Towards a True Energetically Sustainable WSN: A Case Study with Prediction-Based Data Collection and a Wake-up Receiver

Alessandro Bogliolo*, Valerio Freschi*, Emanuele Lattanzi*, Amy L. Murphy† and Usman Raza‡
*DiSBeF, University of Urbino, Urbino, Italy
Email: {alessandro.bogliolo, valerio.freschi, emanuele.lattanzi}@uniurb.it
†Bruno Kessler Foundation, Trento, Italy
Email: murphy@fbk.eu
‡University of Trento and Bruno Kessler Foundation, Trento, Italy
Email: raza@fbk.eu

***Abstract*—One of the most challenging goals of many wireless sensor network (WSN) deployments is the reduction of energy consumption to extend system lifetime. This paper considers a novel combination of techniques that address energy savings at the hardware and application levels: wake-up receivers and node-level power management, plus prediction-based data collection. Individually, each technique can achieve significant energy savings, but in combination, the results are impressive. This paper presents a case study of these techniques as applied in a road tunnel for light monitoring. Preliminary results show the potential for two orders of magnitude reduction in power consumption. This savings of 380 times allows the creation of an energetically sustainable system by considering integration with a simple, photovoltaic energy harvester.**

I. Introduction

One of the main objectives in wireless sensor networking (WSN) research over the past decade has been the reduction of energy consumption to the point that, in combination with energy harvesting, the resulting system is energy-sustainable. Techniques to save energy range across the full stack from the hardware to the application, offering novel designs and tunable parameters to fit a variety of standard application settings. This work considers three techniques from across this spectrum, namely, a wake-up receiver and whole-node power management at the hardware side, plus prediction-based data collection at the application side. The first and third techniques aim to decrease the cost of data collection by focusing on the use of the radio, as it represents one of the most power hungry components on the node, while the second technique achieves further savings by exploiting the low-power modes of all node components.

Although novel medium access control (MAC) techniques such as low power listening (LPL) [1], [2] achieve significant reductions in consumption by putting the radio to sleep for extended periods, nodes still waste energy in two primary ways. First, nodes must periodically wake up and listen to the channel in case a node is attempting to transmit to it. If there is nothing to receive, this energy spent listening is wasted, leading to idle listening. On the other side, a sender with data to communicate must transmit until the receiver wakes up, often leading to long transmission times among unsynchronized nodes. Wake-up receivers are a novel hardware approach to eliminate these two main sources of overhead. Specifically, they provide an ultra low power receiver which is always on and listening to the channel, either the same channel used for communication or a dedicated, out-of-band channel. When a packet is to be transmitted, a preamble is generated by the transmitter to trigger the wake-up of the data radio on the receiving node. This eliminates the idle listening by only turning on the main radio module when there is a packet to be received. Further, it reduces the transmission time by ensuring that the receiver is ready to receive immediately after transmission of the preamble, thus avoiding the repeated transmissions typical of duty cycling protocols.

While exploiting a wake-up receiver leads to significant lifetime gains in many scenarios, applications that collect data at high frequency still incur significant costs to transmit the raw data. To ameliorate these costs, one technique is to develop a model that predicts the application data. In this approach, each node calculates a model for its data and communicates this model to the sink. The sink then uses these models to predict the data samples at each node. As long as the real samples closely match the model, no data is communicated, however as soon as the real data deviates significantly from the data estimated by the model, a node generates a new model and transmits it to the sink. Such approaches have the potential to eliminate 90 to 99% of the transmissions, depending on the type of data being sampled and the sophistication of the model.

In addition to controlling radio usage, the energy efficiency of WSN nodes strongly depends on effective, dynamic management of the power modes of all hardware components. Both the microcontroller unit (MCU) and the radio transceiver usually offer multiple modes that provide different trade-offs between time spent in an idle mode where power consumption is well below μW for many state-of-the-art components, and productive, awake time. In addition, most of the ancillary components of ultra-low-power motes, such as flash memories and real-time clocks, can be opportunistically turned on and off to save power. The actual ability to exploit the low-power modes depends on the node workload, which must be carefully considered by the dynamic power manager.

By combining these three techniques, we exploit the benefits of each to achieve significant savings. Intuitively, prediction-based data collection reduces the radio traffic, leading to long periods between transmissions. The wake-up receiver significantly reduces the power consumption during these idle periods. Finally, as the nodes are only infrequently involved in data forwarding, they remain completely idle between data samples, allowing the exploitation of the low power hardware modes.

Section II offers a brief overview of the related work in each of these techniques while Section III offers a description of how we apply them in this paper. The main contribution of this paper is a proof-of-concept study of the advantages when these techniques are combined in a concrete case study, namely light monitoring in a road tunnel [3]. In this application, light samples are taken from sensors spread along part of a road tunnel. A centralized control system collects these values to determine the light levels of the lamps along the tunnel so that the legal lighting constraints are met. Additional details of the scenario are provided in Section IV. The results, presented in Section V, show savings of up to two orders of magnitude, specifically 380 times, when the wake-up receiver, prediction-based data collection, and node power management are applied in this scenario. Notably, such results are sufficient to enable an energetically sustainable system formed by a simple, photovoltaic energy harvester. The paper concludes in Section VI with a discussion and future directions for this research.

II. Related Work

As the radio represents the most power hungry component on the node, many approaches reduce energy by saving communication costs. We consider two approaches that complement one another. First, prediction-based data collection is a software solution that lowers the data rate of the application without sacrificing data accuracy. Second, a wake-up receiver is a hardware solution that consumes less energy than a standard duty cycling protocol by avoiding the consumption due to idle listening that would occur between the infrequent transmissions generated by the first approach.

A. Prediction-based data collection

Prediction-based data collection maintains the original application sampling frequency, but reduces energy consumption by limiting the amount of data that must be transmitted [4]. This is accomplished by generating a model for the sensed data. This model is used at the sink to approximate the sampled data points. With each new sample, the node verifies that it falls within the allowable error tolerances. If so, no action is taken, but if not, a new model is generated and transmitted to the sink. If the model closely approximates the data trend, the network communication is significantly reduced, up to 99% in some cases [5]–[7]. Various types of models have been studied. Probabilistic models [8], [9] approximate data with a user-specified confidence, but special data characteristics must be encoded by domain experts. Alternate techniques employ linear regression [10], autoregressive models [11] and Kalman filters [12], but with sizeable memory and computational requirements, making them difficult to implement on resource-limited motes. A simpler, linear approach [5], detailed in Section III-A, was recently proposed by some of the authors of this paper, and is adopted for the case study here.

Additional classes of data reduction techniques include data compression, in-network data processing and data aggregation [4]. Although they can be used in conjunction with prediction-based data collection, the additional complexity outweighs the benefit for many WSN applications. Specifically, in the tunnel case study explored here, prediction-based data collection alone achieves sufficient savings, therefore we did not consider the addition of these other techniques.

B. Wake-up Receiver

Wake-up receivers are a viable solution to achieve low-power, asynchronous communication among nodes. The challenge of this approach is the availability of highly energy efficient triggering systems [13]. Several wake-up solutions have been recently developed [14]–[17] focusing on different parameters, namely: working power, sensitivity, distance range, latency, and operating frequency. In general, the adoption of asynchronous wake-up schemes has been shown to be useful in several application scenarios and is therefore deemed an interesting paradigm for the design of energy-efficient WSNs [18].

While using radio waves is a natural solution for the design of wake-up receivers for WSNs, out-of-band signaling solutions also offer viable alternatives. Specifically, an ultrasonic wake-up module for the VirtualSense platform has recently been developed by some of the authors of this paper [19]. Notably, it outperforms state-of-the-art radio wake-up receivers with a sub-μA quiescent current consumption. As a side benefit, ultrasonic wake-up modules can be also exploited to perform pairwise distance measurements [20], [21].

Asynchronous remote triggering is particularly effective for applications with ultra-low traffic as it avoids the idle listening incurred at routing nodes to periodically check for incoming packets. The main contribution of this work is to show that such ultra-low traffic conditions can be achieved in common applications by means of prediction-based data collection, thus enhancing the effectiveness of wake-up receivers.

III. System Architecture

The primary contributions of this paper are the novel combination of technologies from hardware to software to achieve an energetically sustainable system and the concrete evaluation of several configurations of these technologies. Figure 1 offers a very high level overview of the configurable components we consider, dividing them between software and hardware approaches, then further dividing the hardware among those belonging to the VirtualSense platform (the microprocessor control unit (MCU), the data transceiver, and the wake-up receiver) and the energy harvester. This section outlines the primary capabilities and options of each of these components as we apply them in this case study. For simplicity we do not show common components such as the operating system.

A. Derivative-Based Prediction (DBP)

When developing our data prediction technique, first described in [5], we sought to identify a simple model that effectively captures the data trends, thus reducing the amount of data generated at each node and requiring few resources on the constrained nodes. With DBP, we adopted a linear model based on m data samples, the first and last l points we refer

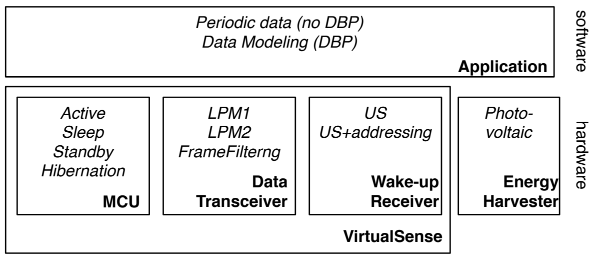

Fig. 1. High level system architecture showing the primary configurable components and the options we consider.

to as edge points. The linear model is calculated as the slope of the line connecting the average of the first l edge points with the average of the last l edge points. This computation resembles the calculation of the derivative, hence the name *Derivative-Based Prediction*.

On system initialization, the first m points are collected, after which the model is generated and sent to the sink. Subsequently, each sensor sample is checked against the value the model predicts. If the reading is within a given tolerance, no action is taken as the sink will also use the model to approximate the sensor sample. However, if the application tolerances are exceeded, a new model is generated and transmitted to the sink.

To offer a brief example, consider a light sensor. At sunrise the linear, DBP model will be an upward sloping line. At some point, however, the light levels will cease to increase and the upward sloping model will be replaced with a flatter model, corresponding to the daytime light levels. While this explanation is over-simplified, it nevertheless offers the intuition of DBP.

It is worth mentioning that we have applied DBP to several real data sets ranging from soil temperature to indoor temperature and light values [22]. In all cases, a reasonably tuned DBP produces data reduction rates above 89% and in most cases above 98%. These savings are sufficient to exploit the combination of technologies explored in this paper.

B. *VirtualSense*

VirtualSense is an open-hardware ultra low-power sensor node featuring a Java-compatible virtual runtime environment. Full details of the node are available in [23] while only the most relevant elements are detailed here. VirtualSense runs the Contiki operating system (OS) [24] and the Darjeeling Virtual Machine (VM) [25], suitably modified in order to make it possible for a Java programmer to fully exploit the low-power states of the underlying microcontroller unit, a Texas Instruments MSP430F5418a.

MCU power management. VirtualSense features four categories of power states: *active*, *standby*, *sleep*, and *hibernation*. In *standby* the CPU is not powered, but the clock system is running and the unit is able to wake itself by means of timer interrupts. In *sleep* both the CPU and the clock system are turned off and the unit is woken up only by an external interrupt. In *hibernation* even the memory system is switched off and there is no data retention requiring a complete reboot of the OS at wake-up, together with a restore of the VM heap.

Power consumption varies significantly across different states of VirtualSense. In the *active* mode, the average power consumption is approximately 13mW when processing and 66mW for transmitting, while the consumption reduces to 14.67μW in *standby*, 1.32μW in *sleep*, and 0.36μW in *hibernation*. We also note that the time to transition from one state to another is non-negligible. Specifically, the transition to *active* is 25ms from *standby* and *sleep*, and 500ms from *hibernation*. In case of memory-less applications, a lightweight hibernation mode is also available to trade-off data retention for wake-up time, which is consequently reduced to 27ms.

Data transceiver. In addition to the low-power modes of the MCU, VirtualSense supports low-power communication by exploiting the inactive modes of Texas Instruments CC2520 radio transceiver (LPM1, LPM2, and an extra low-current RX mode), the hardware frame filtering (FF) capabilities that prevent the reception of non-intended packets, and the ultrasonic wake-up receiver, described next, which allows routing nodes to be opportunistically triggered.

LPM2 is the lowest power consumption mode. In this state no data is retained and power consumption is about 13.5mW, in spite of the complete inactivity and of the need to reboot the embedded controller at wake-up. In LPM1 all data and configurations are retained and power consumption is about 3mW. During transmission, the power consumption of the entire transceiver ranges from 48.6mW (at -18dBm output power) to 100.8mW (at +5dBm output power), while in standard receive mode the power consumption is 69.9mA. Using the low-current RX mode reduces the power consumption in the receiving phase down to 55.5mW at the expense of a decrease in sensitivity from -98dBm to -50dBm.

The frame filtering function rejects non-intended frames and it can be configured to reject frames not matching the local address. When frame filtering is enabled it is possible to shut down immediately the receiver in case of a non-intended frame and to avoid processing the rest of the frame and waking up the MCU [26]. Finally, the data radio transceiver runs a standard ContikiMAC implementation.

Wake-up receiver. The ultrasonic wake-up module of VirtualSense is composed of both transmitter and receiver units whose main component is a piezolectric transducer working in the 2KHz band centered around a 40KHz frequency. Triggering of all nodes within the range of the transmitter is achieved through the detection of the ultrasonic carrier signal (US). Selective triggering of target nodes is optionally achieved by means of an *On-Off-Key* (OOK) modulation of the carrier to encode an 8 bit address (USa) [19]. On the receiver side the overall power consumption, has been experimentally measured at 1640nW in standby, while the power consumed by the transmitter is around 40nW in standby and 37mW during the transmission of the ultrasonic signal. These power consumption values enable the wake-up of motes within a 14-meter line-of-sight range. Such low power consumption values are in stark contrast to the values of the data radio.

C. Energy Harvester

The tunnel case study we use in this paper has an uninterrupted light supply, which can be readily converted by a photovoltaic cell into electrical energy to power the nodes. With direct exposure to sunlight during the day, nodes near the tunnel entrance can harvest a considerable amount of energy. Nodes deeper in the tunnel have only exposure to the artificial, fluorescent lamps, nevertheless with photovoltaic cells optimized for high efficiency at low illuminance, even these nodes have the ability to harvest energy. Therefore, for this study, we consider the Panasonic AM-1816 [27], a cell designed to self-sustain small electronics indoors, even under low-intensity fluorescent lights.

IV. Experimental Setup

Our unique combination of the technologies presented in the previous section exploits the lowest consumption modes of the VirtualSense platform for as much time as possible. This is accomplished by moving to a low-power MCU state between samples. Using DBP further decreases the costs in multiple ways. Consider that in any system, a node transmits its own data, forwards data from other nodes, and unnecessarily overhears packets destined to other nodes. Each of these events requires the node and the data radio to be switched to a high power consumption mode. DBP reduces the total traffic in the network, thus reducing the frequency of all these events, and consequently increasing the time the node can remain in the lowest power mode. Maximally, a node can remain idle between samples.

The results presented in the next section come from a series of simulations with DBP and VirtualSense, performed with actual data collected from a road tunnel and based on the real power consumption measurements of VirtualSense. This section offers details on these measurements and concludes with an estimation of energy harvested from photovoltaic cells.

A. Real-World Road Tunnel Data

The application case study we consider in this paper is based on a pilot deployment in a real road tunnel in Trento, Italy [3]. In this 260 m tunnel, a WSN of 40 nodes is deployed to periodically measure the light levels. Some of the nodes are exposed to sunlight, while nodes deep in the tunnel only detect the artificial light from the lighting system. In all cases, the light levels detected by the sensors every 30 s are transmitted over a multi-hop collection tree to a gateway at the entrance of the tunnel. The values are then used by a control system to gradually adjust the intensity of the lamps throughout the tunnel to meet the legislated light levels. The control system was designed in collaboration with lighting engineers to tolerate a limited amount of data loss and to accommodate some degree of error in the quality of the sensed values.

In this paper, we evaluate the system power consumption with multiple different hardware and software configurations. While the next section addresses the hardware configurations, here we consider the application layer, as it is affected by the case study itself. Specifically, we must consider the amount of data reported by each node with and without the DBP data prediction algorithm. For this, we used actual data traces

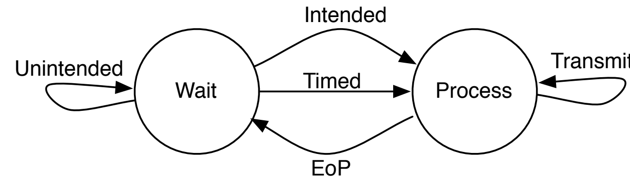

Fig. 3. Reference functional state diagram of VitualSense.

collected from the tunnel over a 47-day period in winter 2010, for a total of $5,414,400$ samples. Without DBP, we assume each sample is transmitted by each node immediately after being measured. With DBP, instead, we assume each node implements the DBP approach, transmitting only when the model changes. For the purposes of this study, we configured DBP to allow the predicted data values to deviate from the actual light values by at most 15 lux. Further, at most two consecutive samples can fall outside this error bound before a new model generation is triggered. These settings allow DBP to reduce the total sent traffic by 99% w.r.t. periodic reporting [5].

Another important element of our case study is the physical layout of the sensors, as this impacts the data collection tree. The 40 nodes are divided evenly between the two tunnel walls, and distributed unevenly along 130 m as shown in Figure 2. Nodes at the entrance of the tunnel, where sunlight influences the light levels, are more densely placed than nodes deeper in the tunnel. Given a 14 m range of the ultrasound wake-up receiver, we used a network simulator to form collection trees based on simulated link quality. Figure 2 shows a sample topology used in our simulations. As system performance is affected by the properties of the collection tree, e.g., depth, amount of data forwarded by a node, etc., we ran simulations with multiple topologies.

B. Power Simulation

Our power simulations arise from the functional state diagram shown in Figure 3, which captures the behavior of a WSN node able to exploit the idle periods during its normal workload to save power. The *Wait* state represents a family of inactive modes exploitable by the dynamic power manager. While waiting the node is sensitive to three types of events: overhearing of an unintended packet, reception of an intended packet, and a timed interrupt which wakes up the node for periodic tasks (i.e., sampling in our case study). The *Process* state represents activity such as sampling a physical quantity, evaluating the need to transmit a sample according to the prediction strategy, or routing an incoming packet toward the sink. When the processing ends (EoP), the node transitions back to the *Idle* state.

From this functional state diagram, we extract five, operating conditions that contribute to node power consumption: *i*) Waiting, *ii*) Hearing an unintended packet, *iii*) Receiving and routing a packet, *iv*) Waking up autonomously to sample then transmit, and *v*) Waking up autonomously to sample but not transmit thanks to DBP. The power consumption of the VirtualSense nodes is detailed in [28]. The average power consumption of each node is then computed by the simulator

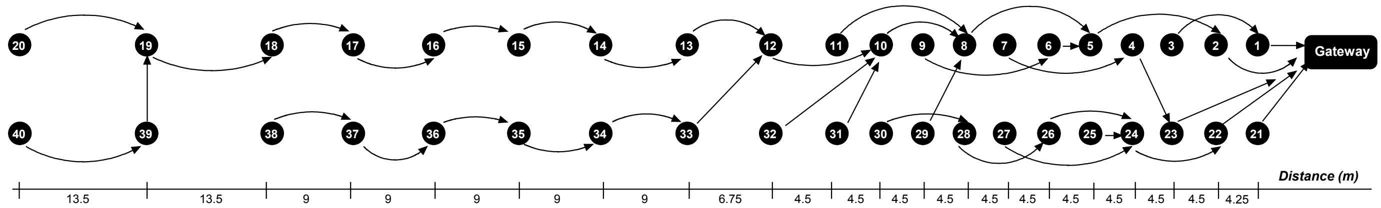

Fig. 2. Layout of nodes in the tunnel and a sample data collection tree with a depth of 15 hops.

as a weighted average, using as weights the actual occurrence rate of each operating condition, as derived from real-world traffic data.

The simulator provides power estimates for all the possible configurations of the dynamic power management options of VirtualSense. In particular, it allows us to: i) select the low-power mode of the MCU, ii) select the low-power modes of the data radio transceiver, iii) enable frame filtering, and iv) make use of the ultrasonic wake-up receiver with or without addressing capabilities. Further, the simulator provides a configuration flag to turn on and off DBP.

C. Estimation of Harvestable Energy

To estimate the harvestable energy for our nodes, we use light values collected from our 40-node WSN described earlier. We assume a fluorescent light source. This choice underestimates the amount of harvestable energy for nodes near the entrance that receive sunlight, but is accurate for nodes deep in the tunnel. From the total harvestable energy, the amount of energy successfully delivered to a node depends on harvester efficiency and battery leakage. For this initial case study, we assume a harvester efficiency of 79% [29] and zero battery leakage.

V. Experimental Results

This section reports the experimental results obtained on the case study according to the described setup. In contrast to previous studies [5], [6] that only consider the cost of radio communication as a proxy for overall node power consumption, thus under-estimating the total cost, the analysis here takes into account the contribution of all hardware components of VirtualSense plus the software overhead, carefully characterized for each operating mode.

A. Energy Savings

Our goal is to evaluate the system power consumption with multiple different configurations of the hardware and software, the options of which are outlined in Figure 1. Table I shows nine different configurations, each of which corresponds to a hardware state previously studied in detail in [28]. The first row shows a standard node configuration that uses the *standby* mode of MCU and the *LPM1* mode of the data transceiver. We consider this hardware configuration, in combination with a software layer that does not use the DBP prediction scheme on top of the ContikiMAC protocol with a 100ms wake-up interval, as the baseline for evaluating the power consumption reductions of all other configurations, shown in the columns labeled "Ratio". The power consumption values, shown in μW, are computed as averages over all 40 nodes and multiple collection tree topologies.

In all configurations, adding DBP results in lower power consumption, in line with our previous results [5]. Additionally, as expected, increasing the use of low-power modes also reduces consumption. Without the wake-up receiver, power consumption reductions are modest, up to eight times, even when considering a configuration that exploits the *sleep* mode of the MCU and avoids unnecessary overhearing with the data radio (FF).

Intuitively the addition of the wake-up receiver should have a significant impact by reducing idle listening in an application such as the tunnel where data is generated only two times per minute. Nevertheless, we see only a small improvement without DBP, specifically between configurations 5 and 6 in the no-DBP case we see only a 0.6 improvement in the power consumption. This is due to the fact that nodes must frequently transition out of the low power state to forward the traffic of other nodes. In other words, while idle listening does occur, the cost is insignificant with respect to the cost to transmit the large amount of traffic in the system.

Instead, combining the wake-up receiver with the extremely low data rate of DBP (configuration 6, DBP) results in a significant 380 times improvement, a result that is much larger than the improvements attainable by each technique in isolation. Specifically, dynamic power management alone achieves at most 2.6 times power consumption improvement and DBP alone obtains at most 7.9 times the baseline. This remarkable result of 380 times is a concrete demonstration of the benefits of eliminating idle listening with the very low transmission rates achieved with DBP. Additionally, as the node no longer needs to forward data on behalf of other nodes, it can spend more time between light samples in the low power mode of the MCU.

Notably, hibernation is not effective in our scenario. Without DBP, the time required to wake-up and restore the state of the node is not compatible with the high packet rate experienced by the nodes close to the sink, which both transmit their own packets and forward those of many other nodes. Fundamentally, a node does not have time to transition to and from the hibernation state. For this reason, we do not report consumption for configurations without DBP. Adding DBP increases the interval between transmissions, allowing all nodes to exploit the hibernation state. Nevertheless, they must still wake up to sample the light level every 30 seconds. The extra energy required to wake up each node from hibernation takes a toll on the power consumption improvements, with the system reaching only 3.5 times the baseline. From these results we infer that with DBP in conjunction with the wake-up receiver, communication is no longer the most power-hungry task.

Finally, we note that the lightweight hibernation mentioned

Configuration ID	Hardware Configuration			no-DBP		DBP	
	MCU	Data Transceiver	Wake-Up	[μW]	Ratio	[μW]	Ratio
1	Standby	LPM1	none	5891	1.0x	3460	1.7x
2	Standby	LPM2	none	3423	1.7x	758	7.8x
3	Standby	LPM2+FF	none	2905	2.0x	758	7.8x
4	Sleep	LPM2	none	3411	1.7x	745	7.9x
5	Sleep	LPM2+FF	none	2893	2.0x	745	7.9x
6	Sleep	LPM2	US	2257	**2.6x**	15.0	**380x**
7	Sleep	LPM2	USa	2750	2.1x	15.2	373x
8	Hibernation	LPM2	US	-	-	1665	3.5x
9	Hibernation	LPM2	USa	-	-	1617	3.6x

TABLE I. SYSTEM-WIDE ENERGY SAVINGS IN THE TUNNEL CASE STUDY. THE GRAY CELLS INDICATE THE BASELINE FOR CALCULATING THE POWER CONSUMPTION IMPROVEMENT RATIO OF ALL OTHER CONFIGURATIONS.

earlier, which does not restore the system memory, is not compatible with DBP. Specifically, with DBP, the application on the node uses the system memory to check if the sensed value fits the model and generate a new model based on the previous samples when the model does not fit. To avoid the full restoration of the system heap, one could consider storing and restoring only this application data to the flash memory, incurring less time and cost. To get an idea of how this kind of *memory-less* application with 30s sampling would perform, we considered a combination of *hibernation* mode, LPM2, and US wake-up. Without DBP, we achieved 2.5 times reduction in power consumption and with DBP we reached a remarkable 367 times reduction.

In summary, these numerical results bolster our argument that the individual techniques of prediction-based data collection and wake-up receivers, while capable of achieving improvements alone, are even more powerful when combined into a single system.

B. Energetic Sustainability

We next turn our attention to understand if DBP, wake-up receivers, and dynamic power management, when applied together, make VirtualSense motes energetically sustainable using a photovoltaic cell in our tunnel environment. To this end, the harvestable power at each node in the tunnel must be compared with its average power consumption when the most effective configuration is used, namely configuration 6. We consider both with and without DBP.

The shaded region in Figure 4 shows the amount of power that can be harvested, on average, at each node. Nodes close to the entrance (e.g., nodes 1 and 21) harvest more energy due to their exposure to natural sunlight at the tunnel entrance, while nodes deep inside the tunnel (e.g., nodes 20 and 40) receive light only from the fluorescent lamps. Values are represented in log scale, with approximately two orders of magnitude difference between the power available at the two ends of the WSNs.

The power needs of each node, operating in configuration 6 (Sleep, LPM2, US), are reported in the bar graph of Figure 4 both with DBP (dark bars) and without DBP (light bars). As expected, the power consumption at each node is much lower when DBP is applied. Also, DBP reduces the link between power consumption and node position in the tunnel. In fact, variation in consumption is mainly due to the data collection tree that produces higher traffic volumes at some nodes, e.g., nodes 2, 23 and 24 in Figure 2. Without DBP

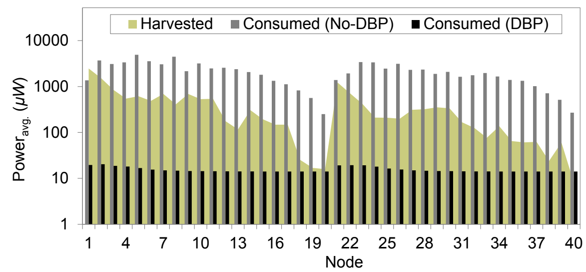

Fig. 4. Comparison of energy consumed and harvested, plotted on a logarithmic scale. Node labels correspond to those in Figure 2.

the effect of traffic forwarding is particularly evident because communication is responsible for the majority of the power budget, while it becomes almost negligible with DBP because it reduces the overall impact of communication on the power budget.

In an energetically sustainable system, the average power demand at each node must be met by the corresponding energy harvester. From Figure 4 it is clear that without DBP the consumption at nearly all nodes exceeds the renewable energy available by more than two orders of magnitude in some cases, resulting in a system that is not energetically sustainable. In contrast, with DBP the peaks of consumption due to data relaying are flattened and the average power consumption reduces to 15μW, a point where a photovoltaic energy harvester can meet the needs of all nodes, even those deep inside the tunnel.

In this analysis we considered only the *best* configurations with and without DBP. From this it is clear that only the combination of the wake-up receiver, dynamic power management and DBP can result in an energetically sustainable system.

VI. DISCUSSION AND FUTURE DIRECTIONS

As shown in Section V, the extremely low packet rates achievable through the application of prediction-based data collection techniques such as DBP combine effectively with advanced dynamic power management techniques that exploit wake-up receivers and the hardware support provided by ultra-low-power motes. Simulations results show that it is possible to decrease power consumption by two orders of magnitude, thus greatly extending the lifetime of battery operated WSNs and enabling energetic sustainability of the workload with energy-harvesting motes. Although this is only a single case study, it is representative of typical WSN applications and provides hints on the potential of the techniques under study.

Nevertheless, our conclusions come with some caveats. First, the ultrasonic wake-up receiver makes use of directional capsules, implying that the topology of the collection tree cannot vary dynamically over time. While this may be reasonable for a static system such as the road tunnel, the implications for long-term reliability under changing environmental conditions must be studied and discussed. Alternate solutions for omnidirectional wake-up receivers offer tremendous opportunities, but current costs and range limitations of such devices make them inapplicable in our scenario.

Additional restrictions come from the range of the ultrasonic wake-up receiver of VirtualSense (about 14m). In a later deployment by our group, distances between nodes in the tunnel reached up to 30m, a value that makes this particular wake-up receiver inapplicable.

To the best of our knowledge, the development of an omnidirectional wake-up receiver covering a range of tens of meters with sub-μW idle power is still an open issue and forms part of our plans for future work. The modular structure of the VirtualSense platform allows us to work on the wake-up receiver without impacting any other layer.

Also, we plan recreate part of the tunnel deployment with real hardware, directly measuring consumption to further validate the experiments presented here.

The enormous potential demonstrated here of the hardware-software co-design of WSN motes with energy harvesting, dynamic power management, and prediction-based data collection motivates further research efforts in this direction.

Acknowledgment

The authors would like to thank Riccardo Dall'Ora for his valuable input to this work.

References

[1] J. Polastre, J. Hill, and D. Culler, "Versatile low power media access for wireless sensor networks," in *Proc. of the 2nd Int. Conf. on Embedded Networked Sensor Systems (SenSys)*, 2004.

[2] A. Dunkels, "The ContikiMAC radio duty cycling protocol," Swedish Institute of Computer Science, Tech. Rep. Technical Report T2011:13, 2011.

[3] M. Ceriotti, M. Corrà, L. D'Orazio, R. Doriguzzi, D. Facchin, S. Guna, G. P. Jesi, R. L. Cigno, L. Mottola, A. L. Murphy, M. Pescalli, G. P. Picco, D. Pregnolato, and C. Torghele, "Is there light at the ends of the tunnel? Wireless sensor networks for adaptive lighting in road tunnels," in *Proc. of the Int. Conf. on Information Processing in Sensor Networks (IPSN)*, 2011.

[4] G. Anastasi, M. Conti, M. Di Francesco, and A. Passarella, "Energy conservation in wireless sensor networks: A survey," *Ad Hoc Networks*, vol. 7, no. 3, pp. 537–568, 2009.

[5] U. Raza, A. Camerra, A. Murphy, T. Palpanas, and G. Picco, "What Does Model-driven Data Acquisition Really Achieve in Wireless Sensor Networks?" in *Proc. of the 10^{th} IEEE Int. Conf. on Pervasive Computing and Communications (PerCom)*, 2012.

[6] E. Gaura, J. Brusey, M. Allen, R. Wilkins, D. Goldsmith, and R. Rednic, "Edge mining the internet of things," *Sensors Journal, IEEE*, vol. 13, no. 10, pp. 3816–3825, Oct 2013.

[7] F. Aderohunmu, G. Paci, D. Brunelli, J. Deng, L. Benini, and M. Purvis, "An application-specific forecasting algorithm for extending wsn lifetime," in *Distributed Computing in Sensor Systems (DCOSS), 2013 IEEE International Conference on*, May 2013, pp. 374–381.

[8] A. Deshpande, C. Guestrin, S. Madden, J. M. Hellerstein, and W. Hong, "Model-driven data acquisition in sensor networks," in *Proc. of the Int. Conf. on Very Large Data Bases (VLDB)*, 2004.

[9] D. Chu, A. Deshpande, J. M. Hellerstein, and W. Hong, "Approximate data collection in sensor networks using probabilistic models," in *Proc. of the Int. Conf. on Data Eng. (ICDE)*, 2006.

[10] D. Tulone and S. Madden, "PAQ: Time series forecasting for approximate query answering in sensor networks," in *Proceedings of the European Wkshp. on Wireless Sensor Networks (EWSN)*, 2006.

[11] ——, "An energy-efficient querying framework in sensor networks for detecting node similarities," in *Proc. of the Int. Conf. on Modeling, Analysis and Simulation of Wireless and Mobile Systems (MSWiM)*, 2006.

[12] A. Jain, E. Y. Chang, and Y.-F. Wang, "Adaptive stream resource management using Kalman filters," in *Proc. of the Int. Conf. on Management of Data (SIGMOD)*, 2004.

[13] E.-Y. A. Lin, J. M. Rabaey, and A. Wolisz, "Power-efficient rendez-vous schemes for dense wireless sensor networks," in *IEEE International Conference on Communications*, 2004, pp. 3769–3776.

[14] G. U. Gamm, M. Sippel, M. Kostic, and L. M. Reindl, "Low power wake-up receiver for wireless sensor nodes," in *Proceedings of the Sixth International Conference on Intelligent Sensors, Sensor Networks and Information Processing (ISSNIP), 2010*, ser. ISSNIP '10, 2010, pp. 121–126.

[15] C. Hambeck, S. Mahlknecht, and T. Herndl, "A 2.4uw wake-up receiver for wireless sensor nodes with -71dbm sensitivity," in *International Symposium on Circuits and Systems (ISCAS 2011)*, 2011, pp. 534–537.

[16] X. Huang, P. Harpe, G. Dolmans, and H. de Groot, "A 915mhz ultra-low power wake-up receiver with scalable performance and power consumption," in *Proceedings of the 37th European Solid-State Circuits Conference, ESSCIRC 2011*, 2011, pp. 543–546.

[17] S. J. Marinkovic and E. M. Popovici, "Nano-power wireless wake-up receiver with serial peripheral interface," *IEEE Journal on Selected Areas in Communications*, vol. 29, no. 8, pp. 1641–1647, 2011.

[18] V. Jelicic, M. Magno, D. Brunelli, V. Bilas, and L. Benini, "Analytic comparison of wake-up receivers for wsns and benefits over the wake-on radio scheme," in *Proceedings of the 7th ACM workshop on Performance monitoring and measurement of heterogeneous wireless and wired networks*, ser. PM2HW2N '12. New York, NY, USA: ACM, 2012, pp. 99–106.

[19] E. Lattanzi, M. Dromedari, V. Freschi, and A. Bogliolo, "A sub-ua ultrasonic wake-up trigger with addressing capability for wireless sensor nodes," *ISRN Sensor Networks*, vol. 2013, 2013.

[20] N. B. Priyantha, A. Chakraborty, and H. Balakrishnan, "The cricket location-support system," in *Proceedings of the 6th annual international conference on Mobile computing and networking, (MobiCom '00)*, 2000, pp. 32–43.

[21] G. Oberholzer, P. Sommer, and R. Wattenhofer, "Spiderbat: Augmenting wireless sensor networks with distance and angle information," in *Proceedings of the 10th International Conference on Information Processing in Sensor Networks, (IPSN 2011)*, 2011, pp. 211–222.

[22] U. Raza, A. Camerra, A. L. Murphy, T. Palpanas, and G. P. Picco, "Practical data prediction for real-world wireless sensor networks," Department of Information Engineering and Computer Science, University of Trento, Trento, Italy, Tech. Rep. DISI-14-004, April 2014. [Online]. Available: http://eprints.biblio.unitn.it/4274/

[23] E. Lattanzi and A. Bogliolo, "Virtualsense: A java-based open platform for ultra-low-power wireless sensor nodes," *International Journal of Distributed Sensor Networks*, vol. 2012, 2012. [Online]. Available: http://www.hindawi.com/journals/ijdsn/2012/154737/

[24] A. Dunkels, B. Gronvall, and T. Voigt, "Contiki - a lightweight and flexible operating system for tiny networked sensors," in *Proceedings of the 29th Annual IEEE International Conference on Local Computer Networks*, ser. LCN '04. Washington, DC, USA: IEEE Computer Society, 2004, pp. 455–462. [Online]. Available: http://dx.doi.org/10.1109/LCN.2004.38

[25] N. Brouwers, K. Langendoen, and P. Corke, "Darjeeling, a feature-rich vm for the resource poor," in *Proceedings of the 7th ACM Conference on Embedded Networked Sensor Systems*, ser. SenSys '09. New

York, NY, USA: ACM, 2009, pp. 169–182. [Online]. Available: http://doi.acm.org/10.1145/1644038.1644056

[26] E. Lattanzi and A. Bogliolo, "Hardware filtering of non-intended frames for energy optimisation in wireless sensor networks." *International Journal of Sensor Networks*, vol. 15, no. 2, pp. 121–129, 2014.

[27] L. Panasonic Eco Solutions Amorton Co., "Am-1816," http://panasonic.net/energy/amorton/en/products/spec/AM-1816.html.

[28] A. Bogliolo, E. Lattanzi, and V. Freschi, "Idleness as a resource in energy-neutral wsns," in *Proceedings of the 1st International Workshop on Energy Neutral Sensing Systems*, ser. ENSSys '13. New York, NY, USA: ACM, 2013, pp. 12:1–12:6. [Online]. Available: http://doi.acm.org/10.1145/2534208.2534214

[29] D. Porcarelli, D. Brunelli, M. Magno, and L. Benini, "A multi-harvester architecture with hybrid storage devices and smart capabilities for low power systems," in *Power electronics, electrical drives, automation and motion (SPEEDAM), 2012 international symposium on*. IEEE, 2012, pp. 946–951.