator provides sufficient power for the base station. The base station enclosure is a traditional *Pelican*



**Figure 4. Base station in 2005 showing the main case open, wind generator and antennas**

case. The main problem on the ice is melting: over 10m of surface disappears annually making it impossible to screw anything down. We rely on the weight of the batteries and the low drag of the aluminium pyramid holding the antennas etc. The weight and sharp edges of the structure cut into the ice and minimise sliding. In 2005 the wind generator was added and this was deliberately kept low to keep the centre of gravity low. We knew we could do this because visits to the glacier in the winter show very little build-up of snow (which would otherwise hinder the generator). However the base station is always moving, preventing the use of directional antennas, good positioning of solar panels and necessitating the use of cable loops allowing the wire to the transceiver to be kept loose. In 2004 the base station fell into a large hole which opened up on the surface nearby, due

to the large amounts of surface water running into it. This probably short-circuited a cable and necessitated a maintenance visit.

### 3.2. Reprogramming Probes and Base station

The probe's firmware is divided into two segments: the program space (5k words) and the user space (3k words). The program space contains the initialisation and control routines that configure the micro-controller's in-built modules (e.g. I2C, ADC) upon starting. It also streams the bytes transmitted to and from the on-board UART (universal asynchronous receiver transmitter), interprets received packets and executes commands. The program space is unalterable and so it was the focus of extensive testing.

The user space, which holds the program that is autonomously executed whenever the probe is *woken up*, comprises the upper 3k words of the program memory. This program was written as a C function and calls the same set of commands that could have been sent via the transceiver. Multiple programs can be stored in the user space and the desired program is selected by issuing a specific command. Programs could also be loaded and removed from the user space via the radio link as the micro-controller has the ability to reprogram its program memory while running. This feature greatly increases the flexibility of the probes since their automated sequence could be remotely altered from Southampton even when they are buried in a Norwegian glacier. Finally, a mechanism was placed to ensure that rogue programs would be terminated if they exceeded some preset timeout, and they would not automatically execute the next time the probe *awakens*. As far as we know this has never happened. During communications tests for example probes were configured to simply send a stream of information.

To reprogram the probes or base station requires waiting for a temporary communications window at noon that involves connecting into the reference station PC via an ssh (secure shell protocol) tunnel and then opening a telnet session from there to the glacier. The base station automatically sleeps after its jobs, so a lock file is created on the compact flash memory card in order to prevent this. The shell scripts and binaries can then be changed easily. This has proved to be useful once trusted modifications are approved for installation. This meant testing on duplicate hardware in the lab first because unpredictable behaviour due to a bug may render the base station inoperable, which in turn requires a 48 hour maintenance trip costing \$1500. Key software or configuration changes are always made with two people present to double-check everything before committing. This combines with a thorough code-checking procedure by a second person (from the CVS repository) greatly reduces the risk of bugs.

## 4. System Performance Evaluation

The design process involved many decisions based on estimates of the risks involved in each part of the system. Now that 1273 probe-days of data have been received it is possible to use the data from sensors and logs to evaluate some of those design decisions.

The risks to the probes are summarised in figure 5. Part of the system design has been to minimise those risks but it can be seen that many risks are actually dealt with construction and management solutions. Implementing enough diagnostics inside the probes so that failures could be identified was not possible. Thus when probes are *lost* completely it is difficult to know exactly why.

The small changes made such as adding wind power effected a significant improvement each time. For example, the base stations suffered major power problems during the winter of 2004 when lack of sun prevented solar power generation. In December 2005, with the addition of the wind power the battery was still in good condition; well above 13V. The probes were designed to store up to 682 days of data in a ring-buffer to prevent data loss in the case of communication failures. This has worked well so far, as base station failures prevented live data recovery for up to three months. Similarly, the base station data archive held data when the long range link was down. Some risks and failures were unpredictable such as the breakdown of the base station in spring 2005, caused by the corruption of the compact flash memory card used in the base station as the main data backup. This malfunction prevented our main code from running on the base station. If this had been foreseen it would have been possible to take necessary actions to avoid such an event from blocking the main code. The base sta-

Risks	Prevention measures
Shorts or damage to sensors	Care in sealing probes
Water ingress	Test seal technique, care when sealing
Communications breakdown	Store and forward protocols
Shorts/Connector problems	Glue, sealant, testing
Faulty software	Testing, dual coding
Wired transceiver damage	Spare transceiver
Fails to go to sleep	Testing. Next day would sleep
Swept away by water	Daily data gathering, more TXR, Ad-hoc network

**Figure 5. Risks faced by probes and preventive measures**

tion's weather instruments suffered frequent failures due to mechanical problems and its need for a continuous power supply. This unreliability sometimes meant that missing records had to be obtained from a Norwegian weather service within the region. Ideally the server in the UK should automatically compensate for this missing data by fetching

the closest alternative data.

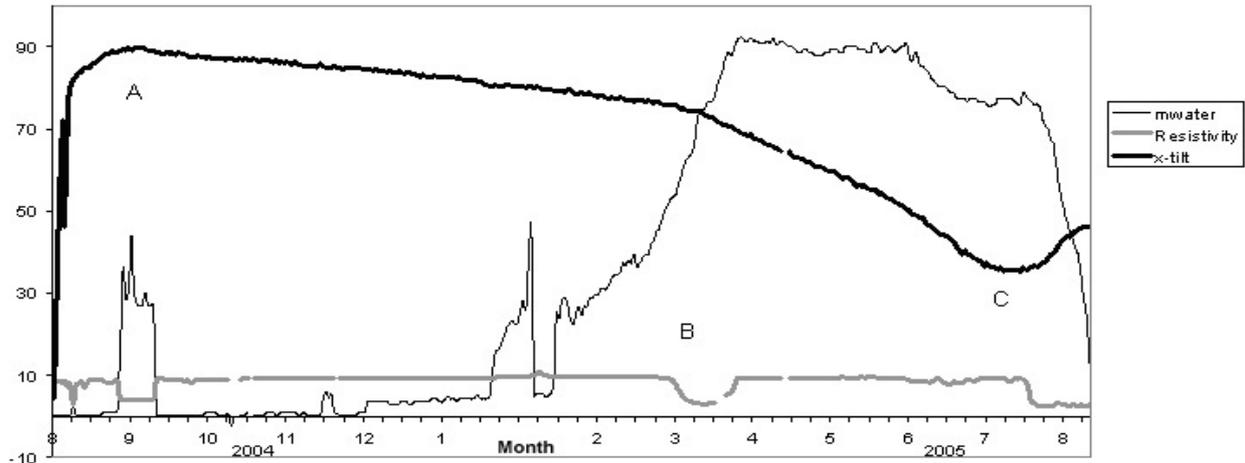
The base station's communications software can log events such as retries, behaviour changes and failures to a file which is sent back to Southampton every day. Analysing these shows several features of the communications system. The decision to limit the number of retries for a packet to only three proves to be generally correct. In the months after deployment, when the probes were still near the bore-hole and lots of water, the communications were less reliable. However there was rarely a need for more than three retries on a packet. Figure 7 shows the retry behaviour of the system, averaged across all nodes. It can be seen that there are more retries in the settling-in period after August than in the winter months when there is less liquid water around. The graph does not show anything from January and February which implies that the base station was down during this period.

## 5. Data Analysis

The project has successfully collected unique data vital for understanding sub-glacial processes. Probes 4 and 8 installed in the summer of 2004 provided particularly useful data as one was in the ice and the other underneath. The data has also proved useful in understanding the system's behaviour. For example rainy and cold periods often affect all the communication systems, including the reference station in the cafe that was cut off due to an avalanche once.

The data from the probe which survived more than a year from august 2004-5 is shown in figure 6. The data can be interpreted as follows. The settling-in period of around a month after deployment shows variations in water pressure as the hole closed-up, together with a drop in resistivity when water was present (event A). The readings are more stable by November and there is no relationship between resistance and pressure, indicating the holes have closed. From January 2005 the pressure steadily increased until it was higher than that possible from a column of water filling the depth of the hole: this is *over pressure* from the ice. The tilt shows the probe moves a lot in the first month then settles into one position. When spring arrives around March 2005 it moves more rapidly showing that the glacier is *waking up*. There is also a *wet* event (event B) around that time presumably from increased melting of ice and snow. Mid July (event C) sees a movement change, rapid pressure decrease and less electrical resistance consistent with the presence of lots of liquid water. This is when the glacier moves the most.

This data is also useful in evaluating whether the fixed sampling rate of six per day was sufficient. During *busy* events the data changes within a few hours and rapid transitions can be seen. This indicates that a higher rate would be advantageous then. However on average the rate was acceptable.



**Figure 6. Sample data from one year. Probe 8 water pressure expressed as depth of water (m), Resistivity (MOhm) and X tilt (degrees)**

## 6. The Future: An Agent-Based Approach

We set out to make the probes smart and autonomous so now that some of the fundamental elements have been tested, our current research advances this direction. The version we are designing for the 2006 deployment will implement a multiple hop, self-organising ad-hoc network that would improve data collection as well as reduce power consumption. Ideally, the probes would be completely autonomous and independent of any manual intervention to perform both reactive tasks (adjusting sensor sampling rates according to change in the sub-glacial environment) and proactive tasks (selecting the most efficient multi-hop route to transmit data to the base station). Presently, there has been an increasing interest in controlling sensor networks using multi-agent techniques and we have designed a protocol [20] tailored for GlacsWeb.

GlacsWeb can be modelled as a cooperative multi-agent system because all the sensor nodes are under our control. The next probes will work towards a predefined system goal of maximising data collection in addition to saving energy through a multi-hop approach in which probes would relay data for one another. Furthermore, at present, sensing is carried out at a pre-determined constant rate which is independent of the natural variations in the environment. This sometimes results in unnecessary sampling and more could be gained in a more dynamic situation given the same energy expenditure. We have developed a utility based sensing and communication protocol (USAC) [11] that introduces the concept of utility among sensor nodes and provides them with incentives to act independently in order to

maximise their goal.

This protocol treats each node as an agent and advances the current implementation of the network in the following ways:

- It initiates a mechanism for adaptive sampling in which each probe adjusts its sampling rate depending on the rate of change of its observations of the environment and a value derived function that they use to assign a value to the data they observe. This function is used in conjunction with a regression model that predicts a confidence interval for a future observation based on its past observations. A greater deviation from this interval results in a greater value being assigned to the observation. The average of all sensed values is then utilised in the communication protocol.
- It instigates a new routing algorithm that finds the cheapest cost route from the probe to the base station. The cost of a link from one probe to another is derived using the opportunity cost of the energy spent relaying the data: i.e. the value that a relay could have gained by using the energy in sensing instead of relaying.

The simulation results of over 200 instances of sensor networks with 20 sensor nodes can be seen in figure 8. The results show that USAC is a significant improvement over the existing GlacsWeb implementation. Whilst USAC not only consumes less energy over time but also collects more valuable observed data than the current GlacsWeb protocol.

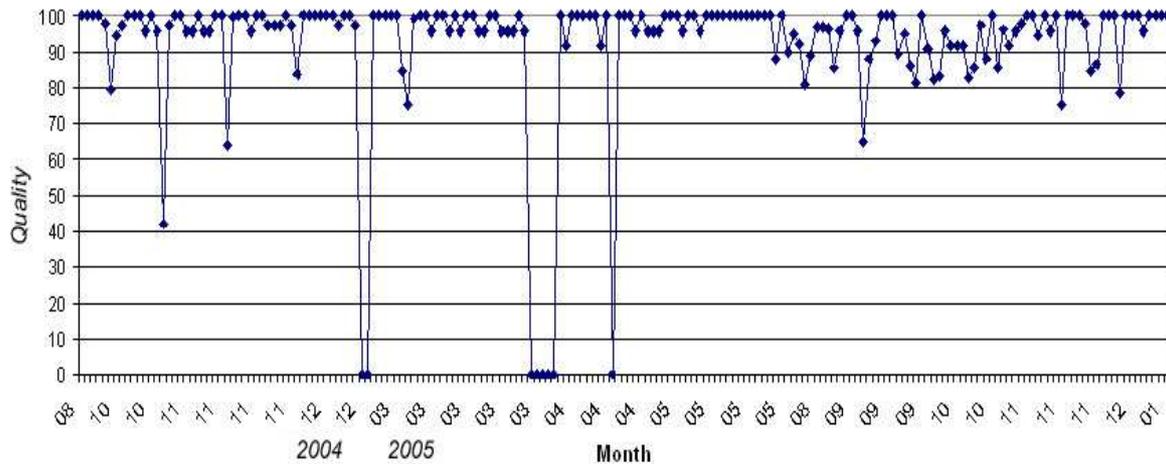


Figure 7. Percentage of good probe packets over 16 months (10000 packets)

## 7. Conclusion

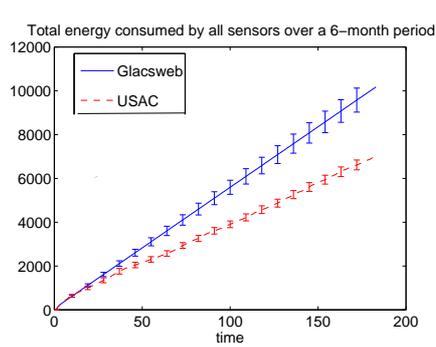
We have designed a working wireless sensor network for an extreme environment, overcoming technical and logistical problems to successfully investigate subglacial processes. However, even with careful design and construction we have learnt there is a difference between the predicted and actual behaviour of sensor networks when they are actually deployed. A combination of unknown factors and unpredicted interactions tend to affect the system in a way which can only be studied later. We have gradually refined our systems to be more fault tolerant and "smarter" and believe that the deployments have proved to be essential to a better understanding of how to make real sensor networks.

## Acknowledgment

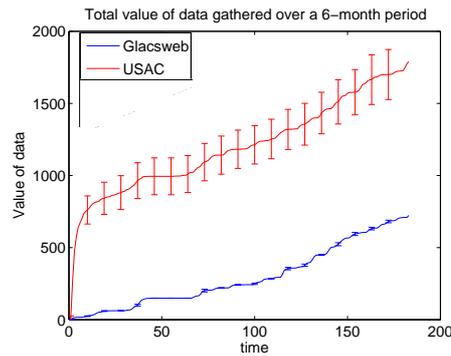
The authors thank Topcon for their GPS support; Hannah Brown, Kim Dowsett, John Hunt, Natalie Jarman, Dan Miles, Dave Robinson, Kathryn Rose, Sarah Stafford, Sarita Ward and Sue Way for field assistance; Harvey Rutt and Joe Stefanov for technical advice; Inge and Gro Melkevol for their assistance and hosting the reference station in Norway; Ken Frampton and Mark Long for making the probe cases; This project was funded by the Royal Society Paul Instrument fund, Department of Trade and Industry Nextwave programme and the Engineering and Physical Sciences Research Council.

## References

- [1] B. Argrow, D. Lawrence and E. Rasmussen, *UAV Systems for Sensor Dispersal, Telemetry, and Visualization in Hazardous Environments*. Proceedings of the 43rd AIAA Aerospace Sciences Meeting and Exhibit, 10-13 Jan. 2005, Reno, NV.
- [2] M. Britton and L. Sacks, *The SECOAS Project: Development of a Self-Organising, Wireless Sensor Network for Environmental Monitoring*. Second International Workshop on Sensor and Actor Network Protocols and Applications, 2004, Boston.
- [3] G.K.C. Clarke, "Subglacial processes," *Annual Review of Earth and Planetary Sciences*, 33, 2005, 247-276.
- [4] K.A. Delin and S.P. Jackson, *The Sensor Web: A New Instrument Concept*. SPIE Symposium on Integrated Optics, San Jose, CA. Jan 2003. <http://sensorwebs.jpl.nasa.gov/resources/sensorweb-concept.pdf>.
- [5] D. De Roure, *Floodnet: A new flood warning system*. *Royal Academy of Engineering Quarterly*, 23, 2005, 48-51.
- [6] R. Aldouri, G. R. Keller, A. Gates, J. Rosillo, L. Salayandia, V. Kreinovich, J. M. Seeley, P. Taylor, and S. Holloway, *GEON: Geophysical Data Add the 3rd Dimension in Geospatial Studies*. ESRI International User Conference Proceedings, San Diego, CA, 2004.
- [7] F.I. Gonzalez, H.M. Milburn, E.N. Bernard and J.C. Newman, *Deep-ocean assessment and Reporting of*



(a) Energy Consumption (Joules)



(b) Value Derived

**Figure 8. The total energy spent and total data value gathered at the end of a 6-month period plotted against number of agents in the network**

*Tsunamis (DART): Brief Overview and Status Report.* Proceedings of the International Workshop on Tsunami Disaster Mitigation, 1998, Tokyo, Japan.

- [8] J.K. Hart, *An investigation of subglacial processes at the microscale from Briksdalsbreen, Norway.* Sedimentology, in press.
- [9] J.K. Hart and K.Martinez, *Environmental Sensor Networks: A revolution in the Earth System Science?* Earth Science Reviews (submitted 2005).
- [10] T. Melodia, D. Pompili, V. C. Gungor, and I. F. Akyildiz, *A Distributed Coordination Framework for Wireless Sensor and Actor Networks.* Proceedings of ACM Mobihoc 2005, Urbana-Champaign, IL, USA.
- [11] P. Padhy, R.K. Dash, K. Martinez and N.R. Jennings, *A utility-based sensing and communication model for a glacial sensor network* Proceedings of 5th International Conference on Autonomous Agents and Multi-Agent Systems, Hakodate, Japan, 2006.
- [12] M. Rahimi, R. Pon, W.J. Kaiser, G.S. Sukhatme, D. Estrin, and M. Srivastava, *Adaptive Sampling for Environmental Robotics.* Proceedings of the 2004 IEEE International Conference on Robotics and Automation, 26 April-1 May 2004, pp. 3537-3544.
- [13] J.P. Ryan, F.P. Chavez, J.G. Bellingham, E. Rienecker, R. Kudela, A. van der Woude, R. Maffione and A. Fischer, *Environmental processes in the Monterey Bay National Marine Sanctuary: Studies integrating AVIRIS and synoptic in situ sensing.* Proceedings of the Eleventh JPL Airborne Earth Science Workshop, 2002, 307-318.
- [14] R. Szewczyk, J. Polastre, A. Mainwaring and D. Culler, *Lessons from a Sensor Network Expedition.* 1st European Workshop on Wireless Sensor Networks (EWSN '04). Berlin, Germany, January 19-21, 2004.
- [15] *United States Geological Survey.* Accessed at <http://water.usgs.gov>, 16th Jan 2006.
- [16] G. Werner-Allen, J. Johnson, M. Ruiz, J. Lees, and M. Welsh. *Monitoring Volcanic Eruptions with a Wireless Sensor Network.* Proceedings of European Workshop on Sensor Networks (EWSN'05), January 2005.