A Flexible In Situ Power Monitoring Unit for Environmental Sensor Networks

William C. Evans, Alexander Bahr and Alcherio Martinoli
Distributed Intelligent Systems and Algorithms Laboratory
École Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland
Email: *firstname.lastname*@epfl.ch

***Abstract*—Wireless radios are a great consumer of energy in sensor networks. Retrieving data from a remote deployment in an energy-efficient fashion is a difficult problem, and while solutions have been proposed in literature, real-world systems typically implement robust though inefficient methods. In an effort to bring efficient monitoring techniques to real-world environmental sensor networks, we seek to now quantify the performance brought by these algorithms in practical terms, i.e., by their resulting reduction in overall station energy consumption. To this end, we have developed a power monitoring extension board that integrates seamlessly with a commercial environmental sensor network platform. The board is capable of measuring all incoming and outgoing power for a station, and can disconnect subsystems, such as the solar panel or sensor bus, as necessary. The board is currently deployed on an outdoor network and is undergoing extensive testing.**

I. Introduction

Wireless radios are a great consumer of energy in sensor networks. Low power networking algorithms for WSNs have received much attention over the last several years, including algorithms for advanced MAC [1] and data transport [2]. Retrieving data from a remote deployment in an energy-efficient fashion is one such difficult problem with particularly strong implications for environmental sensor networks, and while solutions have been proposed in literature [3]–[5], real-world systems typically implement robust but inefficient methods. Simulations over datasets spanning a wide variety of environmental conditions have shown that even simple approaches, e.g., suppressing sensor measurements that have not significantly changed since their last transmission, may reduce the amount of reported data in scientific deployments by more than half [6]. In an effort to bring efficient monitoring techniques to real-world environmental sensor networks, we seek to now quantify the performance brought by these algorithms in practical terms, i.e., by their resulting reduction in overall station energy consumption.

Sensorscope is a commercial platform for environmental monitoring which has seen extensive use in scientific deployments [7], [8]. Individual stations are typically deployed with a number of environmental sensors attached in a daisy chained fashion from a central collection and logging board, henceforth referred to as the *datalogger*. The datalogger features a short-range radio for local communication and optionally a long range GSM transceiver for sending data out of the network. Sensor detection and network configuration is performed autonomously when a station is activated. A complete description of Sensorscope's system architecture can be found in [9], [10].

In order to evaluate the performance of data collection algorithms under real-world conditions, we have developed a power monitoring extension board that integrates seamlessly with Sensorscope's existing architecture. The board logs power consumption data to local storage as well as over the network (as with any other sensor in the Sensorscope system). It is also capable of providing real-time information to the station itself, which has applications beyond performance evaluation, such as power-aware networking algorithms or advanced charge controllers. We currently have ten stations equipped with power boards undergoing outdoor deployment.

II. Related Work

Power monitoring is a highly sought-after feature for WSNs. iCount [11] and SPOT [12] provide energy monitoring functionality. Aveksha [13] is a more complicated extension board developed for the TelosB platform, providing energy monitoring and software tracing. By making minor modifications to the TelosB mote, they are able to take advantage of the microcontroller's debug features to provide this functionality in a non-intrusive fashion.

Our approach is unique in that it is closely integrated with a system designed for use outside of laboratory settings, with the capability to separately monitor incoming and outgoing power for different subsystems.

Many approaches to energy efficient monitoring of WSNs have been proposed. These algorithms can be largely classified as operating on the basis of either spatial or temporal suppression. Both methods attempt to avoid sending redundant data to the sink, preferring instead to have the sink infer these values based on other received data. Temporally focused techniques operate local to a node, examining its history of measured values. In contrast, purely spatial approaches seek to suppress a nodes value if it can be inferred given the value of nearby nodes.

Prorok et al. study hierarchical network topologies based on spatial clustering [5]. In this approach, cluster heads may choose to prune their children if the part of the monitored field they represent is highly isotropic as defined by some statistically computed threshold. Arici and Altunbasak propose using a first-order model to determine the predictability of particular nodes [4]. They define some of the nodes in the

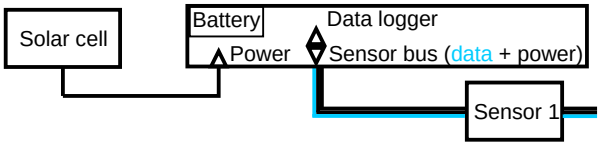

Fig. 2. Block diagram of Sensorscope setup without the power board

network as *macronodes* which attempt to fit a plane over their neighbors' positions and data, commanding easily predictable nodes to stop reporting measurements for some period of time.

Constraint Chaining (CONCH) is a spatio-temporal approach to suppression [3]. CONCH operates by detecting highly correlated neighbors and monitoring a suitably defined edge constraint between them instead of their individual values.

To our knowledge, none of the algorithms listed above have undergone implementation outside of a controlled laboratory environment. We plan to study the performance of these techniques on real systems using the power board described in this paper.

III. POWER MONITORING

Our goals in developing an in situ power monitoring unit for the Sensorscope station, henceforth referred to as the *power board*, were as follows:

- to not require any electrical or mechanical modifications to the Sensorscope system
- to not require any modifications to the code on the datalogger
- to obtain the voltage and current of all supply and load currents of the Sensorscope system (logger and attached sensors)
- to provide the ability to switch on and off any all supply and load currents
- to provide a dedicated processor for intelligent power control algorithms
- to integrate with the Sensorscope bus system
- to integrate mechanically within the existing enclosure

The resulting board as shown in in Figure 1 fulfills all the above requirements. It can be integrated with an existing Sensorscope system by simply opening the cover, removing the two ribbon-cable connectors which provide connectivity to the plugs on the enclosure and reattaching them to the power board. The power board is then connected to the datalogger with two additional ribbon cables. The following section describes the different sub sections of the power board in detail.

A. Interface

In the original Sensorscope system the datalogger provides two connectors (Figure 2). The first provides connection to a solar cell which is used to charge the batteries on the datalogger which provide power to the entire system. It also serves as a switch as two shortened pins within the connector indicate when the connector is plugged in and switch on the entire system. The second is the data and power connection for

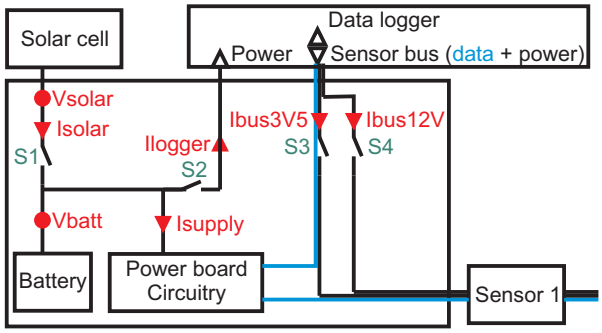

Fig. 3. Block diagram of Sensorscope setup with the power board

the first sensor in the entire chain of sensors which are attached to the datalogger. Subsequent sensors are daisy-chained to the first sensor.

The power board provides its functionality by inserting itself between the datalogger and the two outside connectors. Mechanically this is achieved by providing the same connectors (*Tyco* μmatch 6 pin for power and 8 pin for data bus) on the power board as on the datalogger (top right corner in Figure 1). In order to be directly mounted on top of the datalogger within the same enclosure, the power board has the exact same dimensions. A hole (top left corner in Figure 1) allows for the LEDs on the datalogger to be visible through the power board and a window in the enclosure. When the power board is attached, the original four rechargeable batteries in the datalogger can be removed as they are bypassed and the power storage is provided by the twelve rechargeable batteries on the power board.

As the solar cell is now directly attached to the power board, we can measure the solar cell voltage (V_{solar}) and current (I_{solar}). The power connection to the datalogger allows us to measure the current flowing into the logger (I_{logger}). In addition it monitors the battery voltage (V_{batt}) and the current consumption of the power board itself (I_{supply}).

On the sensor bus side, sitting between the datalogger and the first sensor, serves two purposes. First, the power board is attached to the sensor bus like any of the subsequently attached sensors. It thus appears to the logger board as any other sensor and its measurements can be queried by the datalogger. This implementation was crucial in order to insure interoperability between the two boards without having to modify the dataloggers code. Second, in addition to a data connection, the bus also provides two power rails to the individual sensors 12 V and 3.5 V. We can separately measure the current on both rails (I_{bus3V5} and I_{bus12V}).

B. Microcontroller

The microcontroller on the power board is a an MSP430F1611 by *Texas Instruments* running *TinyOS*. It is exactly the same microcontroller as the one on the TinyNode which controls the datalogger. This facilitates code development and interoperability. Furthermore the MSP430 series has a several features which make it particularly well suited

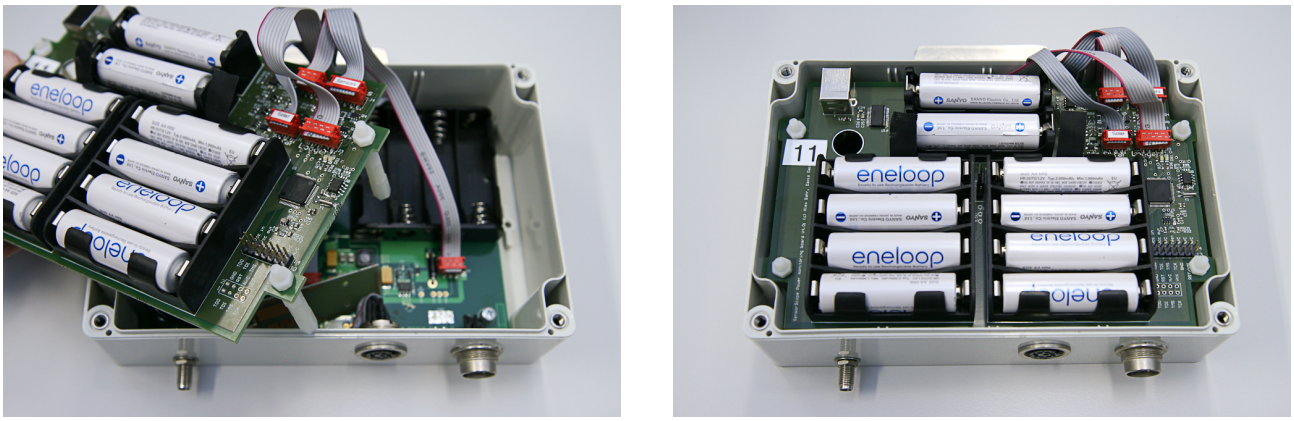

Fig. 1. The power board (top) fits inside the datalogger over the original board.

for WSN applications: it is specifically designed for low power operation with a very efficient core (1200 MIPS/W) and a series of low power modes (down to 1.1 μA) from which it can wake up with minimal overhead. In addition the MSP430F1611 has a very high-resolution (12 bit) A/D-converter.

For data which needs to be permanently stored, such as board-specific configurations, we added a 1 Mbit serial EEPROM (*Microchip* 25AA1024). A unique digital serial number is provided by a *Maxim* DS2401 which can be directly accessed by the microcontroller and provides an identification of the power board to the datalogger and the network. An additional dedicated circuitry (*ST Microelectronics* STM6779) ensures that the MSP430 is reset when its supply voltage drops below 3.08 V. As the precision and stability of the reference voltage has a strong influence on the ADC's performance and we require very high precision on all measurements, the MSP430 is provided with an external 2.5 V precision reference (*Analog Devices* REF192ES).

C. Voltage sensing capabilities

The solar cell voltage V_{solar} and the battery voltage V_{batt} are both measured through a voltage divider between GND and the respective power rail with the center point attached to the MSP430's ADC. In order to avoid a constant leakage current the voltage dividers are separated from the power rail with matched N/P-MOSFET (*ROHM* QS6M3) acting as a high-side switch. The switches can be individually activated through I/O pins of the MSP430 if a voltage is to be measured.

D. Current sensing capabilities

The circuit for sensing all five currents measured by the power board (I_{solar},I_{logger},I_{supply},I_{bus3V5} and I_{bus12V}) is shown in Figure 4. It consists of a shunt R_{sense} in combination with high-side current monitor (*Zetex* ZXCT1010). The sensed current I_{sense} can be computed from the output voltage V_{out} of the current monitor using Eq. 1.

$$I_{sense} = \frac{100 * V_{out}}{R_{sense} * R_{out}} \tag{1}$$

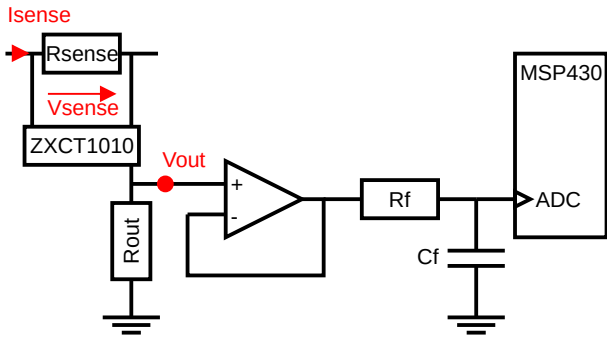

Fig. 4. Current sensing circuit

R_{sense}	1 Ω
R_{out}	1 kΩ
ΔI (Eq.2)	61 μA
I_{max} (Eq. 3)	250 mA
I_{min} (Eq. 4)	$\geq$ 200 μA

TABLE I
COMPONENT VALUES AND CHARACTERISTICS OF THE CURRENT SENSING CIRCUITRY

The values for R_{sense} and R_{out} can be modified such that the losses $P_{shunt} = I^2_{sense} * R_{sense}$ through the sensing shunt are minimized while the resolution is maximized. The resolution of the current measurement ΔI depends on the largest measurable current I_{max} and for our 12 bit A/D converter is

$$\Delta I = I_{max}/2^{12}. \tag{2}$$

Using Eq. 1 we can thus compute the value of the product $R_{sense} * R_{out}$ for a maximum expected current I_{max} and an A/D-converter reference voltage of 2.5 V as

$$R_{sense} * R_{out} = \frac{100 * 2.5 \text{ V}}{I_{max}} \tag{3}$$

For a given $R_{sense} * R_{out}$ we want to minimize the shunt losses by selecting a value for R_{sense} which is as small as possible. As the ZXCT1010's output however is increasingly

non-linear for sensing voltages $I_{sense} * R_{sense} < 200\ \mu$V we also need to make sure that for the smallest currents I_{min} which we would still like to resolve R_{sense} is selected s.t.

$$I_{min} * R_{sense} \geq 200\ \mu\text{V} \tag{4}$$

For all five sensing circuits the same values for R_{sense} and R_{out} (Table I) have been chosen. After we were able to collect sufficient data we modified those values in order to maximize the resolution for each measured current.

To decouple the current sensing from the low-pass filter and A/D-converter an operational amplifier in voltage follower configuration (*Microchip* MCP60X series) is part of each current sensing circuit. The subsequent first order low-pass filter ($R_f = 82$ kΩ, $C_f = 2.2\ \mu$F) with a cutoff frequency $f_c \approx 1$ Hz is matched to the sampling frequency of our A/D-converter $f_s = 5$ Hz $\geq 2f_c$.

E. Switching capabilities

In addition to measuring voltages and currents, the power board is also capable of switching several power rails (see Figure 3 for position of switches). S_1 connects the solar cell to the the power board. If the batteries are fully charged or we want to measure the open-circuit voltage of the power panel, the power board can detach the solar cell from the system. The power supplying the datalogger and the sensors can switched with S_2. Before the batteries are fully depleted the power board can switch off the station by switching off S_2 to avoid deep-cycling of the rechargeable batteries. Another reason to switch off S_2 would be a fault on the datalogger which results in a very large I_{logger}. Similarly S_3 can switch off the attached sensors in case of a fault detection (very large I_{bus3V5}) or just the sensor's 12 V supply in case of a very large I_{bus12V}.

F. Debugging ports

The MSP430's JTAG-port is used for programming and debugging (lower right corner). A second debug port is available through serial USB connection implemented with an *FTDI* FT232R USB-to-serial converter which is directly attached to one of the MSP430's two USART ports.

IV. Results and Conclusion

We have currently deployed ten Sensorscope stations equipped with power boards around EPFL's campus. Preliminary results from this deployment can be seen in Figure 5. When the battery voltage falls below the charging threshold, the power board connects the solar panel and solar current increases dramatically as the batteries recharge. The upper threshold is reached in under an hour, and the solar panel is again disconnected.

We have also used the board to characterize the power consumption of various station activities, as seen in Figure 6. Note that the consumption measured by the power board closely matches that measured by current probe. The power board has limited temporal resolution compared to our oscilloscope, and does not capture the effect of the datalogger's switching voltage regulator, among other factors.

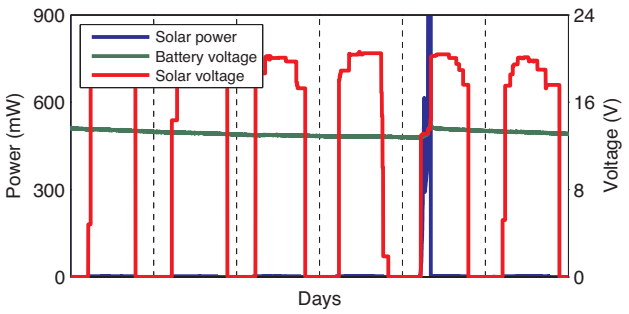

Fig. 5. Power board data from one station during normal deployment. Here we observe the behavior of our voltage-threshold charge controller. The charging period corresponds to the spike in solar power.

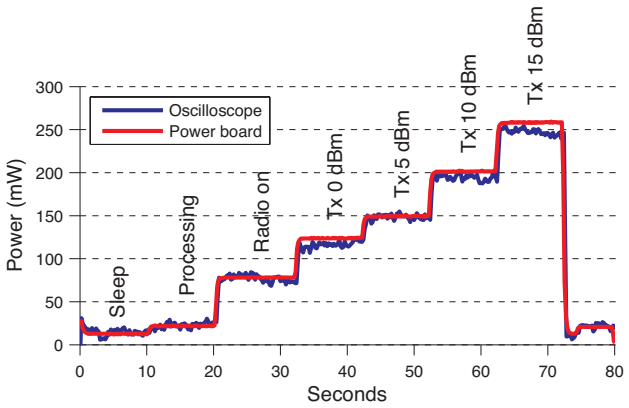

Fig. 6. Power consumed by the datalogger while sleeping, processing, radio on, and cycling through various transmission power levels. The power board has limited temporal resolution compared to the oscilloscope, and does not capture the effect of the datalogger's switching voltage regulator, among other factors.

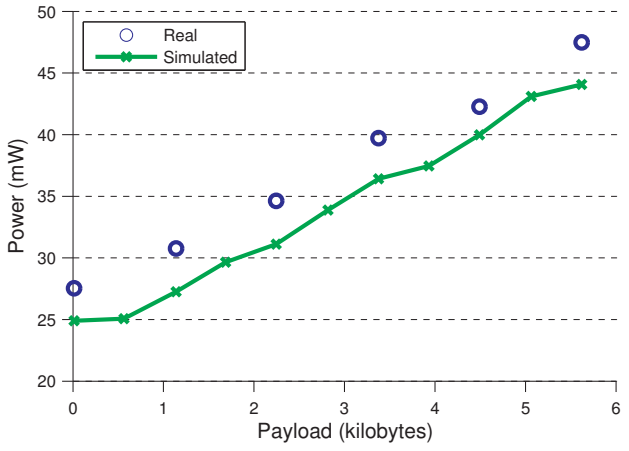

Fig. 7. Comparison of real-world versus simulated power consumption as a function of per-frame payload. Each data point represents the power consumed during one entire radio cycle (two minutes), averaged over a 14 minute period.

While the power board was designed to give quantitative results regarding algorithms deployed in real-world environments, such experiments are time consuming and may yield dramatically different results between runs. It is thus critical that we leverage power board data to estimate power consumption during repeated simulations over the same datasets. The above Sensorscope characterization allows us to add realistic power consumption estimates to TOSSIM [14] simulations. In Figure 7 we present a comparison of simulated power consumption versus data measured by the power board on a real station. The station under measure operated as a typical Sensorscope slave station (i.e., without GPRS), transmitting between zero and six kilobytes of sensor data per cycle. This shows that our simulation environment is able to accurately predict station power consumption.

We are now in the process of using the power board to evaluate the savings brought by various efficient WSN algorithms in both real-world deployments and simulation. We are also exploring its capability for real-time feedback for use in power aware algorithms and advanced charge controllers.

ACKNOWLEDGMENTS

This work was partially funded by "The Swiss Experiment" supported by the Competence Center Environment and Sustainability of the ETH Domain (CCES), and by the NCCR Transfer project "Tamperproof Monitoring Solution for Weather Risk Management" sponsored by the Swiss National Science Foundation and managed by the National Center of Competence in Research in Mobile Information and Communication Systems (NCCR-MICS).

REFERENCES

[1] R. Jurdak, P. Baldi, and C. Lopes, "Adaptive low power listening for wireless sensor networks," *IEEE Trans. Mobile Comput.*, vol. 6, no. 8, pp. 988–1004, Aug. 2007.

[2] T. Le, W. Hu, P. Corke, and S. Jha, "ERTP: Energy-efficient and reliable transport protocol for data streaming in wireless sensor networks," *Computer Communications (Elsevier)*, vol. 32, no. 7-10, pp. 1154–1171, May 2009.

[3] A. Silberstein, R. Braynard, and J. Yang, "Constraint Chaining: on energy-efficient continuous monitoring in sensor networks," in *Proceedings of the 2006 ACM SIGMOD International Conference on Management of Data*, Jun. 2006, pp. 157–168.

[4] T. Arici and Y. Altunbasak, "Adaptive sensing for environment monitoring using wireless sensor networks," in *IEEE Wireless Communications and Networking Conference*, Mar. 2004, pp. 2347–2352.

[5] A. Prorok, C. Cianci, and A. Martinoli, "Towards optimally efficient field estimation with threshold-based pruning in real robotic sensor networks," in *IEEE International Conference on Robotics and Automation*, May 2010, pp. 5453–5459.

[6] W. C. Evans, A. Bahr, and A. Martinoli, "Evaluating efficient data collection algorithms for environmental sensor networks," in *Proceendings of the 10th International Symposium on Distributed Autonomous Robotic Systems*, Nov. 2010, Springer Tracts in Advanced Robotics (2013). Vol. 83, pp. 77–89.

[7] D. F. Nadeau, W. Brutsaert, M. B. Parlange, E. Bou-Zeid, G. Barrenetxea, O. Couach, M.-O. Boldi, J. S. Selker, and M. Vetterli, "Estimation of urban sensible heat flux using a dense wireless network of observations," *Environ. Fluid Mech.*, vol. 9, pp. 635–653, 2009.

[8] S. Simoni, S. Padoan, D. Nadeau, M. Diebold, A. M. Porporato, G. Barrenetxea, F. Ingelrest, M. Vetterli, and M. Parlange, "Hydrologic response of an alpine watershed: Application of a meteorological wireless sensor network to understand streamflow generation," *Water Resources Research*, vol. 47, p. W10524, 2011.

[9] Sensorscope. Official website. [Online]. Available: http://www.sensorscope.ch/

[10] G. Barrenetxea, F. Ingelrest, G. Schaefer, and M. Vetterli, "The hitchhiker's guide to successful wireless sensor network deployments," in *Proceedings of the 6th ACM conference on Embedded Network Sensor Systems*, Nov. 2008, pp. 43–56.

[11] P. Dutta, M. Feldmeier, J. Paradiso, and D. Culler, "Energy metering for free: Augmenting switching regulators for real-time monitoring," in *Proceedings of the 7th International Conference on Information Processing in Sensor Networks (IPSN)*, Apr. 2008, pp. 283–294.

[12] X. Jiang, P. Dutta, D. Culler, and I. Stoica, "Micro power meter for energy monitoring of wireless sensor networks at scale," in *Proceedings of the 6th International Conference on Information Processing in Sensor Networks (IPSN)*, Apr. 2007, pp. 186–195.

[13] M. Tancreti, M. S. Hossain, S. Bagchi, and V. Raghunathan, "Aveksha: a hardware-software approach for non-intrusive tracing and profiling of wireless embedded systems," in *Proceedings of the 9th ACM Conference on Embedded Networked Sensor Systems*, Nov. 2011, pp. 288–301.

[14] P. Levis, N. Lee, M. Welsh, and D. Culler, "TOSSIM: accurate and scalable simulation of entire TinyOS applications," in *Proceedings of the 1st International Conference on Embedded Networked Sensor Systems*, Nov. 2003, pp. 126–137.