

Monitoring in a High-Arctic Environment: Some Lessons from MANA

Marcus Chang
IT University of Copenhagen
mchang@itu.dk

Philippe Bonnet
IT University of Copenhagen
phbo@itu.dk

Monitoring terrestrial high-arctic ecosystems is important because of their great exposure to global warming. Ideally, these ecosystems are monitored continuously to capture the evolution of their characteristics year-round. This requires a pervasive monitoring infrastructure that collects data automatically. A consequence is that measurements that have traditionally been obtained manually should now be obtained with automatic measurement systems. Deploying such systems in a high-arctic environment raises specific challenges due to limited access, extreme weather and the absence of communication infrastructure. In the MANA project, we tackle these challenges in the context of a sensor network based data acquisition system for year-round lake monitoring in North-East Greenland. In this paper, we describe our system design and report the lessons we learnt from the first year of deployment. We emphasize the issues we underestimated initially: i.e., the consequences of operating in a remote region, the impact of extreme weather not only on system design but also on operator activities, and the demands that derive from the absence of communication infrastructure. We also reflect on the supremacy of Murphy's law for unattended systems operating in hostile environments and on what it means for system design.

1 Introduction

Monitoring terrestrial high-arctic ecosystems is important, as they are particularly exposed to climate changes [6], but it is challenging. Indeed these regions are hard to access by humans, they are subject to extreme outdoor conditions, and they are at the very boundary of the global satellite-based communication infrastructure. As a result, year-round manual measurements are impossible, manually tapped data loggers are unreliable [3], and remote supervised control is impractical. How to obtain measurements in such a hostile environment at a cost, resolution, and utility that makes them useful for scientists? This is the challenge we tackle in the MANA project.¹

In MANA, we focus on monitoring limnic parameters (chlorophyll fluorescence, dissolved oxygen, temperature, salinity) in the Zackenberg region, in North-East Greenland (74.5° N).² For ten years, a long-term monitoring program has focused on the water chemistry and occurrence of freshwater biota in two small lakes during summer. These measurements, however, do not adequately describe the processes in the lakes, especially not in the periods when the lakes are covered with ice. In order to increase temporal resolution, we aim at introducing an automatic monitoring system that obtains measurements continuously, all year round.

¹A collaboration between ITU, Fresh Water Biology Lab at U.Copenhagen, Reykjavik University, DanSystem and Arch Rock Corp, funded by the Danish Strategic Research Council. See <http://www.itu.dk/mana>

²See <http://www.zackenberg.dk/>

Ideally, our goal would be to design a data acquisition system that operates in a high-arctic environment at high resolution, low cost, and high utility. In the context of lake monitoring, cost is dominated by sensors (optical sensors to measure chlorophyll concentration cost over \$10,000), which also restricts the potential for high spatial resolution (using a robot to move an expensive sensor instead of deploying multiple cheap sensors). So, we turn our focus on high utility.

Our goal of high-utility cannot be met with the current generation of data loggers. Their shortcomings are twofold. First, traditional data loggers implement a best effort approach: raw measurements are stored as they are received. Scientists need to clean the collected data to find out about outliers, missing values, or mis-calibration [11]. At that point, the best scientists can hope for is to interpolate the valid data points in order to compensate for the unusable ones. Our idea is that the data logger should check the data it collects, and take compensating actions when necessary. This way, the data logger can control the utility of the data it collects.

Second, traditional data loggers implement a static sampling strategy, i.e., the timing (sampling rate) and modality of the sensor measurements are fixed and remain unchanged throughout the data collection campaign. A fixed sampling rate is either set high enough to capture all interesting events – in which case too much energy is spent in the phases where the underlying system behaves as expected, or it is set too low – in which case some interesting events will not be captured. Our idea is to let the data logger decide on whether to reduce the sampling rate to save energy (e.g., the data logger might choose to reduce the sampling rate of the temperature sensors if it can predict the measured temperature), or to improve the sampling rate when it has detected an anomaly in the measurements (e.g., if chlorophyll concentration typically remains constant during the winter season then a slight increase should trigger an increase of the sampling rate to enhance the quantity of data points related to this potentially interesting phenomenon).

In order to collect useful data in a high-arctic environment, we find it necessary to let the data logger check the data it collects and adapt its sampling strategy to meet utility criteria defined by the scientist. We detailed our approach in [2], and we summarize our findings in Section 2.2.

Even if collecting useful data is our overall goal, a first, crucial step is to design a system that is capable of collecting data over the course of a whole year. Our initial system design, detailed in Section 2.1, was driven by one key challenge: the extreme weather conditions, specially in winter (September-July), where the wind is blowing, temperature reaches -40°C , and the lakes are covered with an ice lid of approximately 2 m as well as a layer of snow. A wire between sensors inside the lake and a data logger located on the shore would be exposed to the forces that are applied on the ice forming at the surface of the lake (during freeze-thaw periods or when the wind is blowing). A popular option, in the context of ocean and coastal waters monitoring, is to attach sensors to a buoy that integrates data logging and long-range communication capabilities. This is not appropriate in our case because of the poor satellite coverage and expensive transmissions. We investigate a slight variation of this approach, where one data logger located on the shore is connected via short-range wireless links to a buoy equipped with various sensors.

Our deployment builds on previous experiences and lessons learned in the sensor network community (e.g., [13, 15, 7]). We refine the lessons from the Sensorscope project [1]. Specifically, we insist on the importance of system modularity in mitigating various negative effects of a remote and hostile environment. Compared to sensor networks previously deployed in extremely cold environments such as permafrost monitoring in alpine environment [14], or glacier monitoring in the Swiss Alps [1], in Norway [10] or Iceland [9], our project is characterized by the complete lack of communication infrastructure, the absence of logistics support, and a nine months period where operator intervention is impossible. There is a lot we can learn from the instruments built by physicists for operations on remote planets (e.g., the wind gauge of the Mars pathfinder [12]), or in polar regions (e.g., the AMANDA [5] or IceCube neutrino telescopes [4]). As they do, we aim at a thorough understanding of the instruments we create and at establishing systematic procedures for their deployment and operations. However in MANA, our focus is on leveraging software rather than precise, reliable data acquisition hardware to improve data utility (e.g., a mean-time to failure of 40 years was observed for AMANDA’s optical modules). Also, MANA remains unattended for periods of 9 to 12 months, while, even extremely remote, these instruments are constantly controlled in-situ or remotely.

In this paper, we describe our system and we report on the lessons we learnt from our first exploratory



Figure 1: The Capoh base station (and buoy in the background) deployed at Zackenberg in October 2008 and August 2009.

deployment. We emphasize the issues we underestimated initially: i.e., the consequences of operating in a remote region, the impact of extreme weather not only on system design but also on operator activities, and the demands that derive from the absence of communication infrastructure. We also reflect on the impossibility of achieving high availability for a system that remains unattended for long periods of time in a hostile environment.

2 The MANA Project

2.1 System Design

Our two main objectives were to develop (1) an automatic monitoring system capable of operating in a remote environment under extreme weather conditions and (2) an autonomous system capable of compensating for failures and adapting the sampling strategy to changes in the environment. In order to monitor limnic parameters the sensors have to be submerged, hence we base our system on an anchored buoy, thereby keeping the sensors fixed at a known position, and at the same time protecting them against the strong wind and motion of the ice during thaw and freeze periods. We keep the sensors at a fixed depth of 2 m under the surface to minimize the risk of damaging the sensors against the bottom of the lake and contaminating them with dirt (the lake is 6 m deep). Also, moving the sensors away from the surface allows for year-long monitoring as the lake freezes from the top to the bottom to an estimated peak depth of 1.8 m.

Of course, it is possible to integrate processing and long-range communication capabilities into the buoy itself, however, the massive amount of batteries and energy harvesting needed to provide enough power for such a system, would add significantly to the size and weight. This would add to the cost of deployment and maintenance, due to the lack of infrastructure and due to the man power needed. Instead we opted for a tiered sensor network approach, where the buoy acts as a sensor node with minimal processing and storage capability while a base station at the shore handles the main processing, data storage, and long-range communication (see Figure 2 (b)). We use low-power short-range radios for communication in the sensor network since cables would be vulnerable to ice formations and acoustic modems are inoperable in frozen lakes. The sensor network approach offers several benefits: First, by minimizing the size and weight of the buoy, two people can easily deploy and maintain buoys in several lakes with only one inflatable rubber boat as opposed to a real boat in each lake. Second, a single base station can act as a common gateway for several buoys and other sensor nodes as well, thereby minimizing cost, facilitating data download, and enabling collaboration between sensor nodes. While the system currently deployed relies on a single buoy,




Figure 2: Schematic Diagram of the Capoh buoy and base station components.

our architecture allows for several buoys or sensor nodes around the lake. Third, heavy equipment, such as batteries and energy harvesting, can be better protected when put firmly on the ground and for lower cost than when submerged on a buoy. Also, access becomes easier when equipment is not submerged or buried beneath ice.

The buoy we developed consists of a flotation device and a water proof steel housing connected with a 2 m long steel pipe. The sensors are attached outside the steel housing and protected from debris by a steel cage. By placing the electronics and batteries at the bottom, inside the steel housing, the water acts as insulation against the cold thereby maintaining the batteries' discharge capacity. Another power saving feature is a water proof switch which allows the buoy to be assembled and transported without using power until eventually deployed. Finally, an antenna is connected on top of the flotation device as well as an anchor to keep the buoy in place (see Figure 2(a)).

As our water sensors we used the Wetlabs Water Quality Monitor (WQM),³ which measures conductivity, temperature, depth, dissolved oxygen, chlorophyll, and turbidity. We chose the WQM because of its open protocol, precision, low-maintenance, and sturdiness. Specifically, the built-in anti-biofouling mechanism prevents plants from growing on the sensors. This is extremely valuable in an unattended environment where maintenance and calibration cannot be performed with the same frequency as in other ecological monitoring programs. Another reason, is the WQM's ability to function as a data logger, which means it can automatically take regular samples at fixed intervals and store them internally for later retrieval. Although the purpose of this project is to move beyond simple data loggers, by building on top of these, we add an extra layer of storage redundancy which serves as a backup in case of network outage and catastrophic failure on the base station. Last, the WQM is a high-grade instrument of the same quality and built as the field instruments ecologists regularly use. By building a sensor network around instruments the ecologists are used to and trust we increase the quality of the measurements collected.

We used the Arch Rock IPserial node to build our sensor network infrastructure. The IPserial node uses 6LoWPAN to automatically form a multi-hop routing network and low-power-listening to duty-cycle operations. The defining feature of the IPserial node is the wireless serial (RS-232) connection the gateway can establish with any sensor node connected through the network. Because current state-of-the-art data loggers

³See <http://www.wetlabs.com/>

(including the WQM) and Measurement and Control Systems (MCS) from the industry leader Campbell Scientific⁴ all allow serial connection, establishing a wireless serial infrastructure will effectively turn all data loggers and MCSs within range into potential sensor nodes, all accessible from the base station.

The downside of having a base station on land, instead of in the lake, is the exposure to the environment. Without the water offering insulation against the cold, and protection from the wind and wild life, special care must be taken to protect the base station. We used a water proof flight case with 10 cm insulation to protect the batteries and electronics. Without this insulation the batteries' discharge capacity would be greatly reduced. The flight case itself was grounded by surrounding it with large rocks and painted white to further minimize heat radiation. We calibrated the power subsystem to operate autonomously throughout the winter (discharge only) with recharge cycles throughout the summer. We chose 6V EXIDE⁵ batteries because they maintain a high capacity (240 Ah nominal) at low temperature, even through cycles of deep discharge. For energy harvesting, we opted for a solar panel. This might seem counter intuitive to rely on solar power in the arctic, however, this is the most popular form of energy harvesting in the region (according to PolarPower.org). Indeed, wind is not attractive because frequent storms require a heavy infrastructure for a relative low wind power on average. We placed a motorized solar panel on top of a 2.5 m pole and firmly secured it to the ground, both at the base and with wires, by drilling directly into the frozen rocks. By placing the solar panel above ground, we kept it free from snow cover and protected it from both curious animals and flying debris. Furthermore, changing the angle of the solar panel with the electric motor, reduces the cross section, which the wind can catch on to. Due to the Sun being absent during winter, when the strongest wind gusts are usually observed, this has little impact on the harvested energy. Three antennas (WiFi, sensor network, and GPS) were also put on top of the pole and a reinforced plastic tube was used to protect the cables between the pole and flight case against the local animals, which are notorious for biting in cables. Note that we used an omni directional antenna on the buoy and a directional antenna on the base station as initial tests with an omni antenna were not satisfying. We should also note here that while there is rich information online about designing power subsystems in the arctic, there is virtually no information widely available about wireless communication in the arctic⁶. This should be a topic for future work.

As platform for our autonomous data acquisition system we used the Vexcel MicroServer from the Seamonster Project,⁷ which has been specifically designed to function in remote environments. The water proof MicroServer contains a Single Board Computer (SBC), Power Control Board (PCB), GPS, solar panel charger, and amplified WiFi. The SBC runs Linux and is powerful enough to function as both sensor network gateway and controller. The main advantage of the MicroServer lies in the synergy between the SBC and PCB, where the former can control the power state of all the peripherals and even schedule shutdowns and reboots of itself. This allows very efficient power usage, where components only are switched on when needed and otherwise switched off. This includes the PCB as well, which can enter a low power state with everything turned off except an internal timer, that wakes up both the PCB and SBC when fired. The amplified WiFi is powerful enough to reach the field station 4 km away. Note though that the field station's satellite connection does not have enough bandwidth to sustain a remote download of all the data. To keep time, we use the real-time clock in the SBC for day-to-day use, and to minimize drift we synchronize it with the GPS clock, but only once a month in order to save power.

2.2 Autonomous Data Acquisition

With the sensor network infrastructure in place, we have the means to remotely obtain measurements continually, monitor the status of the system, and upload new commands to the WQM or any other sensor node in the network. In order to increase the utility of the collected data and make the time series more useful for the scientists, our goal is to continuously monitor the collected data and take compensating actions when necessary.

⁴See <http://www.campbellsci.com/>

⁵See <http://www.exideindustrialbatteries.com/>

⁶The most relevant information concerns wireless setups for remote and rural networks, e.g., <http://wire.less.dk/>

⁷See http://robfatland.net/seamonster/index.php?title=Vexcel_Microserver




Figure 3: Voltage readings during Winter 2009




Figure 4: Internal temperature readings during Winter 2009

Ideally, our sensor network controller should take the same set of actions as a scientist would have done if she had been present in the field herself. To accomplish this, we use the scientist's own requirements to the collected time series to drive the decisions taken by the controller. These requirements are stated in the form of ordered collection modes and a target lifetime. By considering the priorities of the collection modes we can recast the controller problem into a Constraint Optimization Problem (COP), since choosing a set of actions which satisfies the collection modes, and at the same time honors the energy and time constraints, is essentially an assignment problem.

Each time new data is collected, different collection modes become active depending on the activation triggers. Together with the current energy and time state a new COP is constructed. By solving this COP we obtain a set of actions which adaptively changes the sampling strategy based on the scientist's requirements. For further details we refer the reader to [2].

2.3 Deployment

The MANA project started on February 1st 2008. The Capoh system components were shipped to Greenland at the end of June 2008. The short development time led us to focus on the survivability of our system rather than its performance. The many unknowns concerning the environment led us to over-dimensioning the power subsystem and the mechanical components. These unknowns prevented us from conducting tests that could cover the range of operating conditions in Zackenberg.

We installed the Capoh system in August 2008. However, due to unforeseen weather conditions the system could not be properly tested (see Section 3). Hence, a second visit in October 2008 was planned to ensure that the system was working as intended. These tests however did not allow us to establish contact between the base station and the buoy. The freezing conditions did not allow us to perform a thorough maintenance operation. We had to wait until August 2009 to maintain the system. We established then that (a) the system was intact from a mechanical point of view, that (b) the buoy had been stopped at the time of the installation, and (c) that the base station had become unresponsive. As a result no data was collected from the buoy in the first season. The tests conducted in August 2009, after a successful maintenance, allowed to collect in two days as many data points than in that last 10 years of manual data collection. The system was fully operational at that point.

The system data, collected from the base station, showed that the SD Card storing the OS from the Microserver had become corrupt, forcing an endless cycle of reboots. According to the last entry in the logs this happened at the end of February. Further examination showed that the internal timer in the PCB had malfunctioned mid-February, causing the PCB to wakeup the SBC more frequently than planned. This error caused the base station to power up more than twenty times a day instead of only four as planned, draining the battery supply to depletion as show in Figure 3. At first, we thought extremely low temperatures had caused the malfunctioning timer, but log of the temperature inside the base station (Figure 4) does not show any significant event. In fact, the lowest temperatures were recorded in the second week of January, where both the PCB and SBC worked flawlessly. After finding bite marks and traces of oxidation on some of the external cables not properly protected, we suspect that the exposed uncut wires, combined with moisture, must have caused a short-circuit in the system, which in turn caused the internal timer to malfunction. This is corroborated by the fact that further testing of the PCB showed it had returned to perfect working condition, while we were not able to reproduce the erratic voltage readings. We subsequently encased all wires in reinforced plastic tubes to avoid similar problems. We suspect that all these factors contributed to causing the SD card to become corrupt.

3 Lessons Learned

Remote Region. An implication of operating in a remote region is that transporting hardware is costly, both time-wise and money-wise. The Zackenberg research station can be reached by a plane (Twin Otter), once a week, during the summer. As a consequence heavy equipment needs to be freighted by ship. For our deployment mid-August 2008, the buoy and base station components were shipped in early July. Beyond the obvious implication in terms of deadline for system design and test, freight had an impact on system design. In our system, a mote is embedded in the buoy. The buoy needs to be safely sealed so that it is water tight. At that point, the mote is no longer physically accessible. We decided that the buoy should be sealed before it was freighted to make absolutely sure that it was water tight. As a consequence, we incorporated an on/off switch in the buoy so that we could activate the mote at deployment time. We underestimated the need for an immediate, low-cost feedback of the on/off switch (through a brief visual or auditive feedback), as making sure that the buoy was actually turned on required the entire system to be setup. This unnecessary dependency turned out to be a strong constraint in terms of our timetable.

Lesson 1 *At system deployment time, it should be possible to test and deploy each subsystem independently.*

Another example of the cost of transport is that the lake we are instrumenting is located 4 km from the Zackenberg research station. How to transport 300 kg of hardware (including 120 kg of batteries, a 1 m³



Figure 5: Helicopter carrying the system hardware components to the lake

box, a 2 m pole and a buoy full of electronics) along those 4 km turned out to be a logistics nightmare. Without snow cover, using an Argo (utility terrain vehicle) is prohibited on the permafrost. Transporting the hardware by hand would have taken several days even with several people working full time and with a high risk of injury, as the terrain is very rocky and the hardware cumbersome to transport manually. We thus opted to use a helicopter (see Figure 5). Bad weather and logistics constraint meant that we had to wait three days before the hardware could be lifted to the lake which also turned out to be a major setback for our timetable (and our budget as a 30 min flight costed \$4,000).

Another consequence of deploying a system 4 km from a research station that can be reached once a week is that we cannot afford to change our plans once we are in the field. We must (a) plan our work meticulously in advance and solely execute it in the field, and (b) carry redundant hardware from our lab to the research station so that we allow for some replanning (i.e., finding solutions to unforeseen issues) there. Agile methodologies are not applicable in this context.

Weather Impact. Our system design was centered around the constraints of the extreme weather of the high-arctic (buoy with a 2 m long pole to tolerate ice, rotating solar panel to tolerate wind, operating range down to -40°C for every component of the base station). We were very much aware that any silent assumption we made would turn out to be wrong. While we did pretty well anticipating pitfalls in terms of mechanical design (the system was intact after a year in the field - the buoy was water tight, the Water Quality Monitor in very good condition, and the solar panel functional), we did not account for the pitfalls we would face in terms of operations (deployment, testing and maintenance). During our first stay at Zackenberg, in August 2008, we could not reach the lake during three days out of a week-long stay because of heavy rains that made the walk from the station to the lake dangerous. During that stay, the helicopter was also delayed for two days because of heavy wind, and an additional day because of the backlog of tasks it should complete. That left us one day to install, test, and deploy the system. We had planned our second stay at Zackenberg, in October 2008, for testing. However, the freezing temperature (around -20°C) made it impossible to manipulate the buoy (since digging it free from the ice would not have helped because the water on the steel housing would have frozen immediately upon reaching the surface), or even to operate our laptop outdoors (as the LCD screen froze). In retrospect, staying only one week for the initial installation, that we had evaluated to (and turned out to be) a 4 day task, was overly optimistic as we did not account for the delays that the weather would incur. More generally, system design should be driven by the limited time window for operation:

Lesson 2 *It should be easy to troubleshoot in the field using observation points to test each component separately*

No Communication Infrastructure The communication infrastructure at Zackenberg is limited to (a) UHF radio that covers most of the Zackenberg valley and the mountains surrounding it, and (b) satellite voice and data services (through private satellite phone or a satellite access point managed at the station) that are expensive and unstable as Zackenberg is at the very boundary of satellite coverage. This lack of communication infrastructure extends the consequences of working in a remote region to also affect data and software access via the Internet. There is no sense in maintaining a repository of our software online for backup or troubleshooting purposes as we would not be able to access it from Zackenberg. Again, we must carefully plan for self-contained deployment, in-the-field testing and troubleshooting of software. The lack of communication infrastructure also impacts system design in the sense that our sensor network infrastructure can only achieve spatial coverage via a mesh constituted of the base stations and sensor nodes that we deploy. There is clearly an advantage in multiplexing the resources of a base station across several sensor nodes as each base station is costly and cumbersome to install due to the batteries needed for its operation.

Murphy’s Law Supremacy Lessons learned from previous deployments [7, 1] (including our own [8]) made it clear that everything that can go wrong will at some point go wrong (the so-called Murphy’s law). For long term deployments in hostile environments, this observation should have a deep impact on system design. When a system operates unattended for long period of time in a hostile environment, then availability cannot be a design goal.

Lesson 3 *The primary goal in terms of both Hardware and Software Design should be maintainability.*

Guarding the system against most problems is of course desirable, but guarding the system against all (combinations of) problems is impossible. Indeed, we do not have a model of all environmental factors that impact our system: they range from polar bears to lemmings, and from snow storm to month-long ice covers. Instead of aiming at availability, we should design our system for maintainability in the short time windows where we can actually perform maintenance.

Lesson 4 *Strict modularity should be enforced at the hardware and software levels: (a) faults should not be allowed to propagate across components, and (b) each component should be replaceable independently*

4 Conclusion and Future Work

The key lesson that we learnt from our first deployment in the context of the MANA project is that maintenance/operations should not be afterthoughts. They should have a deep impact on system design. Indeed, for long term deployments, high availability is an unattainable ideal making maintenance necessary as well as extremely constrained: maintenance must be conducted within short time windows, without outside support, in an unwelcoming environment. Systems deployed in hostile environments should thus be designed for observability and maintainability, enforcing a strict modularity both in hardware and software. There is obviously a trade-off here between modularity and performance or energy consumption. Exploring this trade-off should be a topic for future research.

We are working on a second generation of Capoh system based on the lessons presented in this paper. Our goal is to improve maintainability and to rightsize the energy budget and thus build a system that can be integrated in a long term monitoring program at Zackenberg. We are also planning new deployments of the Capoh system in West-Greenland.

Acknowledgements

Ole Guldbrandsen from Dan-System designed the mechanical components of the Capoh system, and prepared the deployment. He made all the right decisions. Wei Hong and his team at Arch Rock helped a lot smoothe the integration of their product in our prototype. Kirsten Christoffersen and Anders Birk Nielsen from Fresh Water Biology Lab at U.Copenhagen took a major part deploying the system in August 2008. Rob Fatland provided us with the microservers. He has been responsive and helpful throughout the project. The comments from the anonymous reviewers allowed us to significantly improve this paper.

References

- [1] Guillermo Barrenetxea, François Ingelrest, Gunnar Schaefer, and Martin Vetterli. The Hitchhiker's Guide to Successful Wireless Sensor Network Deployments. In *SenSys '08: Proceedings of the 6th ACM conference on Embedded network sensor systems*, pages 43–56, New York, NY, USA, 2008. ACM.
- [2] M. Chang and Ph. Bonnet. Meeting scientific requirements with adaptive data acquisition. Technical Report TR-2009-122, 2009.
- [3] Carsten Egevang, Iain J. Stenhouse, Richard A. Phillips, Aevan Petersen, James W. Fox, and Janet R. D. Silk. Tracking of Arctic terns *Sterna paradisaea* reveals longest animal migration. *Proceedings of the National Academy of Sciences*, pages –, 2010.
- [4] A. Achterberg et al. First Year Performance of the Icecube Neutrino Telescope. *stroparticle Physics*, 26, October 2006.
- [5] E. Andres et al. The AMANDA neutrino telescope: principle of operation and first results. *Astroparticle Physics*, 13(1), 2000.
- [6] Toke T. Høye, Eric Post, Hans Meltofte, Niels M. Schmidt, and Mads C. Forchhammer. Rapid advancement of spring in the high arctic. 17(12):R449–R451, 2007.
- [7] K. Langendoen, A. Baggio, and O. Visser. Murphy Loves Potatoes: Experiences from a Pilot Sensor Network Deployment in Precision Agriculture. In *Parallel and Distributed Processing Symposium, 2006. IPDPS 2006. 20th International*, pages 8 pp.–, April 2006.
- [8] Klaus S. Madsen Philippe Bonnet Marcus Chang, Cécile Cornou. Lessons from the hogthrob deployments. In *Second International Workshop on Wireless Sensor Network Deployments (WiDeploy2008)*, 2008.
- [9] Hart J. K. Martinez, K. and R. Ong. Deploying a Wireless Sensor Network in Iceland. *Lecture Notes in Computer Science, Proc. Geosensor Networks*, 2009.
- [10] Padhy P. Elsaify A. Zou G. Riddoch A. Hart J. K. Martinez, K. and H. L. R. Ong. Deploying a Sensor Network in an Extreme Environment. 2006.
- [11] A. Sharma, L. Golubchik, and R. Govindan. On the Prevalence of Sensor Faults in Real-World Deployments. In *Sensor, Mesh and Ad Hoc Communications and Networks, 2007. SECON '07. 4th Annual IEEE Communications Society Conference on*, pages 213–222, June 2007.
- [12] R. Sullivan, R. Greeley, G. Wilson, P. Smith, and C. Cooper. Imager for Mars Pathfinder Windsock Experiment. In *Lunar and Planetary Institute Science Conference Abstracts*, volume 27 of *Lunar and Planetary Inst. Technical Report*, pages 1289–+, March 1996.
- [13] Robert Szewczyk, Alan M. Mainwaring, Joseph Polastre, John Anderson, and David E. Culler. An analysis of a large scale habitat monitoring application. In *SenSys*, pages 214–226, 2004.
- [14] Gruber S. Talzi I., Hasler A. and Tschudin C. PermaSense: Investigating Permafrost with a WSN in the Swiss Alps. In *Proceedings of the Fourth Workshop on Embedded Networked Sensors, 25–26 June, 2007, Cork, Ireland*, 2007.
- [15] Geoffrey Werner-Allen, Stephen Dawson-Haggerty, and Matt Welsh. Lance: optimizing high-resolution signal collection in wireless sensor networks. In *SenSys*, pages 169–182, 2008.