

GPS-Equipped Wireless Sensor Network Node for High-Accuracy Positioning Applications

Bernhard Buchli, Felix Sutton, and Jan Beutel

Computer Engineering and Networks Laboratory, ETH Zurich, Zurich, Switzerland
{bbuchli, fsutton, jbeutel}@ethz.ch

Abstract. This work presents the design, implementation, and end-to-end system concept and integration of a wireless data acquisition system for high-accuracy positioning applications. A wireless network of GPS-equipped sensor nodes, built from low-cost off-the-shelf components, autonomously acquires L1 GPS data for Differential GPS (DGPS) processing of raw satellite information. The differential processing on the backend infrastructure achieves relative position and motion of individual nodes within the network with sub-centimeter accuracy. Leveraging on global GPS time synchronization, network-wide synchronized measurement scheduling, and duty-cycling coupled with power optimized operation and robustness against harsh environmental conditions make the introduced sensor node well suited for monitoring or surveying applications in remote areas. Unattended operation, high spatial and temporal coverage and low cost distinguish this approach from traditional, very costly and time consuming approaches. The prototype data acquisition system based on a low-power mote equipped with a commercially available GPS module has been successfully implemented and validated in a testbed setting.

Keywords: Wireless Sensor Networks, GPS, Positioning, Geodesy, Monitoring.

1 Introduction

Recent research [11, 20] indicates that the occurrence of destructive terrain displacements in alpine regions has been increasing, threatening critical socio-economic infrastructure. There are indications that permafrost related processes may induce such terrain displacements [13]. Knowledge of the processes affecting terrain stability are crucial for identifying potentially hazardous terrain to be monitored. Therefore, permafrost, related processes and their impacts on terrain stability continues to be an area of active geoscientific research.

To enable modeling of the complex relationships between geomorphological properties, freeze/thaw cycles and ground stability, information about spatial-temporal variations of terrain and the surrounding micro-climate is needed. Traditionally, this geological data is acquired by means of manual campaigns or expensive surveying techniques, such as laser scanning [9], interferometric synthetic aperture radar [28], or unmanned aerial vehicle surveillance [18]. In addition to high cost, these campaigns usually cover only very restricted areas and/or short periods of time. The lack of spatial and temporal resolution severely limits the analysis of large scale terrain dynamics, their origin and relation to other geomorphological processes.

Since the public availability of the NAVSTAR Global Navigation Satellite System (GNSS, commonly called GPS), and especially the removal of selective availability, technological and algorithmic advances have enabled the use of off-the-shelf, single-frequency L1 GPS receivers for terrestrial [22, 14, 26] and aerial positioning [12] applications. The inherent error in GPS positioning caused by atmospheric effects, multipath interference, and clock uncertainty can be mitigated with expensive, high accuracy dual-frequency L2 GPS receivers, which exploit two distinct GPS satellite signals. Alternatively, as in the system described in this paper, the same can be approached with differential post-processing of raw L1 satellite data. It has been shown [10, 21], that appropriately long measurement periods coupled with Differential-GPS (DGPS) techniques can achieve an accuracy well below the nominal GPS positioning accuracy of a few meters [25]. In [4] and [17], the authors further validate that L1 positioning accuracy in the sub-centimeter range is possible.

The X-SENSE project [6], a research collaboration between engineering and geoscience, has the goal of developing sensing technology suitable for highly accurate observation of ground-based terrain movement [27]. To provide the high quality spatial-temporal data that allows modeling of the various terrain movements, recent advances in wireless sensor networks and the aforementioned GPS technology are leveraged. To this end, a Wireless Sensor Network (WSN) equipped with low-cost, off-the-shelf GPS receivers that can be used in combination with differential GPS processing has been implemented and validated in a testbed setting. Monitoring of terrain displacement in high-alpine regions [27] is used as example, but many other scientific or commercial application scenarios are possible, e.g. early warning systems, machine control, precision agriculture, and infrastructure integrity monitoring.

Although equipping WSN nodes with GPS receivers is not new, the limited accuracy of conventional L1 GPS position solutions calculated at the receiver restricts its use to applications where errors on the order of multiple meters are acceptable. While the positioning accuracy can be improved by means of other techniques, such as trilateration using radio signal strength indicators [14], the achievable accuracy still does not satisfy the sub-centimeter accuracy requirements of applications that target monitoring of slow movements, such as supra glacial creep [19]. In these application scenarios, low-cost L1 GPS receivers in combination with receiver built-in DGPS processing techniques, or expensive L2 GPS receiver equipment are employed. Low-cost receivers have also been used in WSNs for establishing a global time base since accurate timing knowledge is essential for low-power wireless connectivity without expending excessive energy [26]. Further, WSNs are often deployed in remote environments where localization of the nodes and identification of the events they are to detect may be challenging. To avoid costly hardware providing such functionality, GPS receivers have been utilized [23]. Finally, an attempt similar to the monitoring approach discussed in this work has been made by the GGPhi project [2], funded by the Galileo Joint Undertaking. Unfortunately, it is not clear to what extent the ambitious goals formulated have been achieved.

This work presents a description of the design, implementation, end-to-end system integration, and evaluation of a low-cost wireless data acquisition system for differential GPS post-processing based, real-time positioning applications. However, details on the double-differencing range and carrier phase processing algorithm for which this system

was designed are beyond the scope of this work. Nevertheless, because of the implications on the data acquisition system design, processing is briefly addressed in *Section 2*.

The remainder of this paper is structured as follows. *Section 2* introduces the GPS based end-to-end X-SENSE monitoring system concept, while *Section 3* provides a brief background on DGPS and its implications on the system design. *Section 4* details the physical implementation of the data acquisition prototype, and *Section 5* discusses the prototype's performance evaluated in a testbed. Finally, *Section 6* concludes this work.

2 End-to-End System Concept

A high level overview of the X-SENSE system concept is depicted in *Figure 1*. For flexibility and ease of maintenance, the system is partitioned into three sub-systems: data acquisition, data handling and processing, and data application. Data exchange between individual sub-systems is achieved with low-power radio and Internet Protocol (IP) communication. These individual sub-systems are briefly explained in this section.




Fig. 1. System concept: The wireless sensor network nodes acquire environmental data at remote deployment sites. The data is collected at a basestation and transferred to the backend servers where it is stored and processed. Finally, the end-user may access raw and processed data through a web-interface.

Data Acquisition. As illustrated in *Figure 1*, the data acquisition system consists of a wireless sensor network with GPS equipped nodes that collect raw GPS satellite data for processing on the backend server. The data acquired on each node is sent via multi-hop communication to a basestation acting as a common sink for the network, and serving as an interface to the backend infrastructure. The data acquisition system implementation is explained in detail in *Section 4*, and its performance evaluation discussed in *Section 5*.

Data Handling and Processing. Data handling consists of storing, processing, and presenting data in a format useful for the respective target application. The tasks of this sub-system range from simple database transactions to complex computations of positions and velocities using the DGPS processing framework, which is briefly explained in *Section 3*. These functionalities are enabled by an enhanced version of the

GSN streaming middleware [1], which provides all-in-one functionality: data storage in relational databases, application of transfer functions to convert raw sensor readings, and publication of data through a web-interface. The plug-in architecture further caters for custom processing modules and functionality tailored to the application's needs.

Data Application. After data acquisition from remote deployment sites and application-specific processing, the data is made publicly available¹ for further integration by the end-users. For instance, the X-SENSE project utilizes position and velocity information fused with other sensor data, such as temperature and high-resolution imagery, to enable geoscientific modeling of high-alpine terrain displacements [27].

Communication. The nodes within the network use a low-power, low-bandwidth wireless network protocol stack for communication. The Dozer [8] network protocol is used for reliable multi-hop data transfer from each network node to a centralized basestation that serves as a data sink. The basestation is connected to the backend infrastructure through a wireless high-speed IP link. Once the data from the low-power network has arrived at the basestation, it is forwarded to the GSN backend using the reliable TCP/IP delivery service and data acknowledgements at the application layer. Server initiated periodic reconnection attempts to the basestation upon link loss allow autonomous and reliable network operation. To cope with intermittent radio communication links, non-volatile storage on each node further ensures operation without data loss.

3 Differential GPS Processing

With the GPS processing playing an integral part in computing sub-centimeter accurate position and velocity from the data acquired, this section discusses the process on a high level to motivate the system design choices. However, details on differential GPS are beyond the scope of this work; the interested reader finds an introduction in [25].

The observation equations describing the relation between raw GPS data collected at the receiver and satellite positions are given in *Equations (1) and (2)* [21], where P_1 and L_1 stand for L1 code and carrier phase measurements, respectively. Superscript i , and subscript k refer to individual satellites and GPS receiver antennas, respectively, while ρ_k^i indicates the pseudo-range between antenna k and satellite i .

$$P_{1k} = \rho_k^i + c \cdot \delta_k - c \cdot \delta^i + I_k^i + \Delta\rho_k^i \quad (1)$$

$$L_{1k} = \rho_k^i + c \cdot \delta_k - c \cdot \delta^i - I_k^i + \Delta\rho_k^i + \lambda_1 \cdot n_{1k}^i \quad (2)$$

There are a number of errors inherent to GPS positioning that can be eliminated with the double-differencing approach. For instance, the error associated with the receiver clock (δ_k) can be approximately eliminated by differencing the phases measured from satellite n with that from satellite $n + 1$ ($\forall n \in \{1, \dots, i - 1\}$) at a given measurement interval. Similarly, the second differencing operation computes the difference of satellite clocks observed by receiver r and $r + 1$ ($\forall r \in \{1, \dots, k - 1\}$). This approximately eliminates satellite clock errors (δ^i). The latter differencing operation is performed between the data collected by a reference station with known coordinates and the individual node

¹ <http://data.permasense.ch>

whose position is to be determined. This requires that the samples are collected in a synchronous manner to avoid non-overlapping samples and data gaps, which negatively affect the solution accuracy [4]. Hence, network-wide synchronized sampling of all the GPS receivers to be included in a position solution is necessary, a requirement that will be addressed in more detail in *Section 4*.

For baselines below one kilometer, errors due to ionospheric (I_k^i) and tropospheric ($\Delta\rho_k^i$) refraction can also be eliminated as it can be assumed that receivers in sufficient vicinity exhibit the same error due to atmospheric effects. Thus, (1) and (2) can be simplified (*Equations (3) and (4)*), where λ_1 is the L1 carrier wave length, and n_{1k}^i represents the ambiguity integer between satellite i and receiver antenna k .

$$P_{1k} = \rho_k^i \quad (3)$$

$$L_{1k} = \rho_k^i + \lambda_1 \cdot n_{1k}^i \quad (4)$$

The GPS data processing is performed with the Bernese [3] software package, utilizing a custom double-differencing algorithm [17] with a least squares approach. To improve data quality [4], the raw data collected is pre-processed (i.e. outlier removal, cycle slip detection, etc.), and static and kinematic solutions are computed using reference data from the International GNSS Service (IGS) repository. A kinematic solution refers to one solution per measurement interval based on latest available data. Kinematic solutions are particularly relevant for real-time applications such as early warning systems. A static solution, on the other hand, incorporates all samples collected on a particular day to compute a single solution for that day. Static solutions yield more accurate results, making this approach particularly well suited for capturing long-term trends, which is needed for scientific modeling and infrastructure integrity monitoring.

Contrary to other GPS applications where typically only few GPS data are acquired between frequent sleep periods of the receiver [14], the algorithm employed necessitates continuous sampling of raw satellite data for extended periods of time. A sensitivity analysis of duty-cycled L1 GPS receivers has shown that a measurement block duration of three hours is necessary and sufficient for the algorithm [17] to achieve an accuracy within five millimeters. Measurement blocks longer than that do not exhibit significant improvements. Interestingly, the sampling frequency, if kept above roughly 10 mHz, has a negligible effect on the solution accuracy, but only increases temporal resolution. However, with the custom double-differencing algorithm, increasing the sampling frequency does not substitute for shortening the measurement block duration in order to save energy. Rather, a higher sampling frequency increases the overall energy budget since more data must be transmitted via the radio.

In summary, the processing algorithm imposes very specific requirements on the data acquisition system. Most notably, the need for network-wide synchronized sampling and long measurement blocks require special attention. However, rather than synchronizing the network nodes using specialized synchronization protocols, we take advantage of the accurate time information provided by the GPS receivers, as explained in *Section 4*. The high energy consumption resulting from the long measurement blocks is mitigated by duty-cycling the power-hungry GPS receiver, and enabling the receiver's power save feature. The energy saving approach is explained in *Section 5*.

4 Data Acquisition System Implementation

Architecture. The architecture of the GPS-equipped wireless data acquisition node (henceforth referred to as “node”) is illustrated in *Figure 2*. The node comprises of an ultra-low power wireless sensor module (Shockfish SA TinyNode184), and a low-cost, off-the-shelf single-frequency L1 GPS receiver module (u-blox LEA-6T).

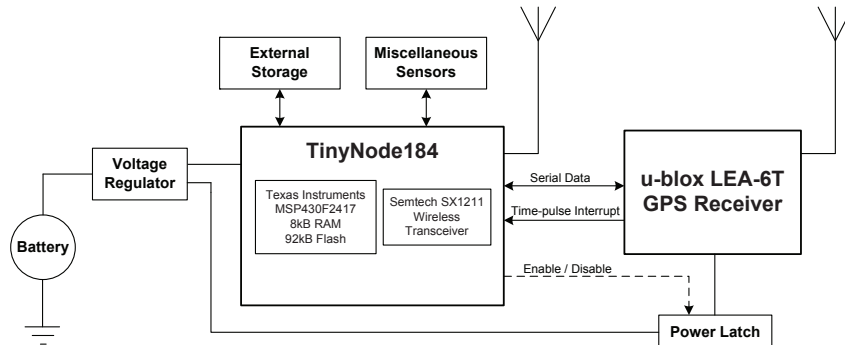


Fig. 2. Data acquisition system architecture. The power-hungry GPS module can be turned on/off by the TinyNode184 through a power latch.

The TinyNode184 consists of a Texas Instruments MSP430 low-power micro-controller and a 868 MHz Semtech SX1211 ultra-low power wireless transceiver. The MSP430 micro-controller has 92kB program flash, 8kB RAM, and a number of external interfaces including GPIO, UART, SPI and I²C. All application software on the node is implemented in TinyOS 2.1 [16].

The node’s micro-controller executes the application logic and communicates with the wireless sensor network using the on-board wireless transceiver. It also interfaces to an external storage device (SD-card) for reliable and redundant storage of measurement data and to miscellaneous application-specific sensors and/or actuators. The interface between the micro-controller and GPS module consists of two control lines and a communication and configuration interface. The usage of the time-pulse interrupt and serial communication is briefly explained in the following subsections. The second control signal enables or disables the power to the GPS module through a power latch.

Prototype. The prototype implementation is depicted in *Figure 3*. All system components are mounted into a water-proof die-cast aluminum enclosure for resistance against the harsh environmental conditions expected in high-alpine environments. The system is designed to be battery operated. Therefore, the lifetime of a node depends mainly on the duration of measurement blocks and sampling frequency (i.e. total number of samples collected, see also *Section 5*). If the lifetime required cannot be satisfied by a feasible battery size, it can be supplemented with an energy harvesting system to enable perpetual operation. Continuous satellite signal reception and extended wireless communication range, even under adverse environmental conditions, is achieved by mounting the system on a pole, thus elevating the antennas. The resulting ambiguity between

rolling, tilting, and actual displacement of the pole base (i.e. the observed agent) can be resolved with inertial sensors interfaced to the micro-controller.



Fig. 3. Prototype implementation housed in waterproof enclosure (60x100x160mm) with power connector, and external radio transceiver and GPS (Trimble Bullet III) antennas

GPS Module. The GPS module is a commercially available single-frequency L1 GPS receiver specialized for precision timing applications. In addition to standard GPS position fix messages, this module supplies raw (i.e. unmodified) satellite measurement data for the set of visible satellites. The module supports a UART interface over which a proprietary UBX binary protocol [24] is used to transfer commands and data between the micro-controller and the module. It further provides a precision UTC (Coordinated Universal Time) synchronized monostable time-pulse with a configurable pulse width and frequency.

The GPS module accommodates two modes of measurement retrieval, namely polled and periodic. The polled mode requires the micro-controller to explicitly request the measurement using a request packet having a zero-length payload, whereas the module in periodic mode supplies the data at a defined interval in coordination with the time-pulse. When the module is powered on, it attempts to search for, and synchronize with all available satellites. Once synchronization is achieved and the ephemeris data has been successfully downloaded and decoded, the module is able to provide polled or periodic measurements. In the case where there are no available satellites, a zero-length-payload packet is sent in polled mode, and no data is returned in periodic mode. The usage of the retrieval modes and the time-pulse is explained in the following subsection.

Data Capture. *Section 2* introduced the system requirements imposed by the processing algorithm. This section now elaborates on the requirement of network-wide synchronized data acquisition. However, rather than depending on a specialized network synchronization protocol, we take advantage of the accurate time information provided by the GPS modules and a measurement schedule broadcasted to all nodes. Apart from reducing the application complexity, and possibly allowing other optimizations, accurate time information on every node allows synchronous duty-cycled operation.

This allows any number of spatially distributed nodes to operate nearly optimally synchronized even during disconnected radio conditions, which is essential for reliable operation.

In order to start and stop the measurement blocks at the correct, and network-wide synchronized time, knowledge of absolute time is needed within the micro-controller. The time accuracy need only be in the order of seconds, since the GPS module may require several seconds to acquire synchronization with the satellites [24]. This is particularly true if the GPS module was powered off for more than four hours between observation periods, causing the ephemeris data to be out of date. However, the quartz oscillator of the micro-controller's internal real-time clock (RTC) is subject to drift, which limits its usefulness for scheduling purposes over very long periods. Instead of relying only on the RTC, the highly accurate time information maintained by the GPS module is leveraged. The micro-controller periodically polls the GPS module to acquire UTC timestamps for scheduling its measurement blocks. The following first explains the scheduling sequence after which the GPS sampling procedure is briefly discussed.

Each time the node is powered-on, or during a power reset cycle, the micro-controller briefly enables the GPS module using the power latch, and requests the current UTC time using the polled data retrieval mode (*Figure 4(a-i)*). Once the GPS module has acquired satellite synchronization, the current UTC time is returned (*Figure 4(a-ii)*). Since a measurement schedule in units of hours is assumed, the year, month, and day values are ignored. The micro-controller then calculates the time until the start of the next measurement block, configures an internal one-shot timer, and disables the GPS module. When the timer event fires, indicating the start of a new measurement block (*Figure 4(a-iii)*), the micro-controller re-enables the GPS module and configures it to commence periodic sampling of the satellite data.

With the nodes' intended deployment in the field spanning long periods of time (multiple months to years), small timing errors at the start of each measurement block can eventually cause the node to sample out of the specified measurement schedule. In order to prevent this measurement drift, the micro-controller reconfigures the one-shot timer at the beginning of each measurement block to indicate the conclusion of the current block (*Figure 4(a-iv)*). Upon firing, the UTC time is requested once again before switching off the GPS module and reprogramming the timer, which will trigger the start of the next block. This process repeats at infinitum and ensures continuous, time-accurate and network-wide synchronized periodic measurements according to a given schedule with minimal overhead.

As previously mentioned, the GPS module provides a monostable time-pulse that is interfaced to the node's micro-controller as an external interrupt. The GPS module is configured to indicate imminent availability of new data by triggering the interrupt. Once the module has acquired satellite synchronization, the time-pulse is always synchronized to within a few tens of nanoseconds of UTC time for *all* nodes in the network irrespective of the physical location, thereby enabling network-wide time synchronized GPS sampling for the entire network. The detailed measurement process using the periodic data retrieval mode is illustrated in *Figure 4(b)*, and explained next.

The GPS module is configured to trigger the time-pulse shortly before periodic GPS data is available over the UART interface, as illustrated in *Figure 4(b-v)*. When the

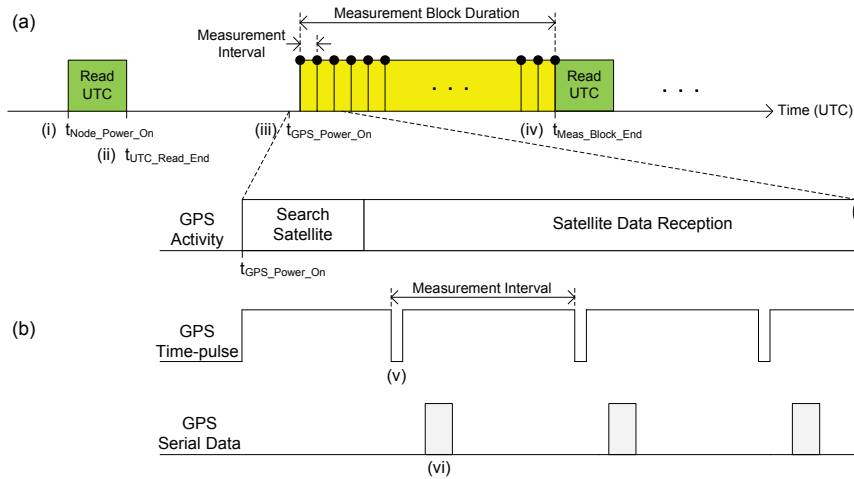


Fig. 4. (a) The GPS measurement block is controlled by the sensor node. (b) GPS module time-pulse and serial (UART) interface communication signals are fed back to the node.

interrupt is received at the micro-controller, the serial interface is enabled and the UBX parser is initialized, ready for the reception of GPS measurement data (*Figure 4(b-vi)*). After the data is received over the serial interface and validated with a CRC checksum, the UART is disabled until the next interrupt announces a new set of raw satellite data. Finally, the data is restructured into appropriate packets and transmitted using the Dozer network protocol as explained in the next subsection.

The interrupt driven approach employed not only guarantees autonomous network-wide synchronous sampling but also optimizes the micro-controller's usage. Rather than periodically spin-waiting for the data, the micro-controller can tend to other tasks, such as sampling of other application specific sensors. Further, the power saving modes of the micro-controller can be fully exploited, thereby reducing its energy consumption.

To provide a fully automated positioning solution for a wide range of application scenarios, a user configurable measurement schedule is implemented. Depending on the application's requirements, the number and duration of daily measurement blocks can be defined. All blocks are implicitly aligned to UTC time, characterized by the measurement block duration and sampling interval specified by the user. For the use-case assumed in this paper, i.e. monitoring and tracking daily variations of relatively slow terrain movement in high-alpine environments [6, 27], these schedules can be simple. In this case, a single measurement block duration of three hours per day is sufficient for sub-centimeter accurate relative positioning. For applications requiring a higher temporal resolution, e.g. capturing sub-daily dynamics, more elaborate measurement schedules based on such multi-hour blocks can be designed.

Network Communication. The GPS module outputs raw satellite information (see *Figure 5*) for the variable number of visible satellites at each measurement interval. The measurement data associated with each satellite is restructured from the binary UBX

format into a linear byte-wise data structure suitable for radio transmission. All unused or irrelevant bit-fields are omitted to reduce costly radio activity.

All satellite data structures are prefixed with the GPS time associated with the sample interval. Although the replication of the GPS time prefix for each satellite packet adds additional overhead, the data post-processing and validation is greatly simplified using this embedded time sequencing information. The satellite measurement data structure has a fixed size of 30 bytes, and is illustrated together with the Dozer and radio packet structures in *Figure 5*. After restructuring the data into the Dozer packets, they are injected into the Dozer protocol stack for transmission over the wireless multi-hop network, and written to non-volatile storage to ensure operation without data loss [5].

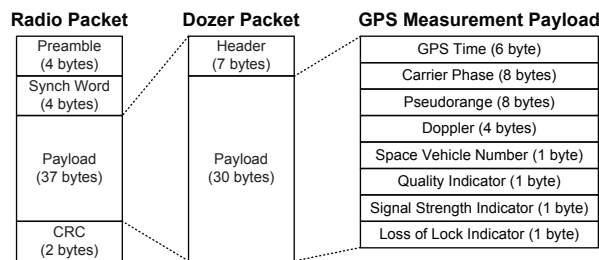


Fig. 5. Packet encapsulation of GPS measurement payload

The Dozer protocol stack also supports receiving command-beacons originating from the backend server. These special beacon packets contain a variety of commands for all or only a subset of nodes within the network. For example, such a command can include an update of the measurement block duration and measurement interval. This feature allows the measurement schedule of each node to be reconfigured remotely, providing great flexibility to a variety of application scenarios. Furthermore, each node is configured to send periodic health status packets. This health information enables remote observability of individual nodes (i.e. core voltage, current drain, temperature, humidity, etc.) and the overall network (i.e. topology, packet count, buffer size, etc.)

5 Evaluation

This section summarizes performance results of the data acquisition system evaluated with a small network deployment situated on the roof of the Electrical Engineering building at the ETH Zurich campus (47.3773909° latitude, 8.5529154° longitude). Although this location exhibits optimal conditions, i.e. excellent view of the sky, and mostly unobstructed radio communication, this scenario is not necessarily to be expected at the intended deployment site. Nevertheless, the performance evaluation presented in this section verifies that the data acquisition process operates according to requirements.

Experimental Setup. The prototype GPS-equipped wireless node is deployed in a 1, 2 and 3-hop line-topology network during a three week period for each topology.

For the data analysis, a measurement block duration of 24 hours per day is configured. Considering the slow movement expected in the use-case and diminishing improvement with higher sampling frequencies (see *Section 3*), a measurement interval of 2 minutes is deemed practical. With respect to the amount of data generated, this configuration represents a worst-case scenario. For the power analysis, two three-hour blocks per day – a realistic configuration for monitoring daily variations of slow terrain movement – are chosen to illustrate the benefits of duty-cycling.

Data Yield and Quality. The solution accuracy achieved by the processing algorithm is directly related to the amount, and quality of consecutive data samples collected. Hence, the data yield and quality are important performance metrics for the data acquisition system. *Table 1* summarizes a collection of statistics gathered from the experiment. Depending on the number of satellites visible at the time of measuring, the UBX packet generated by the GPS module contains a variable amount of satellite data (up to 296 Bytes) [24]. However, the Dozer packet size is only big enough to accommodate one satellite, thus, each UBX packet is restructured into one radio packet.

Data yield of the data acquisition system is defined as the ratio between the number of Dozer packets collected and the expected number of packets during the measurement blocks. To verify that all data collected are transmitted and actually reach the base station, we use the Dozer packet header sequence number (see *Figure 5, Section 4*) in combination with the GPS time associated with each Dozer packet. As shown in *Table 1*, the data acquisition system does not lose a single packet during the evaluation period, and therefore achieves perfect *data yield* for all three network topologies.

Useful yield is defined as the ratio between Dozer packets, excluding those that do not pass a quality check, and the total number of collected Dozer packets. The quality check classifies a measurement interval as useful if and only if at least four satellites are visible, and the *quality indicator* of at least four satellites is greater than or equal to 5. Four satellites is the absolute minimum required for resolving position in three dimensions and performing clock correction [25]. However, the double-differencing algorithm achieves increasingly better results with every additional satellite that can be included in the differential processing. The *quality indicator* value is assigned by the GPS module and a value greater, or equal to 5 indicates that pseudorange, Doppler, and carrier phase values are valid for this particular measurement [24].

Table 1. Summary of data sets collected from the 1, 2, and 3-hop testbed deployment scenarios

Network Topology	1-hop	2-hop	3-hop
Number of meas. intervals	15,120	15,120	15,120
Number of Dozer packets	126,662	133,641	133,085
Data yield	100%	100%	100%
Useful yield	99.99%	99.99%	99.99%

Table 2 tabulates the distribution of visible satellites and their corresponding *quality indicator* observed over the three weeks measurement period for each network topology. While some satellites do not pass the quality check (*quality indicator* less than 5),

they do not adversely affect performance. This is because at least four satellites during the same measurement interval exhibit a *quality indicator* of 5 or greater. *Table 2* also shows that the minimum number of satellites for every measurement interval is greater than four, but only very few measurements contain more than 11 satellites.

Table 2. Distribution of visible satellites and the corresponding distribution of *quality indicator* for the 1, 2, and 3-hop testbed deployment scenarios

Network Topology	Number of Satellites								Quality Indicator			
	5	6	7	8	9	10	11	≥ 12	4	5	6	7
1-hop	0.50	5.01	17.05	31.05	29.46	13.93	2.66	0.28	0.13	1.05	0.01	98.81
2-hop	0.01	1.9	11.38	28.22	27.99	22.35	6.30	1.84	0.13	1.31	0.0	98.56
3-hop	0.02	2.04	11.08	29.61	28.59	20.87	6.16	1.62	0.14	1.04	0.0	98.82

With 30 bytes per satellite, a significant amount of data can be generated on each network node. Empirical evidence from a test deployment at the targeted deployment site [6] with a powerful evaluation platform [7] suggest that an average of seven satellites per measurement interval can be expected under realistic conditions. Samples containing more than nine satellites were found in only 1% of the data collected at that particular site. This finding defines the upper bound of 10 satellites currently supported by the data acquisition software. While this can be easily adjusted if necessary, limiting the number of packets reduces the memory requirement for the node's forwarding buffers, and thus permits larger network deployments.

Table 1 shows that the data acquisition system achieves perfect *data yield* for all three topology configurations. Over the three weeks observation period only a single measurement interval failed the quality check described above, reducing the *useful yield* to 99.99%. The one measurement interval not passing the quality check is likely to be attributed to environmental factors, such as poor weather conditions or an unfavorable satellite constellation, both of which can affect the data quality. Even though the high quality recorded is favorable, it should be noted that the nodes in the testbed have almost semi-spherical visibility of the sky, which is likely not to be the case at an intended deployment site. Nevertheless, these results verify that the system reliably samples the GPS module and successfully communicates the acquired data to the basestation. At the basestation, reliable TCP/IP communication guarantees forwarding of all data to the backend storage server.

End-to-end Latency. While end-to-end latency is not critical for scientific modeling applications, it is an important aspect of other possible application scenarios, e.g. early warning system (EWS). The earlier a destructive event can be detected, the quicker alarms can be generated and threatened areas evacuated. With an EWS implementation in mind, this section briefly discusses the end-to-end network latency.


Using the data sets summarized in *Table 1*, the end-to-end latency is computed based on the GPS time associated with the measurement (i.e. generation time) and the NTP-synchronized time at which the packet is received at the basestation. *Table 3* shows the statistics for the end-to-end latency for all three test deployment topologies. The median latency for the three deployments is found to be 35, 50, and 69 seconds respectively.

Table 3. Statistics for end-to-end latency collected over three weeks from the testbed deployment

Network Topology	1-hop	2-hop	3-hop
Median (in sec.)	35	50	69
Mean (in sec.)	35.54	50.61	71.27
Mode (in sec.)	32	52	82
Variance (in sec.)	82.09	343.13	490.01

The lower bound on the end-to-end latency is based on the parameterized Dozer network protocol configuration, specifically the transmission interval, which is set to 30 seconds in the testbed deployment. Depending on the application's need, the latency can be reduced by adjusting the network protocol parameters accordingly. It is impossible to specify a tight upper bound on the end-to-end latency of the system, as the delay of individual packets is affected by a number of reasons [15]. For instance, a node or the basestation may go off-line due to power loss, poor wireless link quality, node reset, etc., causing an indefinite delay in the data acquisition. Experience with operating three independent networks running the Dozer protocol configured identically to this study suggests that the recorded latencies are representative of real-world deployments.

Power Analysis. The power dissipation of the node is an important performance metric as it defines the longevity of the device in the field. The unique combination of an interrupt driven design, synchronized measurement duty-cycling, and activation of the receiver's Power Optimized Tracking (POT) mode optimizes the overall system power dissipation.

**Fig. 6.** GPS module current drain profile during normal and power optimized tracking modes

Figures 6 and 7 illustrate the current drain of the two main contributors to the total power dissipation of the node: the GPS module and the TinyNode184 radio transceiver respectively. *Figure 6* demonstrates the activation of power optimized tracking of the GPS module. At approximately 5 seconds into the trace the satellite visibility increased such that the POT mode was enabled, drastically reducing the GPS current drain to 22mA. With peaks of up to 57 mA, the current drain during POT mode iterates between 49mA and 22mA with a ratio of 1:5, yielding an average power dissipation of 137mW.

The power profile in *Figure 7* illustrates the current drain due to the TinyNode184 platform only, which consists of a series of periodic radio transmission and reception bursts corresponding to the periodic beaconing schedule and data transmission of




Fig. 7. TinyNode184 current drain: (i) data transmission, (ii) beacon transmission, (iii) beacon reception, (iv) application processing

the underlying Dozer network protocol. Compared to the GPS module and the radio transceiver, the current drain of the application processing is considered negligible as a result of the low-power operation of the MSP430 micro-controller.

The four main contributing components to the TinyNode184 power dissipation (see *Figure 7*) are: (i) GPS data transmission, (ii) beacon transmission to child nodes, (iii) beacon reception from the parent, and (iv) Dozer protocol, TinyOS, and data processing overhead. Periodic beacon transmission and reception is an essential part of the network protocol and ensures network connectivity.

Table 4 tabulates the contribution to the power budget by the GPS module and radio transceiver respectively. The rows labeled *Measurement Active* refer to the period when the system is actively sampling the GPS module, while *Measurement Inactive* shows the contributions during times when the GPS module is turned off. *Table 4* clearly shows that, due to the GPS module's high power dissipation, its energy requirements is orders of magnitude higher than that of the TinyNode184, illustrating the applicability and resulting benefit of duty-cycling this component. Continuous sampling may be of interest for certain applications, as it allows monitoring of sub-daily phenomena, however, the improvement in accuracy over a duty-cycled operation is negligible. Furthermore, although increasing the sampling frequency yields higher temporal resolution, the power dissipation increases due to higher data volume transmitted via the wireless transceiver.

Table 4. Summary of components and their contributions to total average current drain measured at 5V DC. Radio TX and RX refers to radio transmission and reception respectively.

	i	Component	Current in mA (I_i)	Contribution in sec/min (α_i)	Energy in mWh
Measurement Inactive	1	Radio TX	30	0.12	0.3
	2	Radio RX	8	0.1	0.07
Measurement Active	3	Radio TX	30	0.32	0.8
	4	Radio RX	8	0.2	0.13
	5	GPS (POT high)	49	12	49
	6	GPS (POT low)	22	48	88

To approximate the sensor node's lifetime on a single battery charge, *Equations (5) and (6)* are used. *Equation (5)* estimates the node's total average current drain for the cases when the GPS receiver is active (I_{GPS_Active}) and inactive ($I_{GPS_Inactive}$) (in mA) per hour, where I_i is the contributing current of component i (see *Table 4*), α_i is the average proportion of time component i is active, and K is a constant equal to 60 sec/min. *Equation (6)* then yields the lifetime L (in days), where I_{GPS_Active} and $I_{GPS_Inactive}$ are the total average current drain during a 24 hour period when the GPS is active and inactive respectively. N_{Blocks} is the number of measurement blocks per day of duration t_{Active} hours and C is the battery cell capacity in mAh with a discharge efficiency β .

$$I_{GPS_Inactive} = \sum_{i \in \{1,2\}} \left(\frac{\alpha_i}{K}\right) I_i, \quad I_{GPS_Active} = \sum_{i \in \{3,\dots,6\}} \left(\frac{\alpha_i}{K}\right) I_i \quad (5)$$

$$L = \frac{\beta \cdot C}{(24 - N_{Blocks} \cdot t_{Active}) \cdot I_{GPS_Inactive} + N_{Blocks} \cdot t_{Active} \cdot I_{GPS_Active}} \quad (6)$$

Using the parameters in *Table 4* and assuming two measurement blocks of three hours each per day, using a standard battery cell [5] with 14000 mAh capacity having a 75% discharge efficiency (i.e. $N_{Blocks} = 2$, $t_{Active} = 3$, $\beta = 0.75$, $C = 14000$), a node lifetime of 62 days on a single battery charge can be achieved.

6 Conclusions

We have introduced an end-to-end system concept for sub-centimeter accurate positioning applications. Our positioning solution provides high spatial-temporal resolution, cost efficiency, deployment scalability and real-time data acquisition, all of which are significant advantages over more traditional geodetic positioning techniques that rely on costly technologies and/or require extensive manual campaigns.

We have also presented the architecture, implementation and evaluation of a GPS-equipped wireless sensor node custom-built for data acquisition within our end-to-end positioning solution. We show that our sensor node is cost-effective, power optimized, and exhibits applicability for a wide-range of positioning applications.

Acknowledgments. The authors would like to thank the X-SENSE technical staff for their support. The work presented was supported by NCCR-MICS, a center supported by the Swiss National Science Foundation under grant number 5005-67322, and Nano-Tera.ch.

References

1. Aberer, K., Hauswirth, M., Salehi, A.: Zero-programming sensor network deployment. In: In Next Generation Service Platforms for Future Mobile Systems, SPMS (2007)
2. Aguado, L., et al.: A low-cost, low-power Galileo/GPS positioning system for monitoring landslides. In: Navitec (October 2006)

3. Astronomical Institute, University of Bern, Bernese GPS software version 5.0 (January 2007) (accessed August 01, 2011)
4. Beran, T., et al.: High-accuracy point positioning with low-cost GPS receivers: how good can it get? In: ION GNSS (2005)
5. Beutel, J., et al.: PermaDAQ: A scientific instrument for precision sensing and data recovery under extreme conditions. In: Proc. 8th ACM/IEEE Int'l Conf. on Information Processing in Sensor Networks (IPSN/SPOTS 2009), pp. 265–276. ACM (2009)
6. Beutel, J., Buchli, B., Ferrari, F., Keller, M., Zimmerling, M., Thiele, L.: X-SENSE: Sensing in extreme environments. In: Design, Automation Test in Europe Conference Exhibition (DATE), pp. 1–6 (March 2011)
7. Buchli, B., et al.: Demo abstract: Feature-rich platform for WSN design space exploration. In: 10th International Conference on Information Processing in Sensor Networks (IPSN), pp. 115–116 (April 2011)
8. Burri, N., von Rickenbach, P., Wattenhofer, R.: Dozer: Ultra-low power data gathering in sensor networks. In: 6th International Symposium on Information Processing in Sensor Networks, IPSN 2007, pp. 450–459. ACM (April 2010), doi:10.1109/IPSN.2007.4379705
9. Carter, W., Shrestha, R., Slatton, K.: Geodetic laser scanning. *Physics Today* 60(12), 41–47 (2004)
10. Choy, S.: An investigation into the accuracy of single frequency precise point positioning (PPP). PhD Thesis. School of Mathematical and Geospatial Sciences, RMIT University (2009)
11. Crozier, M.: Deciphering the effect of climate change on landslide activity: A review. *Geomorphology* 124(3–4), 260–267 (2010), doi:10.1016/j.geomorph.2010.04.009
12. D'Amico, S., Montenbruck, O.: Differential GPS: An enabling technology for formation flying satellites. In: Sandau, R., Roeser, H.-P., Valenzuela, A. (eds.) *Small Satellite Missions for Earth Observation*, pp. 457–465. Springer, Heidelberg (2010), doi:10.1007/978-3-642-03501-2_43
13. Hasler, A.: Thermal conditions and kinematics of steep bedrock permafrost. PhD Thesis. Department of Geography, University of Zurich (2011)
14. Jurdak, R., et al.: Adaptive GPS duty cycling and radio ranging for energy-efficient localization. In: Proc. of the 8th ACM Conference on Embedded Networked Sensor Systems (Sensys), pp. 57–70. ACM (November 2010), doi:10.1145/1869983.1869990
15. Keller, M., et al.: Comparative performance analysis of the permadozer protocol in diverse deployments. In: Proc. of the Sixth IEEE International Workshop on Practical Issues in Building Sensor Network Applications (SenseApp 2011), pp. 969–977. IEEE, Bonn (2011)
16. Levis, P., et al.: TinyOS: An operating system for sensor networks. In: *Ambient Intelligence*. Springer, Heidelberg (2004)
17. Limpach, P., Grimm, D.: Rock glacier monitoring with low-cost GPS receivers. In: Abstract Volume 7th Swiss Geoscience Meeting (November 2009)
18. Lucieer, A., Robinson, S., Turner, D.: Using an unmanned aerial vehicle (UAV) for ultra-high resolution mapping of Antarctic moss beds. In: Proc. of the 15th Australasian Remote Sensing & Photogrammetry Conference (September 2010)
19. Martinez, K., Ong, R., Hart, J.: Glacsweb: a sensor network for hostile environments. In: First Annual IEEE Communications Society Conference on Sensor and Ad Hoc Communications and Networks, IEEE SECON 2004, pp. 81–87 (October 2004), doi:10.1109/SAHCN.2004.1381905
20. Raveland, L., Deline, P.: Climate influence on rockfalls in high-alpine steep rockwalls: The north side of the Aiguilles de Chamonix (Mont Blanc massif) since the end of the 'Little Ice Age'. *The Holocene* 21(2), 357–365 (2010)
21. Schüler, T., Diessongo, H., Poku-Gyamfi, Y.: Precise ionosphere-free single-frequency GNSS positioning. *GPS Solutions* 15, 139–147 (2011), doi:10.1007/s10291-010-0177-5

22. Squarzoni, C., Delacourt, C., Allemand, P.: Differential single-frequency GPS monitoring of the la valette landslide (french alps). *Engineering Geology* 79(3-4), 215–229 (2005)
23. Stoleru, R., He, T., Stankovic, J.: Walking GPS: a practical solution for localization in manually deployed wireless sensor networks. In: 29th Annual IEEE International Conference on Local Computer Networks, pp. 480–489 (November 2004), doi:10.1109/LCN.2004.136
24. u-blox AG, u-blox 6 receiver description including protocol specification, http://www.u-blox.com/images/downloads/Product_Docs/u-blox6_ReceiverDescriptionProtocolSpec_%28GPS.G6-SW-10018%29.pdf (accessed August 01, 2011)
25. U.S. Coast Guard Navigation Center, NAVSTAR GPS user equipment introduction (September 1996) (accessed August 01, 2011)
26. Werner-Allen, G., et al.: Fidelity and yield in a volcano monitoring sensor network. In: Proc. of the 7th Symposium on Operating Systems Design and Implementation, OSDI 2006, pp. 381–396. USENIX Association, Berkeley (2006)
27. Wirz, V., et al.: Temporal characteristics of different cryosphere-related slope movements in high mountains. In: Proc. 2nd World Landslide Forum. Springer, Heidelberg (2011)
28. Wright, T., et al.: InSAR observations of low slip rates on the major faults of western Tibet. *Science* 305(5681), 236–239 (2004)