

Field deployment of low power high performance nodes

Kirk Martinez, Philip Basford, Joshua Ellul, Richard Clarke
School of Electronics and Computer Science
University of Southampton
Southampton, UK
{km, pjb08r, je07r, rsc106}@ecs.soton.ac.uk

Abstract— When deploying a sensor network into a harsh environment the need for high levels of fault tolerance and maximising the usage of available resources become extremely important. This has been achieved by implementing a highly fault tolerant system based on our Gumsense boards. These combine an ARM-based Linux system with an MSP430 for sensing and power-control. It also allows for dynamic schedule modifications based on the available power and can be synchronised with other systems without relying on direct communication, autonomous behaviour in case of total communications failure is also supported. A deployment on Vatnajökull, the largest ice-cap in Europe, has provided a long-term test for the systems and revealed strengths and weaknesses in the design decisions.

Sensor Networks, Glaciers, Gumstix, Python

I. INTRODUCTION

Wireless sensor networks typically rely on gateway nodes in order to link with external networks. In many cases these gateways are the single points of failure which can make a sensor network difficult to manage [1, 2]. With the emergence of powerful systems capable of sensing as well as gateway tasks it is vital to develop methodologies and techniques to manage them.

Glacsweb is a long term project to monitor the response of glaciers to climate change [3] and has carried out deployments in Norway and Iceland since 2003. This monitoring is achieved by using a hot water drill to place nodes (“probes”) beneath the surface of the ice, in the case of the Iceland deployment the probes are approximately 70 metres beneath the surface. These probes are equipped with an array of sensors chosen to measure changes in conductivity, orientation and pressure. The gateway node on the ice surface also has sensors. In addition to temperature and ultrasonic snow level sensors a differential GPS (dGPS) system is used to record ice velocity changes on both a diurnal and annual scale. This is in order to understand the nature of glacier movement, in particular the relationship of any ‘stick-slip’ motion to changes in water pressure [4, 5]

The data gathered from the probes and dGPS is buffered locally until the scheduled communications window. To reduce power consumption whilst providing timely data this window

is daily, at midday UTC. During this communications window the base station queries the probes for data and retrieves it. If for any reason the communications fail the data is stored locally until it can be sent onwards. From previous deployments this is known to occur frequently, especially in the wetter summer environment. One major challenge faced whilst working on the Glacsweb project is that the system needs to be self maintaining to avoid winter field visits.

In the Glacsweb project [6] a second base station is used to take fixed-location dGPS recordings and meteorological measurements. The hostile nature of the environment means that communications and power generation are highly unreliable. This means the base stations need to be aware of their power and communications states. During the design cycles of the project a considerable amount of work has been carried out on the base station design to make it as reliable as possible while using minimum power. They have to be capable of surviving a long winter (Dec-March) by minimising their tasks. In 2008 a new Gumstix (connex) based base station [7] was implemented and deployed which combined low power sensing and sleep modes (MSP430 driven) with a high performance ARM-based Linux system. This paper outlines the specific design issues and experiences from the long term deployment which are useful for other base station implementers. Fig 1 shows the glacier base station with its aluminium structure to support antennas etc.

II. ARCHITECTURE

The Glacsweb project has been ongoing since 2003 and during that time has gone through a number of different system architecture and hardware configurations. A major contributor to architecture design decisions is the way differential GPS (dGPS) works. In order to dramatically improve the accuracy of the position fix of a mobile object a simultaneous dGPS recording for a known location is needed. This location is known as the reference station, and consists of a dGPS unit connected to some sort of computer control hardware and an external data connection. The constraints surrounding the reference station have differed significantly between the two main deployment areas that Glacsweb has worked in leading to very different designs which have knock on effects for the rest of the architecture.

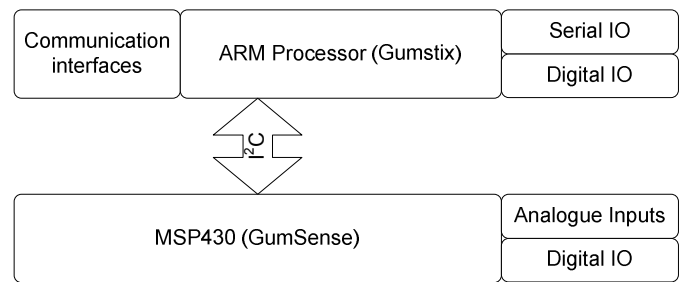


Figure 1 Jeff and Phil carry out repairs to the base station

The previous system deployed in Norway gained a lot of flexibility from running a Linux-based operating system and was more easily managed remotely. It used power management techniques such as only powering peripherals when needed, however, its sleep current was relatively high, which meant it needed a large power reserve in the winter months. The area in which the network was deployed in Norway had very little annual snowfall meaning the wind generator could supply power in winter, whereas in Iceland the expected snow would even stop that source from being useful. The system design in Norway used a point-to-point protocol (ppp) IP link over a radio modem.

The original plan was to base the Iceland architecture on the system used in Norway just replacing the ADSL line with a GPRS modem to provide the internet link via the fixed reference station system on the café 1km away. This option was ruled out for a number of reasons.

The primary reason being power constraints; in Norway the café the reference station is situated at has power available all year. Whilst the Iceland reference station is also attached to a café the power there is only available during the tourist season (April to September) the rest of the time the system needs to be entirely self contained. This increased the complexity of the system significantly, and forced a change in the hardware to be used, as a 'normal' PC would no longer be suitable. It was decided to use identical hardware to the base station – a Gumstix [7] processor attached to a Gumsense [8] board to enable software controlled powering of peripherals and the Gumstix to be woken according to a schedule. The details of



the division of i/o and the communications between the two

Figure 2 Gumsense architecture showing the communication between processors and the distribution of IO

processors are shown in Fig. 2. The system also has a 4GB compact flash card for data storage. The use of a Gumsense board also means that the reference station can monitor its battery voltage and internal temperature and humidity, providing additional data streams from the glacier.

The Gumstix offers an attractive package for sensor network deployments because it provides a lot of processing power in a small footprint (400-600MHz, 80x20mm). However, this processing power comes at the cost of high power consumption (~100mA) and no useful sleep mode. It is for this reason that in these deployments it is combined with an MSP430, meaning the Gumstix is only powered when there is a need for more processing power. This dual processor platform using Gumstix has also been implemented in [9]. The Stargate Netbridge [10] is another gateway platform that is available, however the Gumsense platform is superior for this application. The use of Gumsense is preferable because it has more processing power, RAM and flash as well as including better power regulation.

Whilst in Norway the ppp daemon at the café end of the long range radio modem link could remain running all the time without worrying about power depletion, allowing simpler control logic. With the reference station being battery powered and having to turn off the radio the ability to differentiate between reasons for disconnects becomes vital. This is because the reference stations response to the various reasons is different. If the ppp session disconnected due to interference/temporary failure then the reference station should remain powered for a short period to enable a reconnection attempt. If the termination of ppp was caused by the system having successfully finished its transfer then the radio can immediately be turned off to conserve power.

When testing the long range modems (500mW 466MHz) and the ppp link in the research lab it was found to be very unreliable with frequent drop outs and a very low data rate. It was also observed that the reliability was affected by the time of day which implies that the problems were caused by local interference. Initial testing on the glacier suggested that the modems would be more reliable there than in the lab, however other factors led to the eventual removal of the long range modems from consideration.

Another challenge posed by the fact that the reference station is not powered all the time is the fact that it has to be able to reliably open the ppp communication window at the same time as the base station every day. This synchronization

could be achieved by using the time received from the GPS receiver to synchronize the clocks, however even if the stations started at the same time there is still the potential for the code executed prior to establishing the ppp link to take different times leading to a lack of synchronisation.

In order to give a suitable range for the long range communications it was proposed to use a directional antenna on reference station end. In order to get as close as possible to line of sight to the glacier the antenna would have to be placed on the most exposed side of the café. The area the network is deployed in has very heavy snow fall and high winds which caused damage to the metal frame of the base station pyramid and also to antennas that had previously been mounted on the café. This meant that it was thought unlikely that a directional antenna would survive through the winter on the café. This would mean that communication with the base station would be extremely unlikely after winter.

It was decided to abolish the inter-basestation link and to give each station its own GPRS modem. This approach has many advantages and was the solution that was finally implemented. One advantage of the separation of the systems in this way is that they become independent. This independence means that the failure of one will not adversely affect the other whereas using the previous scheme if the reference station failed in any way then all communication with the base station would also cease. It also means that the tight time synchronisation between communication windows is no longer a requirement for data and command transfers. Synchronisation between dGPS readings is still required and maintaining good time accuracy on the two units is still needed. The dGPS is activated by the microcontroller meaning the execution of software on the Gumstix does not cause drift in the timings of the dGPS. Controlling the dGPS from the microcontroller instead of the Linux system is a change from previous deployments and has been achieved by setting the dGPS to automatically start taking a reading whenever it is turned on.

Before deciding on the dual GPRS modem approach the costs of switching to this method were carefully analysed. The data sent over the GPRS link is paid for per megabyte and so any changes in the amount of data sent would have a cost implication. However the architecture does not dramatically affect the amount of data sent back to Southampton so the cost implication is minimal. There is however a significant power saving over the entire system. This is because the GPRS modems require less current to operate as shown in Table 1. This means that a twofold power saving can be made, both because the hardware is more efficient

TABLE I. CHARACTERISTICS OF SYSTEM COMPONENTS

Device	Transfer Rate (bps)	Power Consumption (mW)
Gumstix	-	900
GPRS Modem	5000	2640
Radio Modem	2000	3960
GPS	-	3600

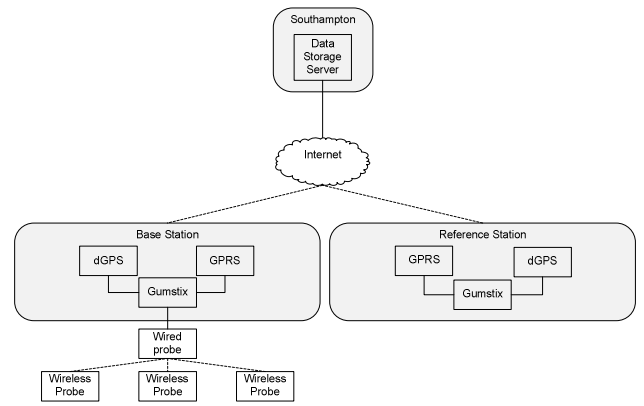


Figure 3 The final System Architecture

and the data from the base station does not have to be sent to the reference station before transmission. The final system architecture is shown in Fig 3.

III. POWER MANAGEMENT

It is known from past deployments that winter conditions reduce the amount of power incoming from solar and wind chargers. The systems must be capable of adapting to preserve their battery power. This is achieved by taking measurements of the battery voltage every thirty minutes and storing these in the microcontroller. Once a day these voltages are downloaded to the Gumstix and a daily average calculated. This averaging is to enable the overall health of the battery to be determined rather than just the health at midday (the time the Gumstix is awake) as the highest voltage for the day is reached at approximately midday as shown in Fig 5.

The power state is then adjusted according to Table 2. The most power intensive jobs that are performed are the GPS and GPRS. Taking sensor readings has negligible cost as it is managed by the MSP430. Retrieving the probe readings has a higher cost because it requires the Gumstix to be powered but it has been shown that radio communication with the probes is better in the winter due to the drier ice conditions so probe communications should always be attempted. Both the reference station and base station have external sources of power available, the base station has a solar panel (10W) and a wind turbine (50W) and the reference station has a solar panel and a mains charger input for when the café has power.

In order to capture ice motion it is desirable to maximise the dGPS reading frequency and duration,

TABLE II. POWER STATES

State	Minimum Threshold (V)	Probe jobs	Sensor Readings	GPS	GPRS
3	12.5	Yes	Yes	12 per day	Yes
2	12.0	Yes	Yes	1 per day	Yes
1	11.5	Yes	Yes	No	Yes
0	-	Yes	Yes	No	No

this has been achieved by allowing the frequency of readings to be varied. Altering the frequency of the readings has a major affect on the amount of power used per day. This is because each dGPS reading is approximately 165KB, although the exact size varies depending on the number of satellites available at the time of the reading. This means that by reducing the amount of dGPS readings taken not only is the amount of power used by the dGPS to take readings is reduced, but also the amount of time taken to transfer the readings from the dGPS's internal compact flash card to the Gumstix and the amount of data that has to be transferred back to Southampton. Although some researchers [12] leave their dGPS recording

full-time in order to obtain high precision, as seen in Table 1, the GPS device uses 3.6W of power use would deplete 36AH of batteries in 5 days, where as in state 3 as described in Table 2 the dGPS unit would deplete the reserves in 117days (for simplicity these figures do not include the consumption of any other component of the system).

The other problem with continuous sampling is the volume or data produced would be too great to transmit off-site in a power-efficient way.

The reference station is the fixed point used for these calculations; this means that the readings from one station are less useful than when readings for both stations are available. There needs to be a way of keeping the number of readings taken per day in sync between the reference station and base station. The new architecture does not allow direct communication between the two stations. In order to overcome this limitation the communications are managed by a server in Southampton, this also allows easy manual overriding of the power states if required. The execution sequence on each station is described in Fig 4. The reason for the upload and download of power states being in different places is to allow for minor variations in timing between the base station and the reference station. The upload of data is known to take a few minutes therefore as long as the time variation in the stations is less than the time it takes for the station which is ahead to upload its data then any changes will be reflected the same day. If the variation in time is greater than this then there will be a one day lag in the states being updated.

When a station requests the override state from the server the server looks up both the existing states from the stations and returns the lowest one to the client. This logic is further backed up by logic running on the stations themselves which does not allow the state to be set higher than the battery voltage allows, or for the station to be forced into power state 0. This is to prevent early depletion of the batteries, or the system being forced into a state in which it does not do communications.

A further safety measure is that if the fetching of the override state from the server fails for any reason then the system will just rely on its local state and take the appropriate readings. The transition between power states can be seen in Fig 5 when in power state 2 the fluctuations in voltage level are due to the wind and solar recharging the batteries and there is no regular pattern. When the system has switched to state 3 regular dips in the battery voltage can be seen, these dips have an interval of 2 hours. This shows that the dGPS uses a lot of power but that the power can easily be recovered when there is sufficient power generation. Although initially the voltage was high enough for the system to be in state 3 it was being held in state 2 by the remote override system.

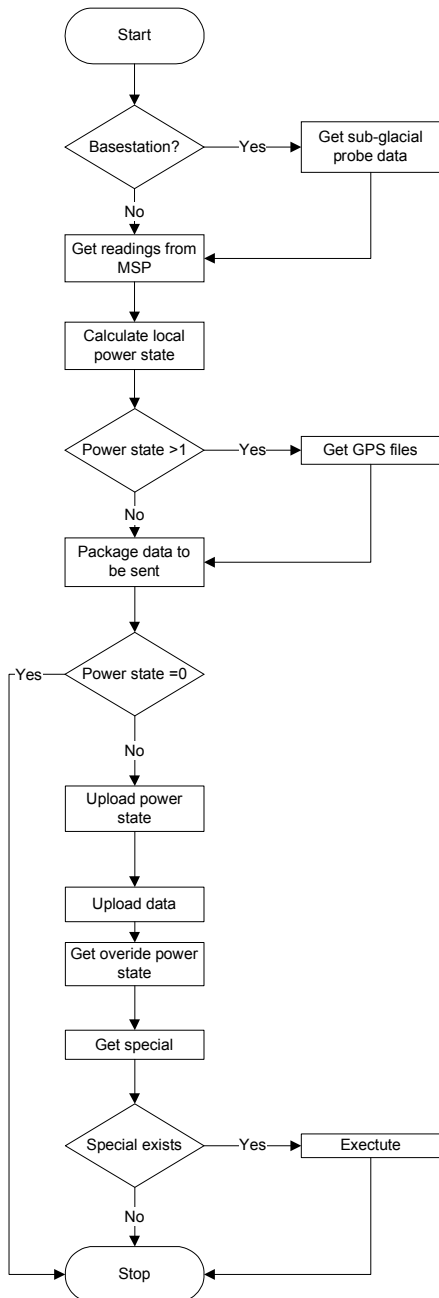


Figure 4 Flowchart showing system operation

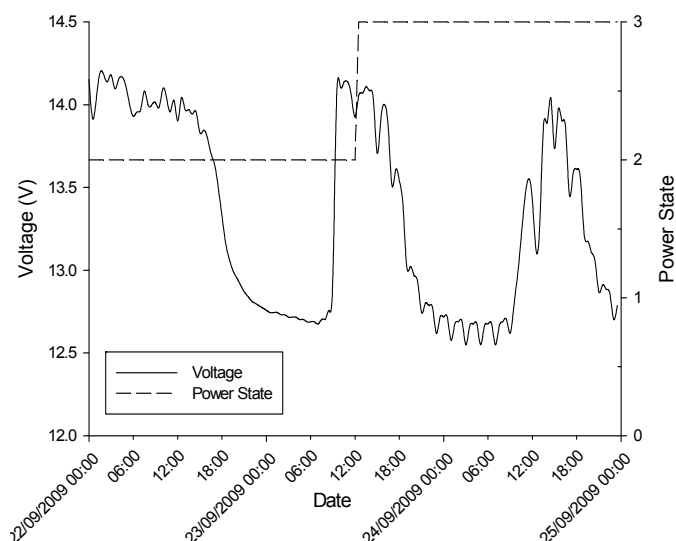


Figure 5 Sample data from Base Station showing Diurnal changes and ripples due to background dGPS task

This system allows for different schedules depending on power levels and could be extended to analyse the data and behave differently depending on the priority of the data as described in [8]. Whilst this power management technique aims to extend the lifetime of the system as much as possible it is however still possible that the system can be totally depleted, this situation has been allowed for.

IV. AUTOMATIC SCHEDULE RESETTING

The systems have external power inputs meaning that their batteries can recover from total exhaustion. However, this has to be detected because the schedule for the microprocessor is stored in RAM so will need to be re-written, a more fundamental issue is that the real time clock will have reset to 0 which is 01/01/1970 00:00. This can be detected because the systems stores the last time that it successfully ran. It then checks that its current time is before the last time the system ran; if that fails it knows that the real time clock (RTC) is not to be trusted. The RTC has to be corrected for synchronisation with the probes, other station and for any of the measured values to have meaning. This can be achieved by turning on the GPS and using that to reset the system clocks. If the system cannot set the time using GPS then the system will sleep for a day and try again. In the future this could also be extended to fall back to getting the time using the GPRS link and network time protocol (NTP). Once the date and time have been corrected the system will set the schedule to state 0 as described in Table 2 and will then proceed as normal.

V. SUBGLACIAL PROBE DATA

The probes deployed in the summer of 2008 survived longer than previous generations (4/7 after one year), with fewer vanishing offline and data is being produced by two after 18 months under the ice. Fig 6 shows a sample of data from three probes towards the end of winter. The electrical

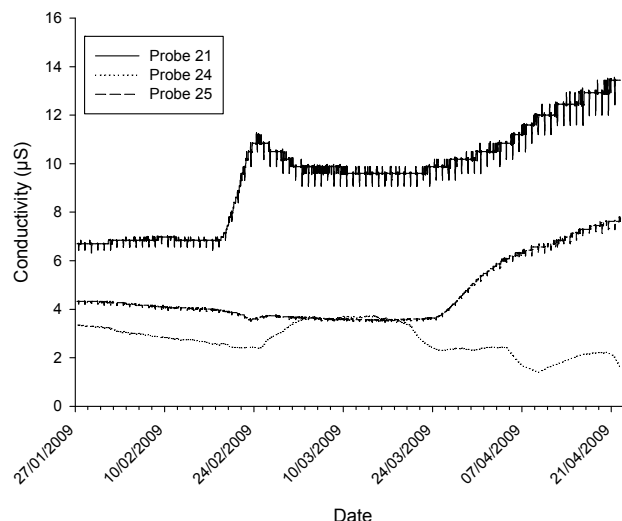


Figure 6 Sample data from three sub-glacial nodes showing electrical conductivity changes at the end of winter

conductivity increases show that melt-water is starting to reach the glacier bed.

However there were lessons to be learnt about base station design due to the large quantity of data they transmitted after months offline. This was due to the base station being damaged by deep snow and the failure of the wired probe. In order to avoid the wired probe being a single point of failure using several wired probes has been considered in the past, but was ruled out in this deployment because of the lack of serial ports.

Although the base station code receiving the data used a new technique, avoiding acknowledge packets, the real-world limitations were revealed by this large data fetch of thousands of readings. The algorithm records missing or broken data packets then later requests individual readings which were missed, unless there were so many that it would be as efficient to request them all again. With 3000 readings being sent in the summer, across the weakest link (due to summer water) 400 missed packets were common. Fetching that many individual readings was never considered in the testing phase and the process could fail. Fortunately the task was not marked as complete in the probes; so many missing readings were obtained in subsequent days.

Small adjustments could be made to the base station behaviour in order to try different strategies for retrieving data and the system is still running successfully in October. One of the many lessons learnt from this deployment is the importance of a reliable robust remote configuration system.

VI. LESSONS LEARNT

During the course of the field work and the weeks following it some valuable lessons about sensor networks have been learnt, these have come about because of differences in predicted behaviour and actual behaviour or because of eventualities that were not considered in advance.

As shown in Fig 4 the data is sent back to Southampton before the execution of the special command shell script sent from Southampton, while this is not a problem on its own when combined with the safety mechanism that is running on the system it causes a problem. This safety mechanism prevents the system from running for more than two hours at a time. This is to make sure that if something crashes in the system – for example a SCP transfer hangs – the system does not remain running until its batteries are depleted. The most common reason for this situation to occur is a communication drop out, which means the next day it is possible it will succeed. This does however cause a problem in some situations. These situations occur when either the data from the GPS has not been successfully downloaded for approximately 21 days whilst in state 3 or 259 days in state 2. As in this case there will be more data than can be downloaded from the GPS in 2 hours. Another more likely cause for this 2 hours maximum to cause problems is if the GPRS has not worked for a few days. In this case there will be more data than can be sent within the time frame. In both these situations the data will be processed file by file, and so over the course of a few days the backlog will be cleared. A more serious problem is caused when the GPS has not been downloaded for a few days because this could cause a single file to exceed the amount that can be transferred in a single window, meaning that no progress could ever be made. In order to solve this problem without removing the emergency time out it is suggested that the execution of remote code is performed before the data is transferred. Fortunately this risk is minimal as it could only be caused by an intermittent RS232 cable or dGPS unit, which has never been encountered.

Another problem that has been encountered is that the output from the special file downloaded from Southampton just goes into the normal log files. This means that there is a 24 hour delay between executing the code and getting the results back. This then leads to a 48 hours delay between the code being sent and the results from it being acted upon.

The inaccessibility of the stations in the system means that any changes made to the code have to be carefully verified. This verification process includes testing on similar hardware in the lab before the code or binaries are sent to the stations. In order to make sure that the code has arrived at the station without corruption the code then has to have a checksum calculated. This process has been automated by having scripts on the system which automatically download the program, calculate a checksum and if it is correct replace the old file with the new one. Additionally to overcome the delay in the results being accessible in Southampton the script that performs this verification uploads the MD5sum that it has calculated using a HTTP GET (the version of wget in use does not support POST) when the verification is performed, this enables researchers to know immediately if the transfer was successful.

Another issue that has been encountered is that the amount of output from the binaries whilst working on the systems is excessive for remote debugging, because of this care needs to be taken to make sure that the final binaries left on the system still have enough output but without causing excessive amounts of output. An example of this is when a probe is

communicated with for the first time in a few months then over 1 megabyte of log data can be produced, which then takes time/power/money to transfer but is of little use. To minimise the amount of redundant data transfer careful consideration of the log output needs to be made before the final deployment.

Although Linux-based base stations can be programmed in almost any language we have used Python for all high-level code, with pre-existing C binaries either being executed directly or libraries linked-in using Cython. This isolates the lower level tested and trusted code which communicates with modules such as the GPS or probes. All decision making, most time-outs and state-handling is written in Python so that it is easily modified in the field or with remote updates. All messages or errors are redirected to a standard logfile which is sent back daily with the data.

The deployment in Iceland has provided a valuable test bed for the algorithms discussed in this paper and in the post deployment phase of the project new lessons are regularly being learnt.

VII. CONCLUSIONS

The design of base stations and gateways requires as much consideration as the individual sensor nodes because they act as single points of failure. They also carry out important sensing and data buffering tasks. The systems described have improved longevity and sampling through the use of techniques normally reserved for smaller WSN nodes which take into account not only communication but also sensing costs. This field deployment has shown that the implemented power management design works well and provides not only high reliability but has improved sensing capabilities without compromising system lifetime. Further system safety enhancements including system clock and schedule recovery after power failures will improve fault-tolerance.

This work could be extended by enabling the base station to analyse the data collected and prioritise it forcing communication even if the available power is marginal if the data warrants it. This system also shows that data collated from the base station can provide useful insights into the condition of the system. This information could be further improved by adding more sensors to the base station; examples of possible additional sensors include; pitch and roll so that the enclosure's movement as the ice melts can be tracked. There are many applications where a set of nodes as powerful as these base stations would be useful so the lessons learnt are generally applicable in other GPS or meteorological sensing systems for example.

A further issue that was discovered whilst on the field trip is that the CF card used to store the readings from the previous year had become corrupted. The exact cause of the corruption is unknown and it proved possible to recover the data from the card, however it prompts investigation into whether a more suitable file system format can be found for the storage card.

These deployments have been highly successful not only because data has been continuously received but also due to the

resultant generalisable techniques which can be used elsewhere to improve system performance.

ACKNOWLEDGMENT

Our thanks to Royan Ong for probe PCB design and Jeff Gough for tireless hardware debugging. Also the 2007 Glacsweb team for their dedicated efforts: Kathryn Rose, Robert Spanton, Tom Bennellick, Stuart Rimmer and James Cheshire.

REFERENCES

- [1] Werner-Allen, G., Lorincz, K., Johnson, J., Lees, J. and Welsh, M.: Fidelity and Yield in a Volcano Monitoring Sensor Network. Operating Systems Design and Implementation. *Proceedings of the 7th symposium on Operating systems design and implementation*, vol. 381 – 396. (2006)
- [2] Martinez K, Hart JK, Ong R.: Environmental Sensor Networks. *Computing*, vol. 37 No 8: pp 50-56. (2004)
- [3] Hart, J. K., Rose, K. C., Martinez, K. and Ong, R.: Subglacial clast behaviour and its implication for till fabric development: new results derived from wireless subglacial probe experiments. *Quaternary Science Reviews*, vol. 28, pp. 597–607, (2009).
- [4] Bahr D. B. and Rundle, J. B.: Stick-slip statistical mechanics at the bed of a glacier. *Geophysical Research Letters*, vol. 23 (16), pp. 2073-2076. (1996).
- [5] Fischer, U. H. & Clarke, G. K. C.: Stick-slip sliding behaviour at the base of a glacier. *Ann. Glaciol.* Vol. 24, pp 390–396, (1997)
- [6] Martinez, K., Hart, J.K., Ong, R.: Deploying a Wireless Sensor Network in Iceland. *Lecture Notes in Computer Science, Proc. Geosensor Networks*, vol. 5659, pp. 131-137 (2009)
- [7] Gumstix, <http://gumstix.com>.
- [8] Martinez, K., Basford, P., Ellul, J., Spanton, R.: Gumsense - a high power low power sensor node. 6th European Conference on Wireless Sensor Networks, Cork; Ireland (2009)
- [9] Keller, M., Beutel, J., Thiele, L.: Mountainview – Precision Image Sensing on High-Alpine Locations. 6th European Conference on Wireless Sensor Networks, Cork; Ireland (2009)
- [10] Stargate NetBridge Datasheet. Crossbow Technology Inc. (2007)
- [11] BitsyX - PXA255, http://www.applieddata.net/products_bitsyX.asp.
- [12] Geirsson, H., Árnadóttir, T., Völkens, C., Jiang, W., Sturkell, R., Villemin, T., Einarsson, P., Sigmundsson, F., Stefánsson, R.: Current plate movements across the Mid-Atlantic Ridge determined from 5 years of continuous GPS measurements in Iceland. *Journal of Geophysical Research*, vol. 111, (2006)