

Structural monitoring with wireless sensor networks: Experiences from field deployments

Glauco Feltrin¹, Olga Saukh², Reinhard Bischoff^{1,3}, Jonas Meyer^{1,3}, and Masoud Motavalli¹

¹ Empa, Swiss Federal Laboratories for Materials Science and Technology, Duebendorf, Switzerland

² ETHZ, Swiss Federal Institute of Technology, Zurich, Switzerland

³ DecentLab GmbH, Duebendorf, Switzerland

ABSTRACT: In the last years, wireless sensor networks have emerged as a promising technology that is inducing a deep innovation in the field of structural monitoring. The main advantages of wireless sensor networks are fast deployment, little interference and self-organization. However, since wireless sensor nodes are battery powered, in long term monitoring applications the power management influences significantly the operation of a wireless sensor network. In data intensive applications, e.g. vibration based monitoring, low power hardware, duty cycle operation and efficient communication policies are not sufficient for achieving a sustainable system lifetime. Since data communication is the most energy consuming task, long system lifetimes can only be achieved by a significant data reduction in the nodes. This data reduction is a challenging task, since it has to be performed with very limited computational and memory resources and in competition with tasks providing the basic network functionality. The objective of the paper is to provide a brief overview of the wireless sensor network technology and to present our experience over the past three years with data intensive structural monitoring using wireless sensor networks. Deployments on two bridges are illustrated and specific aspects of sensing, data quality, stability, availability, and system lifetime are analyzed.

1 WIRELESS SENSOR NETWORKS

1.1 Introduction

A wireless sensor network (WSN) is a computer network consisting of many small, intercommunicating computers equipped with one or several sensors (Culler & Wei 2004). Each small computer represents a node of the network and is commonly called sensor node. The communication within the network is established using radio frequency transmission. All sensor nodes are equipped with specific sensors tailored to their measurement tasks. One or several sensor nodes act as root nodes and represent the data sink in the network.

The typical hardware components of a sensor node are the sensors, a signal conditioning unit, an analog to digital conversion (ADC) module, a central processing unit (CPU) with random access memory (RAM), a radio transceiver and the power supply (see Figure 1a). A physical implementation of a sensor node is displayed in Figure 1b. Various hardware wireless sensor networks platforms are commercially available today, which offer the basic functionality together with specific sensing capabilities (Bischoff et al. 2009). The diversity of platforms offers the possibility to choose a platform which best fits the needs of a specific application.

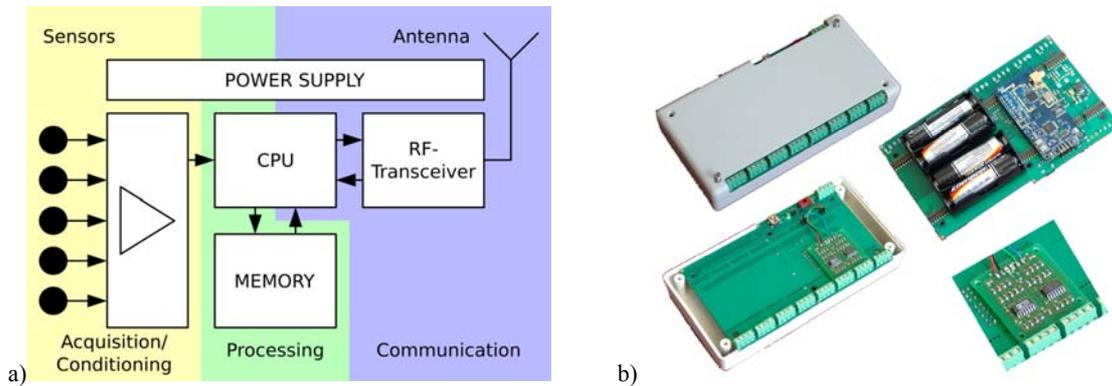


Figure 1: a) Hardware architecture of a sensor node. b) Physical implementation of a sensor node.

Contrary to conventional monitoring systems, which usually have a centralized system configuration and data acquisition software, a WSN is a distributed data acquisition system. Each sensor node executes software that provides the functionality for performing many tasks:

- scheduling and execution of the measurement tasks
- signal conditioning and data acquisition for different sensors
- temporary storage of the acquired data
- data processing
- self monitoring (e.g. supply voltage, communication link quality)
- time synchronization of the network
- management of the data acquisition and processing configuration (e.g. changing the sampling rate, reprogramming of data processing algorithms)
- reception and forwarding of data packets
- coordination and management of communication and networking.

The sensor node software of a WSN for civil structures that supports different sensors and data processing algorithms is described in Feltrin et al. 2010. Because of the very limited memory resources, the software is usually tailored to the specific sensors and tasks of a sensor node in order to keep its size as small as possible. Often, the software is written in NesC (Gay et al. 2003), an extension to the C language. The basic functionality like time synchronization, multi-hop functionality etc. is provided by TinyOS (Levis et al. 2005), a widespread operating system for WSNs that has been ported to many WSN platforms. TinyOS is highly tailored to the limited resources of the node hardware. NesC as well as TinyOS are both Open Source projects (www.tinyOS.net).

1.2 Power management

Since the sensor nodes have to be operated with batteries, the power supply is very limited. Power saving is therefore of utmost importance in designing, implementing and operating WSN based monitoring systems for medium and long term deployments, since too frequent battery changes increase the maintenance costs severely compromising one of the main advantages of WSN monitoring systems.

Energy consumption is reduced by using low power hardware (sensors, microcontrollers, radio chips) for implementing sensor nodes that consume typically significantly less than 100 mW (Bischoff et al. 2009). Another means to reduce energy consumption is to operate the network in

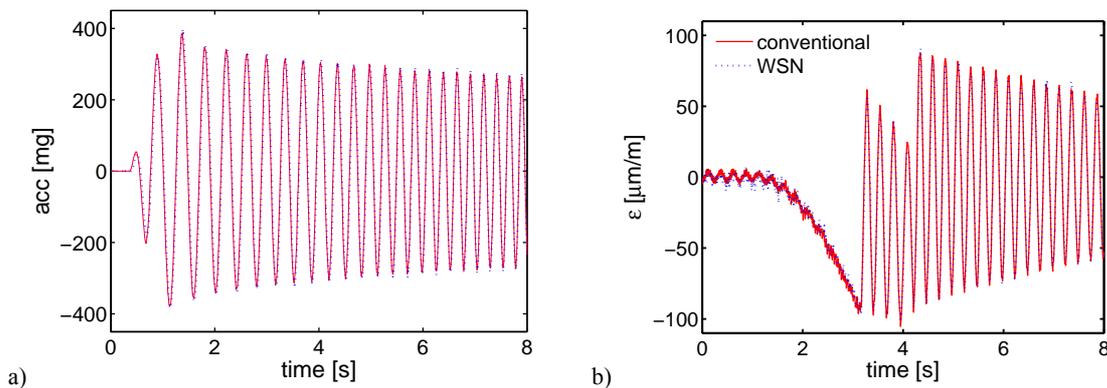


Figure 2: a) Acceleration records obtained with a WSN equipped with a low cost MEMS sensor and a conventional data acquisition device with a conventional piezoelectric accelerometer. b) Strain records obtained with a WSN and a conventional data acquisition device.

switched-off state for a significant amount of time and by switching-on hardware components only for the time period required to complete the scheduled tasks.

Communication is the most energy consuming task. Since the energy consumption increases exponentially with the transmission distance (exponent greater than 2), power can be saved by operating WSNs as multi-hop networks (Culler & Wei 2004). By establishing communication links only to neighbor nodes, thus reducing the transmission distance, the data are sent to the data sink through several nodes, which act as relaying stations, receiving and forwarding data from adjacent nodes.

Finally, power can be saved by processing the raw data in the node with the goal to reduce the amount of data that need to be transmitted (sending of information instead of raw data), since sending 1 bit costs as much energy as executing about 1000 instructions of a low power microprocessor (Culler et al. 2004). When monitoring vibration based processes, which produce large samples of raw data, this strategy is the most powerful energy saving method and is mandatory if a system lifetime of several months should be achieved (Straser & Kiremidjian 1998; Feltrin et al. 2006; Lynch et al. 2006; Spencer & Nagayama 2006).

1.3 Sensing

Several commercial WSN platforms offer specific low power data acquisition boards with integrated MEMS sensors and signal conditioning circuits for sensing and acquiring temperature, humidity, light intensity, gas pressure, or accelerations. MEMS sensors have several advantages compared to conventional sensors: They are small, generally low power, highly integrated and, if no high end accuracy and resolution is required, rather inexpensive.

Several commercial MEMS accelerometers have the characteristics to be suitable for structural monitoring applications. MEMS accelerometers with an amplitude range of several g have typically a sensitivity of 0.5...1 V/g, can be powered with a low DC voltage of 3 to 5 V, which can be provided with standard batteries, and their power consumption is of the order of several mW.

Figure 2a displays a comparison of the performance of a 10 € MEMS (LIS2L06 of ST Microelectronics) accelerometer and a 500 € conventional piezoelectric accelerometer (PCB M393 A03). The accelerometers were mounted on a shaker that performed a sweep excitation between 2 and 50 Hz. The accelerations of the MEMS accelerometer were recorded with a sensor node equipped with a 12 bit ADC and those of the piezoelectric accelerometer with a

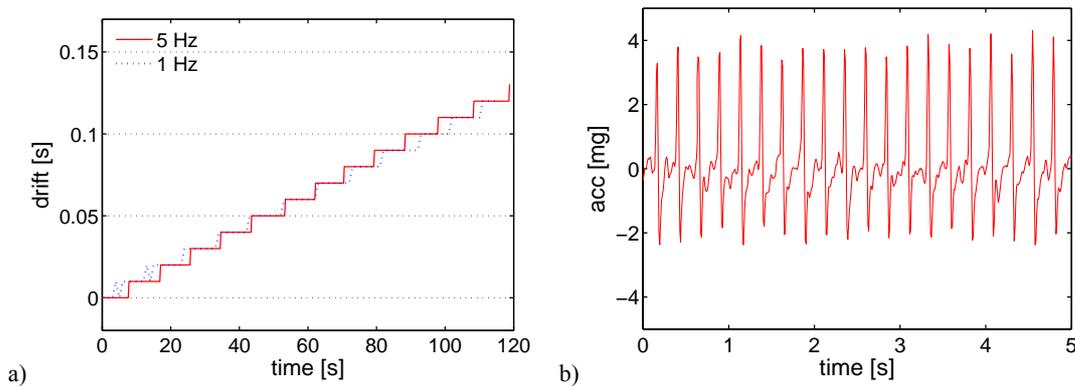


Figure 3: a) Time drift of a sensor node. b) Signal corruption by the duty cycle of the radio.

24bit high end data acquisition device. The agreement displayed in Figure 2a is very good. In the investigated amplitude range both accelerometers are essentially equivalent.

A sensor node can be designed to support also conventional sensors. A typical widely used sensor in structural monitoring is the electrical strain gage. Figure 2b displays strain records obtained with a WSN and a conventional high end data acquisition device. The quality of the WSN record is comparable to the record obtained with the conventional system. The average difference between the two records has a magnitude of a few percent. Unfortunately, because of its low resistance (typically 120 Ω), strain gages are rather power consuming (approx. 40 mW).

1.4 Data acquisition

A critical aspect of data acquisition is the analog to digital conversion (ADC) since it has an important impact on the resolution of the digital data. Modern wired data acquisition devices are typically equipped with 24 bit AD converters. In contrast, mainly for reducing the power consumption, WSN platforms are provided with 8 or 12 bit AD converters. This represents a quite stringent limitation that requires a careful balance between resolution and amplitude range. At low amplitudes the quantization effect of the AD converter usually limits the resolution of the whole data acquisition process. This effect manifests itself as noise and is displayed at small amplitudes in the record displayed in Figure 2b (0 to 2 seconds). Contrary, at high amplitudes, Figure 2b does not display a significant difference between the data obtained by the sensor node and the high end data acquisition device.

A second aspect that has to be considered is time drift due to an inaccurate setting of the sampling rate of the WSN platform. The inaccurate setting of the sampling rate is an effect of the quantization of time due to the finite frequency of the quartz oscillator (32768 Hz). The effective sampling rate of the WSN ADC was tested by monitoring the change of the phase difference of a sine signal between the WSN and conventional data acquisition. Figure 3a) displays the time drifts computed by two sine signals with 1 and 5 Hz frequency. For both sine signals the time drift is very similar and reaches 0.13 seconds after 120 seconds. The relative time drift of the WSN data is therefore approximately 1.1%, which is a very small figure but still has appreciable effects on measurements that last for several minutes or have high frequency oscillations. Since the time drift is linear, it can be easily corrected by an adequate post-processing.



Figure 4: The Keräsjokk Railway Bridge in Sweden near Haparanda.

1.5 Duty cycle operation

Duty cycle operation is one of the main mechanisms for reducing power consumption in WSNs. The duty cycle describes the fraction of the switched-on time of a hardware component within a given period of time. A small duty cycle means that most of the time the hardware component is switched-off or in sleep mode. The smaller the duty cycle, the less energy is consumed.

Since communication is very power consuming, operating the radio with a low duty-cycle is a standard strategy to save power. However, because of the high power consumption, each time the radio is switched-on, the power supply voltage drops by a small amount (e.g. 0.1 V) due to the internal resistance of the power supply. If the sensors are connected to the same circuit that powers the WSN platform, which is the standard situation in commercial WSN platforms, the periodical voltage drop affects also the power supply of the accelerometer and, as a consequence, produces a variation of the output signal, which reduces the quality of the sensor signal. Figure 3b displays such a periodical acceleration signal noise. For ambient vibration measurements, the noise induced by duty cycle operation is very likely to severely degrade the quality of the data.

The signal corruption can be eliminated by separating the power supply of the sensors from the power supply of the WSN platform (Feltrin et al. 2009) or by powering the sensors with a voltage regulator circuit that keeps a constant output voltage, e.g. 3V, independently of the current supply voltage of the batteries. Furthermore, if a particular sensor requires a supply voltage that is greater than the supply voltage provided by the batteries, a voltage regulator boosts the battery output voltage to the correct voltage level. The latter solution cancels also the effect of the fading of supply voltage due to power consumption that is typical of batteries.

2 FIELD DEPLOYMENTS

The performance of wireless sensor networks in structural monitoring were demonstrated in several short-term field deployments (Glaser 2004; Kim et al. 2006; Lynch et al. 2006; Mechitov et al. 2006; Gangone et al. 2007; Pakzad et al. 2008). However, very little experience exists with long-term deployments (several days and more) that feature data intensive applications or in-node data processing. Two field deployments with these characteristics are presented in this section.

2.1 Strain monitoring on a railway bridge

The bridge over the Keräsjokk River is on the Haparanda railway line (single track, non-electrified) in northern Sweden. It is a single span, simply supported riveted steel truss bridge built in 1911 with a length of 31.6 m (Figure 4). The railway tracks lies on wooden sleepers

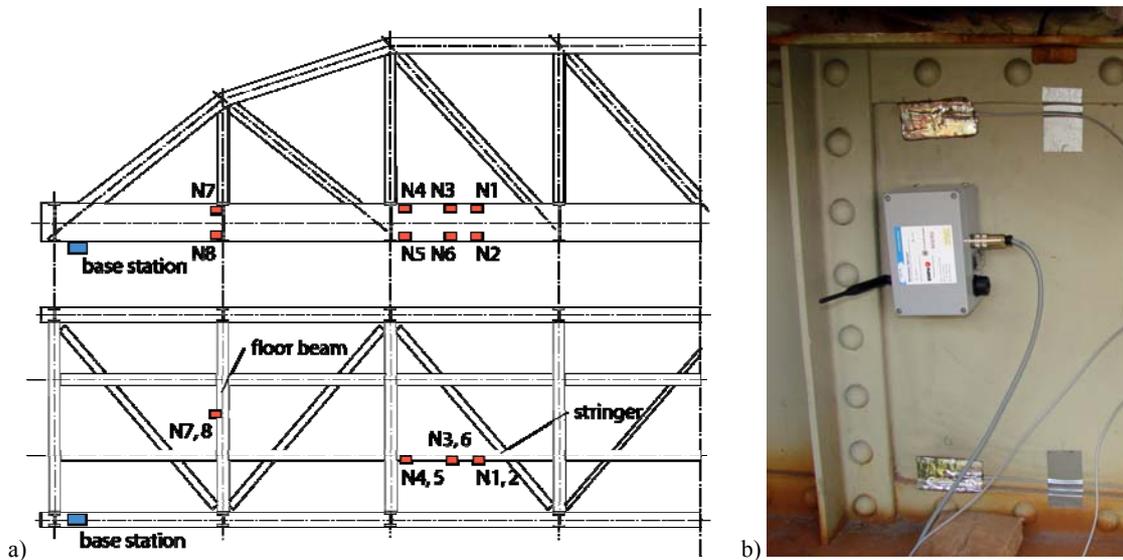


Figure 5: a) Location of strain measurements on the bridge. b) Sensor node and strain gage mounted to the bridge.

directly on the stringers (no ballasting). It was assessed for investigating their fitness for an axle load increase from 22.5 to 25.0 tons (Enochsson et al. 2008). The initial assessment showed that the fatigue capacity of the web joints in the floor beams would exceed their capacity. This bridge was chosen to demonstrate some of the assessment and monitoring methods developed within the EU-funded project Sustainable Bridges (www.sustainablebridges.net, Kiviluoma et al. 2007). The bridge was crossed by freight trains driven by diesel engines (one or two per day).

2.1.1 Monitoring system

The wireless sensor network deployed on the bridge consisted of 8 nodes and the root node connected to the base station. The location of the strain measurements, all mounted on secondary elements, are displayed in Figure 5a. Six strain measurements, labeled as N1 to N6, were performed on a longitudinal stringer and two measurements, labeled as N7 and N8, were mounted on a floor beam. The strain was measured with soldered strain gages featuring a resistance of 120 Ω . Figure 5b shows a deployed sensor node and a strain gage mounted on the stringer (positions N4 and N5). The strain gages were connected to the nodes of the wireless sensor network by cables, which were plugged into an external connector. The housings of the nodes were equipped with four magnetic footings, which allows a simple and fast mounting.

The sampling rate of data recording was 100 Hz and the record size was 30 seconds. Since one data item was stored in 16 bit word the total memory requirement of a complete time history was 6 kB or 60 % of the RAM memory of the microcontroller. The radio communication was periodically switched on and off to save energy. This duty cycle was synchronized over the whole network. 90 % of time the radio was shut down and only 10 % of the time the nodes had enabled the energy intensive radio communication.

All the data recorded by the eight nodes was sent to the base station. To avoid packet collision, a sending policy was chosen where each node has a time slot for transmitting its data. The wireless monitoring system recorded also temperature, humidity, supply voltage and network tree data with a time interval of 2 minutes.

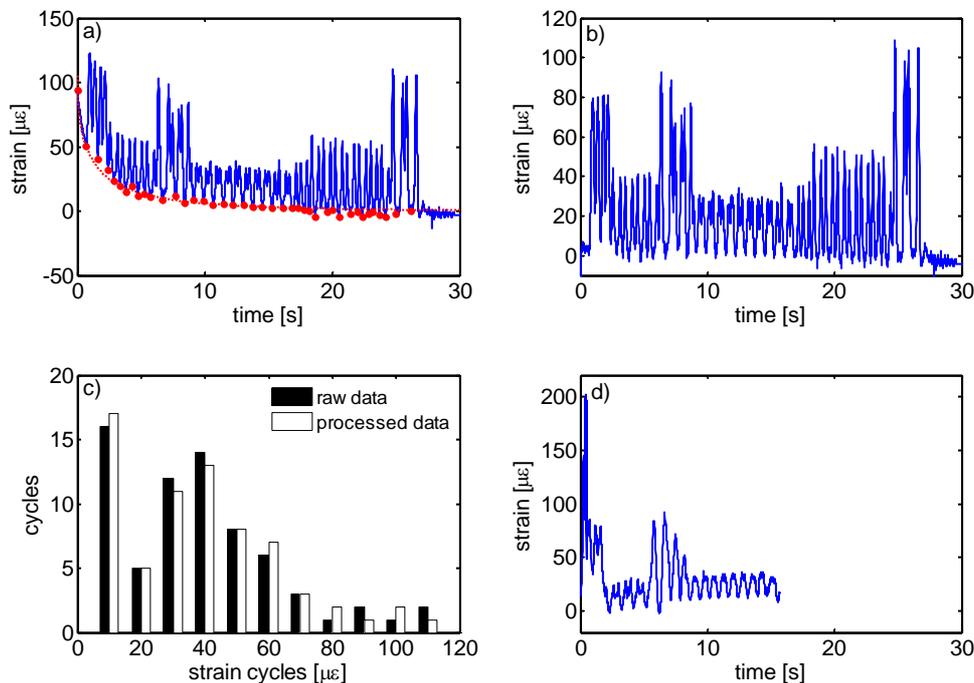


Figure 6: a) Raw strain data recorded at node 2 with bias and curve fitting for bias removal. b) Strain after removal of switch-on bias. c) Cycle counting using raw and post-processed strain data. d) Strain recorded at node 8.

The wireless monitoring system was installed over a period of 8 days. Since the strain gauges had been installed in advance, the installation was quite fast and took approximately 4 hours. The most time consuming work was the connection between strain gauges and the nodes (soldering of wires, checking of functionality).

2.1.2 Triggered measurement

Since electrical strain gages are quite power consuming, a 120 Ω strain gage consumes approx. 40 mW, a long-term operation require the implementation of a power saving mechanism. Because of the low traffic on the bridge, switching on the strain sensing hardware during train crossing would allow a significant power saving. This policy, however, requires a triggering mechanism that consumes significantly less power than a strain gage. We opted for equipping the sensor nodes with an ultra low power acceleration sensor that was running all the time and identified approaching trains through vibrations. We used the accelerometer Bosch Sensortec SMB380, a low cost tri-axial sensor for consumer market applications. Two features make this sensor particularly useful in this application context: good resolution, power saving capabilities with configurable duty cycle and on-chip signal processing capabilities. The power consumption is 500 μW at full operation mode and 3 μW at sleep mode. The sensor has a resolution of about 4 mg at a measurement range of $\pm 2\text{g}$. The wake-up time is 1ms.

The sensor has a built-in signal processing capability for analyzing the acquired accelerations. This feature was used to generate a hardware interrupt if a specific threshold is violated. The interrupt switch-on the strain sensing module (strain gage and signal conditioning) and starts the data acquisition. This wake up process lasted 100 ms. Each node was triggered independently.

2.1.3 Results

A typical recorded time series of the raw data at the measurement point N2 at the stringer (bottom, middle) is displayed on Figure 6a. Each axle of the train is clearly visible. The resolution of the strains was found to be approximately $1 \mu\epsilon$. The accuracy, which could not be tested, was estimated to be approximately $\pm 3 \mu\epsilon$. The achieved accuracy is good enough for assessment purposes and in particular for fatigue assessment using cycle counting based methods. The time series do not show significant dynamic effects, which is due to the small dynamic amplification factor.

The raw data, however, are biased by a significant time-dependent signal. The bias was generated by switching on the strain signal conditioning device shortly before the measurement was started. Due to the resistance of the strain gage, the current flow heats the strain gage increasing its temperature and consequently its resistance. The signal conditioning (Wheatstone bridge) translates this resistance change in a decreasing voltage signal. Since the heat produced in the strain gage flows into the surrounding metal and air, the temperature of the strain gage increases with increasing time until eventually an equilibrium state is achieved. Figure 6a shows that after 30 seconds from switching-on the strain gage board the equilibrium state was not completely achieved. Conventional monitoring systems do not show this bias because they are always operated in the equilibrium state since switching-on occurred long time before the first measurement started and the devices are never switched-off between two measurements.

The bias can be removed by adding a dummy gage in the Wheatstone bridge for achieving temperature compensation (2 gage system). Since the dummy gage is not bonded to the structure, it has to be designed to provide the same thermal characteristics of the primary gage for achieving good temperature compensation. This solution, however, is expensive since the temperature compensation must be adapted for each specific application and complicates the deployment process.

An alternative method is to remove the time-dependent bias by post-processing the raw data. The goal is to fit the bias with a suitable time-dependent function and to remove it by subtracting the fitted function from the raw data. The principle is depicted in Figure 6a. The bias is defined selecting the local minima of the raw data. This approach is justified since significant nonzero strains occurred only when the stringer is directly subjected to the axle loads. Hence, the minima of the unloading phases of the raw data are driven mainly by the non-stationary heating process. Independent measurements performed with conventional strain measurement devices confirm this approach (Kiviluoma, 2007). A satisfactory fitting of the minima was achieved with a simple exponential function that depends on the square root of time. The fitting was performed by a least square fit. The fitted function is shown in Figure 6a by the dotted blue curve. Figure 6b displays the strains after removal of the bias.

The effect of bias removal on strain cycle counting is displayed on Figure 6c. The most important differences concern the cycles with large strains. The differences are due to the first 10 seconds of the raw data because in this interval the bias gradient is high and tend to overestimate the amplitude of the cycles. Since the amplitude of cycles are defined by the difference of two values, the cycle counting histograms of raw and processed data do not differ significantly. If the accuracy requirements are not particularly tight, cycle counting using the biased raw data may already provide sufficiently good results for fatigue assessment.

Figure 6d demonstrates that the conservative sending policy with separated time slots did not prevent data loss. Only 16 of 30 seconds of the data recorded at node 8 reached the base station. A possible reason for the data loss was the electro-magnetic shielding due to the lateral main girders. Data loss is avoided by retransmitting the lost packets. This policy, however, consumes

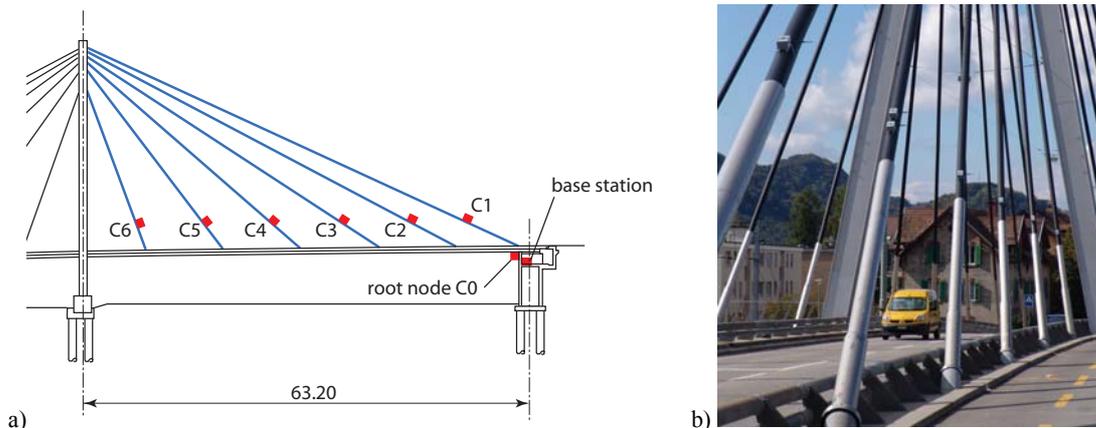


Figure 7: a) Set-up of the WSN monitoring system. b) Stay cables of the Stork Bridge with mounted sensor nodes.

more power, shortens the life-time and increases the latency period until the system is ready for a new record. Instead of sending the full record, the raw data can be processed in the node to extract the cycle counting. For the record displayed in Figure 6a, this data processing reduces the communicated data to 0.4% of the original data size and takes 1.2 seconds for completion. Such small data sizes can be retransmitted without significantly affecting the power consumption and the latency period for the next measurement.

2.2 Cable stays monitoring

The long term performance of the WSN monitoring system was studied with a stay cable tension monitoring deployment. Cable tension is usually estimated by correlating the measured natural frequencies, which are extracted from ambient vibration recordings, with natural frequencies predicted with a cable model (Casas 1994; Feltrin et al. 2006). Since the natural frequencies are the only information needed to estimate the cable tension, they are extracted by processing the raw data in the nodes and transmitted to the network sink, while the raw data can be discarded. From the initial amount of data, e.g. thousands of samples, the data to be communicated is reduced to several natural frequencies. This approach results in a drastic reduction of the power consumption.

2.2.1 Monitoring system

The WSN monitoring system has been deployed on a cable-stayed bridge (Stork Bridge) in Winterthur. The network consists of 6 nodes (C1 to C6), which are mounted on 6 stay cables, a root node (C0), which is located at the northern abutment under the bridge deck, and a relay node (C7), which is mounted beside the bridge and has a line of sight to all other nodes. Figure 7a illustrates the set-up and Figure 7b displays the sensor nodes mounted on stay cables of the Stork Bridge. The root node is attached to the base station, an industrial PC, powered from the mains power supply. The base station establishes a communication link with a computer in Empa's laboratory over the cell phone network using standard Internet communication protocols. The laboratory computer hosts the data base of the monitoring system and provides the remote configuration tools.

Vibrations are measured with a LIS2L06 MEMS capacitive accelerometer from ST Microelectronics. The accelerometer is integrated into an electronic circuit containing a low pass filter with a cut-off frequency of 20 Hz and a signal amplifier with an amplification factor

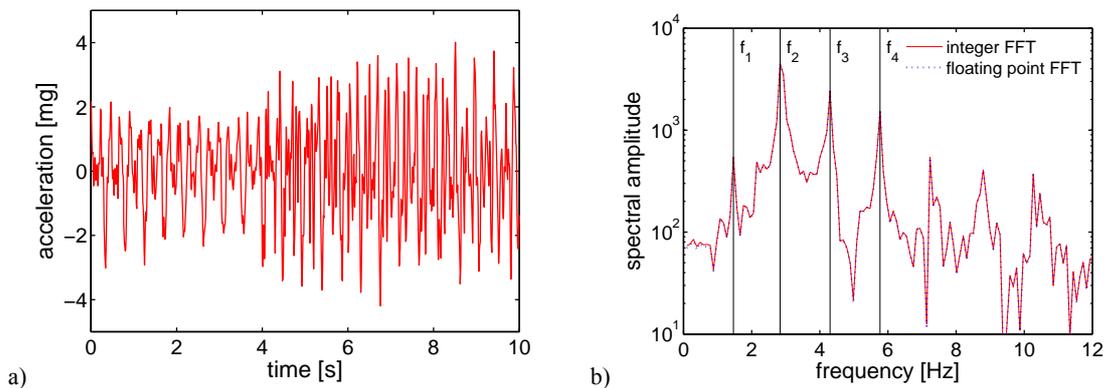


Figure 8: a) Ambient vibrations recorded on a stay cable. b) Fourier spectra computed with fixed and floating point operations.

of 50. The power consumption of the accelerometer and signal conditioning is approximately 5 mW. The ambient vibrations were recorded with a sample rate of 50 Hz for 20.5 seconds and contained 1024 samples. A typical time series of ambient vibrations is displayed in Figure 8a.

In addition to the ambient vibrations, temperature and humidity are periodically measured with the single chip Sensirion SHT11 sensor. The sensors are mounted inside the enclosure to investigate the time evolution of humidity as well as into an opening of the housing to measure the ambient temperature and humidity. Power supply voltage, routing and link quality information are also monitored. All data was acquired periodically with a time interval of 5 minutes.

2.2.2 Computation of natural frequencies

Typically, the estimation of natural frequencies via frequency spectrum is performed in three steps. First, the Fourier transform of the acquired vibration data is computed with a FFT algorithm. The second step computes the frequency spectrum with the real and imaginary of the Fourier transform. The third step identifies the peaks of the spectrum that correspond to the cable natural frequencies by using a peak detection algorithm.

Implementing this algorithm on a microcontroller with 10kB RAM is not a straightforward exercise since a time series of 1024 16 bit samples takes 2kB of memory. Since the RAM stores also the sensor node software, little memory is left for additional memory requirements of the data processing algorithm. Fortunately, the FFT algorithm, which represents the most demanding data processing step, can be implemented by using essentially only the memory of the time series. The intermediate and final results are stored in the same memory, overwriting the data from the previous stage.

Usually, the FFT algorithm is performed in floating point arithmetic. However, the microcontroller performs the floating point operations very slowly since they have to be emulated with integer operations. This overhead results in a greater code size and longer execution time. Less memory demanding, much faster and less power consuming is to perform the FFT algorithm with integer operations.

Evidently, an integer FFT implementation may introduce significant errors. A typical error source is overflow. By scaling the recorded time series with a suitable factor, however, the overflow can be avoided and the approximation error can be maintained within an acceptable range. Figure 8b) compares the frequency spectra of the time series displayed in Figure 8a) that are computed with a 16 bit integer approximation FFT and with the standard 32 bit floating

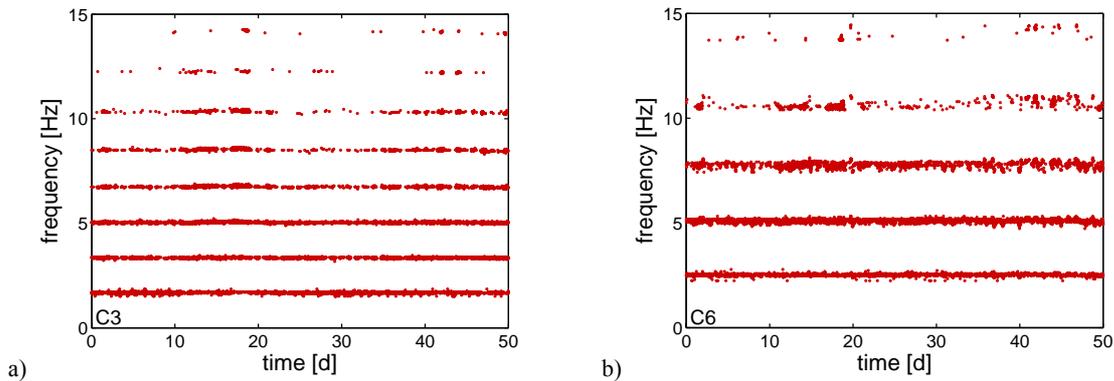


Figure 9: a) Natural frequencies of cable C3. b) Natural frequencies of cable C6.

point FFT. As can be observed, the result of the integer approximation FFT matches very well with the floating point FFT. The net execution time of the FFT was less than 0.55 seconds. The FFT computation with floating point operations took approximately 10 times longer and namely 5.8 seconds.

Computing the frequency spectrum (ℓ_2 norm) took 0.19 seconds and 8 natural frequencies additional 0.21 seconds. Hence, the final result, the natural frequencies, was computed in less than 1 second. Compared to the recording time of 20.5 seconds the data processing is very fast.

Figure 9 shows the time evolutions of the captured natural frequencies of cable C3 and C6. 6 natural frequencies could be regularly monitored for the cable C3 and 3 natural frequencies for the cable C6. Natural frequencies with a magnitude greater than 11 Hz are more difficult to detect since the associated vibration modes are scarcely excited. The standard deviation of the natural frequency estimations of cable C3 is less than 0.05 Hz and therefore approximately of the same magnitude of the frequency resolution of the spectrum, which is 0.05 Hz. A more pronounced scattering is displayed by natural frequency estimations of cable C6 that range between 0.05 and 0.13 Hz for the first three natural frequencies. This effect is due to the shortness of the cable C6 that limits the magnitude of the ambient vibrations. The 1h average of the natural frequencies (12 estimations) of the cables C3 differs from reference measurements by less than 2%. For cable C6 the difference between the average value and the reference measurement is smaller than 4%.

2.2.3 Stability and reliability

From the very first day on, the most challenging issue was to achieve system stability. The software that integrated data acquisition, data processing, time synchronization, low duty cycle, task scheduling etc. into one packet turned out to be very sensitive to many tiny details regulating the various activities. Sensor nodes disappeared and spontaneously reappeared after some time without any evident reason. When operating the network with TinyOS 1.x/Boomerang the routing tree was changing continuously. In this highly dynamic situation quite often nodes failed to choose a parent node correctly and lost their link to the network.

The system stability and reliability improved significantly after porting the software to TinyOS 2.x. This progress is shown in Figure 10a that displays the hourly delivery ratio during 60 days of two test periods. The delivery ratio is defined as the percentage of natural frequencies received by the remote control unit with respect to the theoretical maximum.

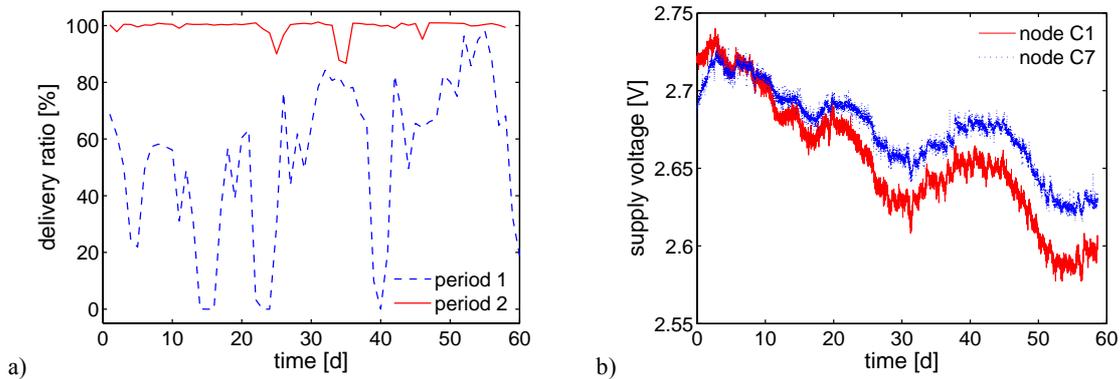


Figure 10: a) Delivery ratio of two test periods. b) Supply voltage evolution of nodes C1 and C7.

The period 1, April-May 2007, displayed rarely a delivery ratio higher than 80% and several total break downs of the whole network. Total failures of sensor nodes to deliver data occurred very often. In the second period, January-February 2010, this pattern was not observed. The delivery ratio was 100% for most of the time and was only briefly interrupted with partial break downs that did not go below 80%. The mean delivery ratio of the two test periods was 53% and 99.5%, respectively.

2.2.4 Energy Consumption

The average power consumption of the WSN node is determined by the time needed for sensing and data processing, the amount of data to be transmitted, the duty cycle period, the power management of the hardware and the network topology, since it determines the number of hops for reaching the data sink. Sensors and signal conditioning boards were switched-off after completion of the data acquisition. Furthermore, since the nodes were not equipped with a voltage regulator, the radio was switched-on during data acquisition in order to avoid signal corruption by the duty cycle. Figure 10b shows the battery voltage drop of the sensors node C1 and C7. Due to ambient temperature variations, the observed voltage curves are not decreasing monotonically but are oscillating significantly. The voltage drop of C1 was approximately 0.15 V in 60 days. Since a sensor node can be operated correctly provided the supply voltage is higher than 2.4 V, a lifetime estimation based on the observed volt-age drop predicts a battery lifetime of approximately 240 days. The relay node C7, which was not equipped with sensors, had half the voltage drop of sensor node C1 and thus twice the life-time of node C1. This figure indicates that roughly half of the energy of a node is consumed by data acquisition, data processing and the extra radio-on time during data acquisition.

3 CONCLUSIONS

This 3 year of field test experience demonstrates that medium and long term data intensive monitoring of civil structures with wireless networks is feasible and that the produced information complies with the quality requirements in civil engineering. In the beginning, operating the wireless network reliably over a period of months turned out to be non-trivial task. The problems relied basically on balancing the requirements of a data intensive application with the requirements of minimizing power consumption for achieving a sufficiently long battery lifetime. Data reduction, a powerful method to save power and therefore a key factor for achieving the targeted lifetime, destabilized seriously the WSN by disturbing the basic operations like time communication and routing.

Nevertheless, with many hard- or software improvements, the problems could be substantially solved. The availability of the most recent test phase was better than 99%. Considering that the data loss of the most recent test phase was due to malfunctioning of the link between the root node and the base station, the availability figure is very encouraging. Despite the severe hard- and software limitations, relatively complex in-node data processing could be performed and the accuracy of the generated information was better than 5%. This figure is very close to comply with the quality requirements in civil engineering, which usually do not require high precision information. The simultaneous application of several power saving mechanisms, in particular in-node data processing, allows to easily achieve node lifetimes of several months. Further significant improvements are still possible, since the duty cycle of the radio was not optimized. The hardware limitations, however, imposes a tight specialization to the monitoring task. This implies a detailed analysis and specification of the monitoring goals.

Although not all aspects were investigated in depth, the results obtained by the field deployments demonstrate that there are no fundamental obstacles preventing the application of long term monitoring systems based on wireless sensor networks to civil structures. This technology is very close to be mature for practical application. In the near future, progress in low power hardware will increase the computational resources without increasing or even decreasing the current power consumption. This progress will allow to perform more complex monitoring tasks and in-node data processing with less compulsion to specialization.

4 ACKNOWLEDGMENT

The authors express their gratitude to the Swiss State Secretariat for Education and Research, the Swiss Federal Roads Office, the Gebert RUF Stiftung and the directory board of Empa for the financial support.

5 REFERENCES

- Bischoff, R., Meyer, J. and Feltrin, G. 2009. Wireless sensor network platforms. In *Encyclopedia of Structural Health Monitoring*. Chichester, UK: John Wiley & Sons Ltd.
- Casas, J. R. 1994. A combined method for measuring cable forces: The cable-stayed Alamillo bridge, Spain. *Structural Engineering International*, 4: 235-240.
- Culler, D., Estrin, D. and Srivastava, M. 2004. Overview of sensor networks. *IEEE Computer*, 37(8): 41-49.
- Culler, D. E. and Wei, H. 2004. Wireless sensor networks. *Communications of the ACM*, 47(6): 30-33.
- Enochsson, O., Elfgren, L., Kronborg, A. and Paulsson, B. 2008. Assessment and monitoring of an old railway steel truss bridge in northern Sweden. *Fourth International Conference on Bridge Maintenance, Safety and Management (IABMAS '08)*, Seoul, Korea, July 13-17 2008, 3617-3624.
- Feltrin, G., Meyer, J. and Bischoff, R. 2006. A wireless sensor network for force monitoring of cable stays. *IABMAS'06 - Third International Conference on Bridge Maintenance, Safety and Management*, Porto, July 16-19, 2006, on CD.
- Feltrin, G., Meyer, J., Bischoff, R. and Motavalli, M. 2009. Modular low-power wireless sensor network for structural health monitoring. *4th International Conference on Structural Health Monitoring on Intelligent Infrastructure (SHMII-4)* Zurich, Switzerland, 22-24 July 2009, on CD.
- Feltrin, G., Meyer, J., Bischoff, R. and Motavalli, M. 2010. Long-term monitoring of cable stays with a wireless sensor network. *Structure and Infrastructure Engineering*, 6(5): 535-548.

- Gangone, M. V., Whelan, M. J., Janoyan, K. D., Cross, K. and Jha, R. 2007. Performance monitoring of a bridge superstructure using a dense, wireless sensor network. *Structural Health Monitoring 2007: Quantification, Validation, and Implementation*, 493-500.
- Gay, D., Levis, P., von Behren, R., Welsh, M., Brewer, E. and Culler, D. 2003. The NesC language: A holistic approach to networked embedded systems. *Proceedings of the ACM SIGPLAN 2003 Conference on Programming Language Design and Implementation*, 1-11.
- Glaser, S. D. 2004. Some real-world applications of wireless sensor nodes. *Proceedings, SPIE Symposium on Smart Structures & Materials/NDE 2004*, San Diego, CA, March 14-18, 2004, on CD.
- Kim, S., Pakzad, S., Culler, D. E., Demmel, D., Fenves, D., Glaser, S. and Turon, M. 2006. Health monitoring of civil infrastructures using wireless sensor networks. *Technical Report No. UCB/EECS-2006-121*, University of California, Berkeley.
- Kiviluoma, R., Feltrin, G., Meyer, J., Aho, T., Kilpelä, A., Lyöri, V. and T., R. 2007. Sb-8.2 demonstration of bridge monitoring. Prepared by Sustainable Bridges - a project within EU FP6. Available from: www.sustainablebridges.net.
- Levis, P., Madden, S., Polastre, J., Szewczyk, R., Whitehouse, K., Woo, A., Gay, D., Hill, J., Welsh, M., Brewer, E. and Culler, D. 2005. Tinyos: An operating system for wireless sensor networks. In *Ambient intelligence*. New York: Springer-Verlag.
- Lynch, J. P., Yang, W., Loh, K. J., Jin Hak, Y. and Chung Bang, Y. 2006. Performance monitoring of the geumdang bridge using a dense network of high-resolution wireless sensors. *Smart Materials and Structures*, 15(6): 1561-75.
- Mechitov, K., Kim, W., Agha, G. and Nagayama, T. 2006. High-frequency distributed sensing for structure monitoring. *Transaction of the Society of Instrument and Control Engineers (SICE)*, E-S-1(1): 109-114.
- Pakzad, S. N., Fenves, G. L., Kim, S. and Culler, D. E. 2008. Design and implementation of scalable wireless sensor network for structural monitoring. *Journal of Infrastructure Systems*, 14(1): 89-101.
- Spencer, B. F. and Nagayama, T. 2006. Smart sensor technology: A new paradigm for structural health monitoring. *Proc. of Asia-Pacific Workshop on Structural Health Monitoring*, Yokohama, Japan, December, 4-6, 2006, on CD.
- Straser, E. G. and Kiremidjian, A. S. 1998. A modular, wireless damage monitoring system for structures. 128, John A. Blume Earthquake Engineering Center, Stanford University, Stanford, CA.