

Energy Neutral Operation of Vibration Energy-Harvesting Sensor Networks for Bridge Applications

Andrea Gaglione^{* 1,2}, David Rodenas-Herraz^{* 1}, Yu Jia³, Sarfraz Nawaz¹, Emmanuelle Arroyo¹,
Cecilia Mascolo¹, Kenichi Soga⁴, Ashwin A. Seshia¹

¹University of Cambridge, UK, ²Digital Catapult, UK, ³University of Chester, UK, ⁴UC Berkeley, USA

Abstract

Structural monitoring of critical bridge structures can greatly benefit from the use of wireless sensor networks (WSNs), however energy harvesting for the operation of the network remains a challenge in this setting. While solar and wind power are possible and credible solutions to energy generation, the need for positioning sensor nodes in shaded and sheltered locations, e.g., under a bridge deck, is also often precluding their adoption in real-world deployments. In some scenarios vibration energy harvesting has been shown as an effective solution, instead.

This paper presents a multihop vibration energy-harvesting WSN system for bridge applications. The system relies on an ultra-low power wireless sensor node, driven by a novel vibration based energy-harvesting technology. We use a receiver-initiated routing protocol to enable energy-efficient and reliable connectivity between nodes with different energy charging capabilities. By combining real vibration data with an experimentally validated model of the vibration energy harvester, a hardware model, and the COOJA simulator, we develop a framework to conduct realistic and repeatable experiments to evaluate the system before on-site deployment.

Simulation results show that the system is able to maintain energy neutral operation, preserving energy with careful management of sleep and communication times. We also validate the system through a laboratory experiment on real hardware against real vibration data collected from a bridge. Besides providing general guidelines and considerations for the development of vibration energy-harvesting systems for bridge applications, this work highlights the limitations of the energy budget made available by traffic-induced vibrations, which clearly shrink the applicability of vibration energy-harvesting technology for WSNs to low traffic applications.

Categories and Subject Descriptors

C.3 [Special-Purpose and Applications-Based Systems]: Real-time and embedded systems

General Terms

Design, Simulation, Experimentation

Keywords

Wireless Sensor Networks, Energy Harvesting, Vibrations, Supercapacitors, Data collection

1 Introduction

Cost and long-term performance of remote monitoring technology are key barriers to large-scale adoption in civil engineering industry. Critical infrastructure assets, e.g., suspension bridges, transport systems, are built to last for decades. The UK road/rail network includes more than 40,000 bridges and over crossing structures for road, rail, and energy networks [3]. Many of the smaller bridges may be invisible to most people as they travel around but the big iconic bridges are easily recognizable. A failure in any of these structures can also have catastrophic social and economic consequences. Some of the most critical assets already use some form of remote monitoring, but the maintenance of over 95% of UK bridges relies on visual inspection on a scheduled basis [25].

With these premise, there is a pressing interest in monitoring applications of these structures to check structural movements or damages, e.g., displacement and deformation of key structural elements due to temperature changes, stress of the bridge deck under long-term vehicle load, tilt of steel bridge members. Traditional wired methods are expensive, with much of the cost derived from cabling and installation. Furthermore, wiring would require integration in the structure which might not be feasible, would disrupt appearance or disturb normal operation. In contrast, wireless sensor networks (WSNs) have the ability to provide the same functionality at a lower price, while facilitating the ease of deployment. Although WSNs have been used for monitoring bridges and other civil structures [14, 22, 21, 6], the conservation of energy to prolong network lifetime is a crucial aspect that has been stimulating the research in this field for over a decade.

As battery replacement is cumbersome and usually infeasible, the idea of powering sensor nodes with energy harvested from the environment (e.g., vibrations, sunlight, wind, temperature gradient) has become a key factor to enable the development of energetically autonomous WSNs, by achieving *energy neutral operation* which refers to a mode of operation

*These authors contributed equally to this work.

where the average power consumption of the nodes is less or equal than the average power harvested from the environment. Designing energy-harvesting WSNs is challenging because of the complex tradeoffs arising from the interaction between energy sources, energy storage device used, communication protocol, and application requirements. Energy-harvesting power supplies based on sunlight [28, 19], vibrations [31], and temperature differences [23] have been presented for sensor nodes. Although multihop data collection protocols have also been proposed, they are designed for solar cell operated sensor nodes [36, 29].

This paper presents the design, implementation, and evaluation of a vibration energy-harvesting WSN for bridge applications. As these latter often require nodes to be placed under the bridge deck—where only little energy can be harvested from sunlight or wind—we based our system on a sensor node powered by traffic-induced vibrations. Compared with other energy sources (e.g., sunlight), vibrations provide lower performance in terms of power density [34] and this has a great impact on the set of applications that can be targeted by our system as well as on the design and calibration of the multihop communication protocol. While other works have studied single node based vibration energy-harvesting sensor systems [31], to our knowledge, *this is the first comprehensive work targeting vibration energy-harvesting sensor networks*.

In detail, the paper introduces the following contributions:

- We present a vibration energy-harvesting sensor node, powered by a harvester able to be driven into both direct resonance and autoparametric resonance. The sensor node is based on an ultra-low power Ferroelectric RAM (FRAM) platform and uses a supercapacitor as energy buffer.
- We develop a comprehensive framework including a power profiling methodology and an energy-harvesting software to conduct realistic simulated experiments before on-site deployment. The power profiling methodology leverages real traffic data collected from a bridge during a preliminary acquisition campaign to derive the charge and discharge rates of the supercapacitor. To ensure multihop data communication, the energy-harvesting software includes a simple yet effective routing scheme. The latter is based on an existing receiver-initiated MAC protocol, opportunely adapted to work under the challenging restrictions of the limited vibration power.
- We achieve energy neutral operation of a 4-hop sensor network by simulating low traffic applications and using the charge models derived from the power profiling methodology. We rely on COOJA/MSPSim [12] environment: to this end, we ported the Contiki [9] operating system to the sensor node platform used in this study and also extended COOJA/MSPSim accordingly. Our hardware design and energy considerations are validated through laboratory experiments on real hardware.

The ultimate goal of this study is to achieve energy neutral operation and enable the development of energetically

autonomous WSNs for bridge applications, using vibration harvesters to power sensor nodes. Although we rely on real vibration data collected from a single bridge to design the system and the software architecture, the proposed framework, including the power profiling methodology, can be easily applied to other settings. In fact, both the harvester and its model can be configured to suit the unique vibrational responses of different bridges. Therefore, this paper is meant to provide general guidelines and considerations for the development of vibration energy-harvesting systems based on traffic-induced vibrations on bridges. The study highlights the limitations of the energy budget made available by traffic-induced vibrations, which clearly shrink the applicability of vibration-based harvesting technology for WSN to low traffic applications.

We discuss potential applications and the challenges introduced by vibration energy harvesting in Section 2. We provide a concise survey of related work in Section 3. Section 4 describes the hardware and system design, including a power profiling methodology to infer energy intake from real vibration data. The results of the power profiling methodology are leveraged in a comprehensive framework, including an energy-harvesting software for WSNs, which is introduced in Section 5. Section 6 describes the multihop communication protocol. Our evaluation, in Section 8, is divided in two parts: (i) we demonstrate the ability to achieve energy neutral operation in a 4-hop WSN over one-week simulation time; (ii) we validate the viability of multihop communication in a laboratory test on real hardware. Finally, Section 9 offers brief concluding remarks.

2 Background and Motivation

Various harvesting solutions can be used to power sensor nodes. Solar, wind, and vibrations are the most accessible ambient energies on a bridge. Among the different harvesting options, solar is the most mature and most commonly used solution in bridge deployments [32]. The amount of energy obtained depends on the context (e.g., solar exposure). Solar panels produce sufficient energy on sunny days, but energy is often limited or non-existent during the night or cloudy periods. The amount of harvested solar energy can also be dramatically diminished when excessive dust accumulates on the solar panels [8]. In addition, solar panels might not be usable if sensor nodes have to be placed on the underside of the bridge. In these settings, complex electrical wiring must be routed from the solar panels on the bridge surface to power the sensor nodes, which can be expensive and labor-intensive. Wind energy harvesting has also been investigated [26]. However, installation in bridges may likewise be a complex task and not always feasible due to the different wind flow patterns and potential requirements for wiring.

Vibrations provide an alternative source of energy [30, 27, 17]. A bridge can experience vibrations through passing vehicle traffic, wind, and even people. The vibration amplitude and its frequency are the main factors that determine the amount of energy harvested. The energy also depends on the absolute attainable power that a vibration harvester can generate. In addition, environmental changes can affect and diminish the energy that the harvester produces over time. An

additional challenge is that vibration energy is usually neither constant nor continuous along a bridge. The vibration amplitude can vary considerably from one location to another, and will depend on factors such as the type of structural member which the harvester is attached to, and its proximity to the abutments and to the supports of the bridge. Likewise, the excitation produced by vehicle traffic is known to be location dependent, non-stationary, and substantially transient in nature [20].

Most of the research on structural monitoring of bridges with WSNs primarily focuses on continuous (near-) real-time monitoring of structural physical aspects, such as operational loadings and structural responses [37, 21]. Such WSN deployments require vibration data to be sampled and transmitted at very high rates, which is problematic in networks powered solely by harvesters (and even by batteries) as these requirements lead to a fast depletion of the stored energy.

In contrast, this study aims to provide an effective system for long-term monitoring of temperature and structural movements, e.g., displacement, deformation, tilt, in key structural elements of bridge structures. These applications are delay tolerant and typically require a sampling rate of a few samples per hour, or even per day, as changes in the parameters being measured are usually expected to occur very infrequently [16]. Consider for example thermal loads due to temperature changes, which are an important factor driving the serviceability limit state design of steel bridges. Current engineering practice relies on standards for the design of structures, such as the Eurocode 1 [1], as well as on recommended safety factors, either on expected temperature ranges or on the estimated thermal movements. However, this practice is overly simplistic as it only considers ambient air temperature, which can lead to excessively overestimated thermal movements or, even worse, serviceability failures (e.g., locked expansion joints [5]). The temperatures of the main structural components of bridges may considerably differ from each other, and from the air temperature, at any point in time. Understanding the relationship between the movements of key structural elements, such as deck expansion joints and girders, and their temperature throughout their service life is crucial to assess the thermal behavior of bridges against their serviceability limit state design. Furthermore, assessing the thermal behavior of bridges is needed to enable better design practice, as well as to early detect potential failures [7].

In this context, temperature and displacement measurements can be collected using a WSN consisting of nodes mounted onto the expansion joints and girders. Besides temperature sensors, linear potentiometers or extensometers can be used to measure displacements of the expansion joints as well as to identify any possible quasi-static rotation of the main span deck about the vertical axis. Instead, foil strain gauges can be used to measure horizontal and vertical deformations of the girders. As expansion joints and girders are located *under* the bridge deck, a vibration energy-harvesting system is regarded as a more suitable alternative than other harvesting solutions. As changes in temperature and displacement are fairly static, data reporting on an hourly or even a daily basis would suffice. Collected data are processed off-line and used to derive models which relate temperature and

movements at each element in order to assess the thermal behavior of the bridge.

This work is also geared towards long-term monitoring of the inclination (tilt) of bearings underneath the bridge deck. Visual inspection is still the most common practice in assessing the safety of bridges, however having inspectors visiting the bridge to determine the progression of these deficiencies over time (many years) is cumbersome. Even in this case, a vibration energy-harvesting WSN is a most sought-after solution and typical applications require a low sampling rate [16]. Besides periodic data transmissions, the application at hand may also consider to send notification messages to any key stakeholders (e.g., bridge operators) whenever the tilt exceeds a structural margin of safety.

3 Related Work

Vibration energy harvesting introduces additional challenges to communication protocol design in WSNs. Popular medium access communication (MAC) protocols, such as XMAC [4], ContikiMAC [9], and Time-Slotted Channel-Hopping (TSCH) MAC [11] rely on the use of radio duty cycling to save energy by removing unnecessary overhearing and idle listening. Radio duty cycling introduces communication latency and therefore such protocols are designed to turn the radio on at regular intervals to limit latency while achieving a reasonable throughput. To regulate energy consumption according to the availability of harvested energy, sensor nodes will have to interleave very short transmission periods with long sleep periods necessary to accumulate energy. This will result in a higher delivery latency, which is acceptable to support the general low-rate applications targeted by this study. Receiver-initiated duty-cycling protocols have been found to be appropriate for energy harvesting WSNs [13, 29]. Such approach is adopted in this work to develop an opportunistic multihop solution that provides reliable data transmission while preventing node failures due to energy depletion.

The use of vibration energy from bridges to power wireless sensors has been investigated in a limited number of real-world experiments [31, 24]. Sazonov et al. [31] describe an acceleration-based WSN system positioned under the bridge deck that harvests vibrations induced by road traffic. The study uses a charge-sample-transmit strategy, which in practice may suffice for single-hop topologies, while is not enough for multihop networks, which the study presented in this work aims to target. McCullagh et al. [24] report results from a long-term single-hop deployment of a number of wireless accelerometers powered by vibration harvesters (together with other battery-powered sensors) at New Carquinez suspension bridge. Given the high sampling rate and transmit power requirements, the vibration harvester required solar panels in order to increase the production of harvested energy. Interestingly, most problems experienced in that work were found to be caused by degradation of the rechargeable batteries. In this work, we use supercapacitors for buffering energy as suggested in [38, 29]. They are relatively small, have virtually unlimited number of charge-discharge cycles and long shelf life, thus enabling perpetual and maintenance-free operation for energy-harvesting sensor networks.

This paper presents a novel comprehensive study on vi-




Figure 1. Hardware architecture of the energy-harvesting sensor node.

bration energy-harvesting sensor networks for bridge applications. Besides achieving energy neutral operation of a multihop sensor network in simulation for a certain set of applications, it provides general guidelines for designing and developing similar systems based on traffic-induced vibrations on bridges.

4 System Design

In this section, we first present the hardware architecture of our vibration energy-harvesting sensor node. To show the practicality of the harvester, we then describe the power profiling analysis used to infer energy intake using real vibration data collected from a bridge. The results of the analysis will be further leveraged into a comprehensive framework targeting a whole sensor network, as detailed in Section 5.

4.1 Hardware

The hardware architecture of our vibration energy-harvesting sensor node is shown in Figure 1. It includes the following main components: a vibration energy harvester, a power supply unit, and the sensor node. The harvester uses vibrations induced by passing traffic on a bridge to generate power. The harvested energy is then stored and regulated by the power supply unit, which in turn drives the wireless sensor platform.

Vibration energy harvester

The vibration energy harvester employed [18] in this study has been designed to match the frequency vibrations measured at Tamar Bridge that were found to provide the higher energy intake. The mechanical subsystem used to amplify the mechanical energy inside the harvester has the ability to be driven into both direct resonance and autparametric resonance, which have shown the potential to fundamentally enhance both the recoverable power amplitude and the operational frequency bandwidth of the harvester. The conversion of the mechanical energy into raw electrical energy is performed by an electromagnetic transduction system. The packaged prototype has a total volume of 126 cm³.

Power supply system

The power supply system consists of three main components: the power conditioning circuit, a voltage multiplier, and a storage supercapacitor. The power conditioning circuit

aims to control the charge of the supercapacitor, which acts as an energy reservoir for the harvested energy. It also provides a regulated constant voltage supply to the sensor node, which is only operational when the supercapacitor is sufficiently charged. We used the LTC3588-1 circuit from Linear Technology to provide a constant output voltage of 2.5 V. The power conditioning circuit by itself is unable to burst the energy charge. Therefore, a six-stage Dickson voltage multiplier is used upstream to rectify the AC input signal from the vibration harvester, and boost the voltage level from its input. When the voltage across the supercapacitor reaches the start-up voltage of the circuit, namely V_{on} , which is 4.04 V for the LTC3588-1, the latter is brought into regulation and the node is switched on. If the voltage falls below the cut-off voltage of the circuit, namely $V_{cut} = 2.7$ V, the node is switched off until the voltage across the supercapacitor reaches V_{on} again.

Sensor node

The wireless sensor platform relies on the ultra-low power Ferroelectric RAM (FRAM) based MSP430FR5969 microcontroller coupled with an off-the-shelf TI CC2520 radio transceiver (IEEE 802.15.4-compliant). FRAM technology offers faster write operations, much higher endurance for read and write operations, and much lower power consumption compared to equivalent flash memories. The TI MSP430FR5969 also has very fast wake-up time, low active current consumption (2.51 mA @ 8MHz), and ultra-low current consumption in lower power modes (0.3 μ A in LPM4). This allows to perform software optimizations to substantially minimize energy consumption. Although the amount of available RAM is limited to 2 KB, the built-in 64 KB FRAM memory can be opportunely used when needed (e.g., to implement a packet queue).

4.2 Preliminary Data Acquisition

In collaboration with the Tamar Bridge and Torpoint Ferry Joint Committee, we conducted an acquisition campaign on Tamar Bridge—a major road bridge in the South West of England—for two days to acquire vibration data under real traffic conditions. The goal of the campaign was twofold. First, we wanted to collect enough vibration data under real traffic conditions in order to be able to process them offline and power profile our hardware in simulation, thus understanding the energy limitations for the development of energy-harvesting WSN applications. Second, we aimed to do an on-site inspection to find the best locations for vibration energy harvesting, and also evaluate the communication link quality for a future real WSN deployment.

Tamar Bridge is illustrated in Figure 2. It spans across the River Tamar, between the city of Plymouth and the town of Saltash, providing four vehicle lanes and one pedestrian walkway. It is a symmetrical suspension bridge having a total length of 563 m over three spans. Reinforced concrete towers, both 73 m above their caisson foundations, connect a main span of 335 m to two side spans of 114 m each. To acquire real traffic-induced vibration data, we deployed two triaxial wireless accelerometers (BeanDevice[®] AX-3D) in two different locations. The BeanDevice[®] AX-3D offers a very low noise density (45 μ g/ \sqrt{Hz}), enabled by a 5th order Butterworth anti-aliasing filter, and a $\pm 2g$ measurement



Figure 2. Tamar Suspension Bridge (UK): view from the Devon bank.




Figure 4. South elevation diagram showing accelerometer locations.

range. Data acquisition happened in parallel by positioning the accelerometers at the two locations to capture the vibrations. The fundamental frequencies of most civil structures are below 10 Hz [21] and according to the Nyquist theorem, sampling rate should be at least twice than that. However, the natural frequencies of individual elements of the bridge, e.g., stiffeners, hangers, were unknown in advance and likely to have higher vibration modes. Therefore, a sampling rate of 200 Hz was chosen.

Data was stored locally at the device and downloaded in bulk at the end of the data acquisition period. The locations were chosen by trial and error aiming to find structural members showing high amplitude vibrations in the time domain—high power spectral density (PSD) in the frequency domain—which could then lead to a better energy harvesting. Best locations seemed to be at the stiffening members under the bridge deck, as shown in Figure 3. With respect to the main motivating application scenarios described in Section 2, we believe that even though necessary wiring could be required between harvesters on the stiffeners and sensor nodes on structural elements located under the bridge deck, e.g., expansion joints, girders, bearings, that would be minimal as compared to solar-based harvesting solutions. In fact, the latter would imply solar panels to be placed on the outer side of the bridge and wiring would be overly complex.

To understand the feasibility of multihop communication, range tests were carried out to determine the most appropriate distance between locations. The variety of structural members including girders, bracings, stiffeners and walkway made of steel under the bridge deck greatly impeded wireless communication between the accelerometer nodes. A final distance of approximately 27 m was chosen. Figure 4 shows




Figure 3. Accelerometer attached to a stiffening member.




Figure 5. Average vehicle traffic data, February 2016. The shaded areas indicate our acquisition times.

the accelerometer locations on the bridge, close to the west (Saltash) tower. We will refer to these locations as L_1 and L_2 hereafter, where L_1 is the closest to the Saltash tower.

We collected vibration data over one-two hours at different times of the day on two weekdays. The datasets collected from the two locations refer to the following four time intervals: (I) 7.34AM - 8.51AM, (II) 9.56AM - 11.04AM, (III) 1.33PM - 2.36PM, and (IV) 3.53PM - 4.55PM. To correlate the four datasets with the vehicle traffic load at the bridge, the average traffic count data per hour is shown in Figure 5. Such data was collected by the Tamar Bridge and Torpoint Ferry Joint Committee in the direction of the Tamar Bridge tolls, the same side of the bridge where the accelerometers were deployed. Figure 5 also shows the mean and the variance of the total number of vehicles (*traffic counts*) crossing per hour at weekdays and weekends, and averaged for the period relative to February 2016. As observed, dataset I corresponds to the morning rush hour, when the highest traffic load flows from Saltash to Plymouth, while dataset IV corresponds to the afternoon rush hour, when the highest traffic load flows in the opposite direction.

The PSDs of the recorded vibration data have been investigated to identify which acceleration axis and natural frequency at each location have the highest power content. PSDs of the dataset I at both locations are presented in Figure 6. The latter shows the signal power in a joint time and frequency domain by using the short-time Fourier transform (STFT). This analysis is useful to determine the correct orientation of the harvester for a future WSN deployment, as well as to design and tune the harvester to an optimal frequency for harvesting as much energy as possible. For both locations and for all the four datasets, vibrations are more significant




Figure 6. PSD (left) and STFT spectrograms (right) of the vibration dataset I at two stiffening members of Tamar Bridge.

in the axis perpendicular to the ground (Z). Similar results have been obtained for datasets II, III, and IV, whose plots are omitted for the sake of brevity. For all datasets, the highest peak of power at location L_1 was found to be 9.1 Hz (fundamental frequency of that specific stiffener), while the highest peak at location L_2 was found to be 18.7 Hz (second natural frequency of that specific stiffener).

4.3 Power Profiling

We now describe how we used the vibration data to infer energy values which could be harvested on Tamar bridge. A numerical model of the vibration energy harvester was developed in MATLAB based on an experimental prototype deployed on the Forth Road Bridge [18]. The model was tuned and validated against real vibration and power output data from both field site testing and lab-based experiments. To accommodate the vibration data collected from Tamar Bridge, the natural frequency of the previous prototype model was shifted from 13 Hz—the frequency used for the experiments at Forth Road Bridge—to 9.1 Hz to reflect a harvester at location L_1 , and to 18.7 Hz to reflect a second harvester at location L_2 . While some realism is lost through the shift in natural frequency, it is still based on a near identical model that has been experimentally validated. By using the numerical model,




Figure 7. Charge of the supercapacitor obtained from LTspice simulation experiments.

the vibration data from Tamar Bridge was transformed into AC output voltage data.

A hardware model of the architectural components introduced in Section 4.1 was then built in LTspice IV to empirically quantify and characterize the charge-discharge behavior of the supercapacitor, using the AC output voltage data generated by the numerical model of the vibration harvester. A key issue in the design of the power supply system relates to the supercapacitor to be used. The energy E stored in a supercapacitor is expressed as:

$$E = 1/2 \cdot C \cdot V_{cap}^2,$$

where V_{cap} is the voltage across the supercapacitor and C is its capacitance. A supercapacitor with a capacitance of the order of mF would charge very quickly, but the amount of energy it can store would be limited. Bigger supercapacitors are able to store more energy, although accumulated at lower rates. We ran LTspice simulations to find a supercapacitor able to provide a good trade-off between size, capacitance, and charge-discharge rate as well as to understand the average harvested power from the vibration data collected from Tamar Bridge. We observed that supercapacitors between 1 F and 3 F can be suitable to operate a sensor node and sustain the beaconing process of a communication protocol for low traffic applications. For such process, it may suffice to only transmit a few small packets, namely *beacons*, per minute. For this reason, we chose a supercapacitor with a capacitance of 2 F.

Figure 7 shows the charge voltage curves of a 2 F supercapacitor obtained from LTspice, while Table 1 reports the amount of average harvested power available from our vibration datasets.

We observe different charge rates of the supercapacitor due to different harvested power from the datasets related to the two locations and four datasets. The morning rush hour dataset provides the fastest charge rate at location L_2 , followed by the non-rush hour datasets and lastly by the afternoon rush hour (high traffic flow in the opposite side of the accelerometers locations) dataset. Interestingly, location L_1 presents a slightly different behavior with non-rush hour datasets providing a faster charge, followed by the morning rush hour and the afternoon rush hour datasets. This can be explained as follows: as described above, the natural frequency of the harvester model has been tuned to 9.1 Hz and 18.7 Hz for locations L_1 and L_2 , respectively. However, the vibration spectrum for L_1 related to the morning rush hour generates low power density, also showing a relatively low power content at 9.1 Hz, as illustrated in Figure 6.

Table 1. Average harvested power available.

Dataset	Avg. power (μW)	
	L_1	L_2
I (7.34AM - 8.51AM)	401	661
II (9.56AM - 11.04AM)	562	403
III (1.33PM - 2.36PM)	459	450
IV (3.53PM - 4.55PM)	249	138

The LTspice model is also used to obtain the actual power consumption of the sensor node when connected to the power supply system. The current consumption values used in simu-




Figure 8. Toolchain for hardware modeling and power profiling.

lation are taken from the data sheets of the MSP430FR5969 MCU and the CC2520 radio. In our system design, the low power mode LPM4 of the MSP430FR5969 MCU and the power mode LPM2 of the CC2520 radio are used to save energy when the node is in sleep mode. Simulation showed an average consumption of $300 \mu\text{W}$ when the MCU is active and the radio is either in idle mode, transmitting or receiving. Almost null consumption is reported when both the MCU and the radio are in their above mentioned low power modes. Approximately 3 minutes are enough to make V_{cap} fall from 4.04 V (V_{on}) to 2.7 V (V_{cut}) when using a 2 F supercapacitor, and when the MCU is active and the radio is in reception mode, which gives an idea of the lifetime of the sensor node when no duty cycle scheme is adopted to save power.

5 Toolchain and Software Architecture

We leverage the results of the power profiling in a comprehensive framework that allows developers to conduct realistic simulated experiments to evaluate the system before on-site deployment. The toolchain used for hardware modeling and power profiling is summarized in Figure 8.

As described in Section 4.3, we use a two-step modeling process to characterize the voltage charge curves of the supercapacitor. First, we feed a MATLAB model of the vibration harvester with real acceleration data to obtain the corresponding output voltage values. Second, we use a LTspice model of our energy-harvesting sensor node, including the power supply unit, in order to obtain the voltage charge curves across the supercapacitor.

To simulate various charging patterns in several times of the day, we fit second-order polynomial models to the voltage charge curves and use the models to feed a simulated energy-harvesting software, whose main architectural components are shown in Figure 9.

We rely on the COOJA simulator [12] and specifically MSPSim, a hardware emulator for the MSP430 MCU which enables to directly reuse the code simulated into real world. Our reliance on COOJA/MSPSim is also motivated by the ability to use its scripting language to import the models




Figure 9. Energy-harvesting software architecture.

above into a custom *script*. The latter leverages the models to periodically compute V_{cap} , the voltage across the supercapacitor that is supposed to power a (simulated) sensor node. The *energy monitor* updates the V_{cap} values by estimating the power consumption of the node, reusing the Energest [10] component of the Contiki distribution [9]. It also ensures that V_{cap} never falls below V_{cut} , the cut-off voltage of the power conditioning circuit, which would imply to wait a substantial amount of time for it to switch back only when V_{cap} reaches V_{on} . The *application* is a simple program to trigger sensor readings on a node at a given sampling rate. Finally, a *multihop data communication* component implements a routing protocol that ensures connectivity between the nodes and the sink, while its *queue management* module maintains a packet queue at the node and trigger packets transmission. Details of the routing protocol are presented in the next Section.

6 Multihop Data Communication

State of the art in vibration-based energy-harvesting WSNs is restricted to one-hop networks [31], however in most bridge applications, such as those described in Section 2, multihop communication towards the sink is required.

Although many routing protocols for sensor networks have been designed with a focus on energy efficiency (e.g., CTP [15]), their consumption overhead for link quality assessment and route maintenance is not acceptable in energy-harvesting WSNs. As explained in Section 4.3, power-hungry operations, such as transmission and reception, need to be reduced in order for the network to remain energetically sustainable. Recently, it has been demonstrated that a novel class of opportunistic routing protocols addressing those needs can be used for data collection in solar-powered WSNs [13, 35, 29]. In the following, we describe how such research findings have been leveraged and adapted to design a scheme for WSNs powered by traffic-induced vibrations, with very limited harvested power available, as reported in Section 4.3.

The routing protocol developed in this work relies on RIMAC [33], a receiver-initiated asynchronous duty-cycle MAC protocol for WSNs, which has been tweaked to allow for multihop data communication. The motivation behind the




Figure 10. Simplified representation of RI-MAC with routing capabilities. Consecutive transmission of two data packets triggered after checking the hop count sent within the beacon.

adoption of a receiver-initiated scheme is that these protocols minimize the amount of time any arbitrary pair of sender and receiver nodes within communication range occupy the wireless medium in comparison with sender-initiated approaches. This is crucial for vibration energy-harvesting WSNs, where periods of communication activity must be limited as much as possible.

In RI-MAC, every node periodically wakes up according to its own duty cycle, defined by the *off time* period, and immediately broadcasts a beacon frame to inform potential sender nodes within communication range that it is ready to receive data. Similarly to [29], we enhance RI-MAC to become a routing protocol and enable multihop communication towards the sink. In fact, the beacon frames in our protocol include information about the node hop distance to the sink (*hop count*). Any sender node with pending data to be transmitted will stay awake listening for the beacon from any potential receiver for approximately the duration of one wake-up interval. If after such listening period the sender does not receive any suitable beacons, it will switch off. Otherwise, upon reception of a beacon frame, the sender checks whether the hop count of the potential receiver which sent the beacon is less than the sender's hop count. If so, the sender updates its hop count and then transmits the data. Once data transmission is over, the sender keeps listening for the acknowledgment frame from the receiver, confirming the successful reception of data.

We leverage the capability of RI-MAC to reliably send streams of data packets in a single wake-up interval. In fact, the queue management module introduced in Section 5 allows a node to store its own packets in a queue, along with the packets received from its neighbors, and forward them according to a configurable queue size. The goal is twofold. First, nodes can save much more energy transmitting data in bulk in a single wake-up interval rather than employing several wake-up intervals to transmit the same number of packets. Second, the communication bandwidth can be fully utilized, while reducing in turn the overall communication activity in the network and the chances of packet collisions.

The price for this is an increase in the *packet delay*, but this is not an issue for the applications targeted by this study. An example of data transmission of two consecutive packets in the same wake-up interval using our routing protocol, with packet queue size = 2, is illustrated in Figure 10.

Finally, we adapted the encounter optimization mechanism of the Contiki implementation of X-MAC [4] to further reduce the energy wasted in idle listening. After a first encounter, a sender node learns about the wake-up interval of its neighboring nodes. Therefore, instead of listening to the wireless medium for an entire wake-up interval, the sender selects a node from its neighbor table which has a hop count less than its own and wakes up only right before it expects such neighbor to wake up, by also accounting for the rate of hardware clock drift.

7 Implementation

All the components of the vibration energy-harvesting application, including the routing protocol are implemented for the Contiki OS [9]. As introduced in Section 5, to simulate various charging patterns for several times of the day, we fit second-order polynomial models to all the voltage curves resulting from the power profiling, shown in Figure 7. According to traffic statistics and real vibration data from Tamar Bridge, illustrated in Figure 5a, we are able to model four different charging patterns for each of the two locations considered to acquire the acceleration data. Specifically, for each location, we assume that the models fitting the voltage curves resulting from datasets I - IV cover the time intervals: 6AM - 9AM, 9AM - 12PM, 12PM - 3PM, and 3PM to 8PM, respectively. We also assume no traffic over the night, which is an extreme case and an underestimation of the real conditions. Figures 11 and 12 show the voltage curves and related models for the two locations.

We implemented a script for the COOJA simulator to compute the values of the voltage across the supercapacitor V_{cap} once every 30 s using the aforementioned models, as detailed next. Initially, $V_{cap} = V_{on}$ for all the nodes, where V_{on} is the start-up voltage of the power conditioning circuit. Such values are then sent to the nodes via their emulated serial port and used by the energy monitor to assess their energy status. Specifically, the energy monitor reuses the Energest component of the Contiki distribution, which provides a mechanism to estimate the power consumption, and updates V_{cap} according to the discharge curve of the supercapacitor. The updated value is fed back to the script, which exploits the models to compute the new V_{cap} value after a period of 30 s.

According to the implementation of the queue management module of the routing protocol, every node maintains a packet queue and only triggers transmissions if the number of packets in the queue is at least $N \geq S$, where S is a user-defined queue size. We store the queue in the 64 KB FRAM memory of the TI MSP430FR5969. The maximum queue size is limited by the size of the binary program image uploaded on the node. However, $S \leq 64$ ensures a sufficient amount of FRAM for a typical WSN program. The binary program image of the current implementation for the TI MSP430FR5969 has a size of 25.6 KB (not including the packet queue) and uses 1.5 KB RAM.




Figure 11. Charge of the supercapacitor obtained from LTspice simulations and related second-order polynomial models. First location (L_1).

We adapted the encounter optimization mechanism of the Contiki implementation of X-MAC. As we target applications that typically require very low sampling rates, e.g., few samples per hours, and leverage a packet queue that delay their transmissions, encounters are likely to happen very rarely, e.g., few times per day. For this reason, to record encounters and optimize the sender wake-up times, we calculate clock differences with the granularity of 1 s, by also accounting for the rate of hardware clock drift. Although this substantially increase the idle listening, it is acceptable as nodes are not supposed to transmit very often.

As an essential part of this work, we ported the Contiki OS to the TI MSP430FR5xxx MCU series, which was unavailable at the time of writing this paper. We have also extended COOJA/MSPSim to enable instruction level emulation of the wireless sensor platform used in this study. All the software components, the hardware model, and the collected vibration datasets are available online [2].

8 Evaluation

In this section, we present the results of our evaluation. Section 8.1 assesses the effectiveness of the framework and the sustainability of the routing protocol through a simulation energy analysis, considering a 4-hop WSN and using the charge models obtained from real vibration data. The benefits of the protocol come at the cost of increased packet delay, which is acceptable by the particular class of applications targeted by this study. Section 8.2 presents a laboratory experiment to validate both the hardware design and the beaconing process of RI-MAC on real hardware, by using a real unit of the vibration energy harvester considered in this study—described in Section 4—and real vibration data collected at Tamar Bridge.

8.1 Simulation Experiments

To concretely illustrate the use of the framework, we resort to low traffic applications, which motivated the research presented in this paper, as described in Section 2. We use the COOJA/MSPSim time-accurate simulator to demonstrate energy neutral operation of a sensor network, and show that the




Figure 12. Charge of the supercapacitor obtained from LTspice simulations and related second-order polynomial models. Second location (L_2).

system maintains functionality over time, preventing energy depletion.

Setup and parameters

We consider a relatively small WSN, namely a chain of 5 nodes, including the *sink*, with non-overlapping ranges, as illustrated in Figure 13. The nodes are based on the TI MSP430FR5969 FRAM MCU, coupled with a TI CC2520 transceiver, and the network is supposed to monitor one side of the main span of Tamar Bridge (see Figure 2). We use the first set of charge models described in Section 7 for nodes A and C, whereas the second set is used for nodes B and D. We set a 2 F supercapacitor on every node. The latter run the software introduced in Section 7.




Figure 13. WSN topology in the simulated experiments.

We configure the routing protocol with an off time period of 16 s on all the nodes except the *sink*. As the latter is not energy constrained, we set an off time period of 128 ms. We use a 32-bit integer value to represent a single sensor reading of a physical quantity. A packet is then created with the data. No compression or in-network data aggregation is performed by the application. According to the datasheet of the LTC3588-1 power conditioning circuit, $V_{cut} = 2.7$ V and $V_{on} = 4.04$ V were chosen. The time interval between two consecutive simulated readings (*sampling period*) is set to 10 min. Each node in the network create a data packet at each sampling period and push it in the packet queue S , whose size was set to 16. The packet queue is stored in the program memory (FRAM) and the chosen size ensured a sufficient amount of memory for the program code. The queue allows to save energy by streaming the stored data packets in a single wake-up interval and the transmission of each data packets is retried at most 4 times. Even if the queue introduces a packet delay in the system, this is acceptable for the delay tolerant




Figure 14. V_{cap} traces over one-week simulation time.

applications targeted by this study.

Results

To evaluate the effectiveness of the system and assess its ability to maintain energy neutral operation, we run a one-week simulation test by tracing the values of V_{cap} at each node. The results are shown in Figure 14. Simulation starts at 6am of the first day, when $V_{cap} = V_{on} = 4.04$ V for all the nodes, and terminates after seven days at the same time.

Overall, the chart shows a regular alternating pattern of V_{cap} on all the simulated nodes. During the first part of each day, the balance between energy intake and energy consumption is mostly positive at every node, in which an increase of V_{cap} can be observed. It then decreases in the afternoon (after 3PM) and at night, when the energy balance is negative. This can be explained by the fact that subtle energy intake, due to very small vibrations, is observed in the afternoon between 3PM and 8PM in both the collected data sets (with the second one being the worst)—as shown by the related charge models in Figures 11 and 12—against a much higher energy consumption, which makes the contribution of the energy intake to the balance negligible. After 8PM, the negative gradient of V_{cap} becomes even more steep as we assumed no traffic/vibrations, namely no energy intake at all at night. The latter assumption allowed to perform a worst-case energy analysis at night, but was also dictated by the fact that we did not collect real vibration data at night during our preliminary acquisition campaign on Tamar Bridge because of access restrictions to the site.

Interestingly, the pattern of V_{cap} for node A is above the patterns of all the other nodes, thus exhibiting an overall better energy balance. In the first instance, this might seem against the flow as node A is the closest to the *sink* and more transmissions are required to forward the data coming from nodes that are longer hops away. However, it is due to the low off time period of the sink, which minimizes the idle listening of node A while waiting for a beacon after it turns its radio on. As discussed in Section 6, we adapted the encounter optimization mechanism provided by the Contiki implementation of X-MAC to reduce idle listening on nodes that are more than one hop away from the *sink*. However, idle listening

still highly contributes to the depletion of nodes B, C, and D because of the 1-second granularity of the long-running timer used. Nonetheless, initially nodes are not synchronized and the first idle listening period could be a random interval between 0 s and 16 s. This issue alongside the retransmissions of a few packets explains the energy depletion in the first part of day 1 for nodes B and D. Node C was less affected by the initial synchronization issue in the considered simulation experiment, as its first idle listening period was ~ 1 s. Retransmissions added up further packet delay in the system. In fact, according to the protocol, they are not triggered until the next wake-up time. However, they ensured a reliable packet delivery as we did not observe any packet loss.

The pattern of V_{cap} for node B exhibits slightly worse performance than the other nodes (except node A) mainly because it acts as a relay node for both of them. The pattern of V_{cap} for node D—the leaf of the topology—exhibits better performance than node C from 6AM to 3PM every day, except day 1 because of the above mentioned synchronization and packet retransmission issues. Every day from 3PM to 8PM, the same pattern ends below the pattern of node C because of the poorer energy intake, as discussed in Section 4.3, and keeps staying below that at night.

On all the nodes, V_{cap} is constantly sufficiently above V_{cut} —reaching a minimum of 2.89 V on node B at the end of day 2—thus preventing total energy depletion. The results prove the effectiveness of the system as well as its ability to maintain energy neutral operation over time preserving energy with careful management of sleep and communication times. Nonetheless, energy analysis conducted in this simulation experiments, coupled with the power profiling analysis described earlier, could be replicated in other studies targeting bridges with different shapes, sizes, and materials. Once determined the resonant frequencies of specific members of a bridge structure that would host vibration energy-harvesting sensor nodes, the harvester model can be tuned accordingly. The charge models can then be calculated using real vibration data collected from the bridge and used to feed the COOJA simulator to execute a specific embedded application. Wake-up interval of the routing protocol, sampling period, and packet queue size should be set according to the application requirements and the available harvested power.

8.2 Validation on Real Hardware

We validate our hardware design and the sustainability of the beaconing process of RI-MAC, through a laboratory experiment on real hardware. We used 2 sensor nodes based on the MSP430FR5969 MCU and equipped with a TI CC2520 transceiver. One of the nodes is supplied by the vibration energy harvester and run the beaconing process with off period = 16 s. The second node, which is connected to a laptop, acts as a base station and is merely used to verify the transmission of beacons.

Experimental setup

Figure 15 illustrates a recap schematics of this experimental setup, while a picture is shown in Figure 16. The energy harvester is fixed onto an electromagnetic shaker driven by a waveform generator in Arbitrary Function mode and a power amplifier. The vibration profile recorded at Tamar Bridge




Figure 15. Schematics of the experimental setup.

is transferred to the waveform generator using the software Intuilink Waveform Editor. The amplifier, together with an accelerometer placed on the shaker to give feedback on the vibration acceleration, allows adjustment of the vibration amplitude.

The waveform generator in Arbitrary Function mode allows to record up to 30,000 points. Therefore, only a small part of the vibration data—which was sampled at 200 Hz, as described in Section 4.2—can be tested at one time, which are then looped over time to generate the waveform. For this experiment, we use vibration data from dataset II of location L_1 . The experiment, although not representative of the traffic conditions at all times, aims to show that the beaconing process of RI-MAC running on a sensor node can be sustained by the harvester subjected to real vibrations.




Figure 16. Experimental setup.




Figure 17. Experimental V_{cap} and V_{out} over time.

Rectification and regulation of the harvester output voltage are performed by a voltage multiplier circuit connected to the LTC3588-1 power conditioning circuit, as shown in Figure 1 of Section 4.1. The LTC3588-1 integrates a full bridge rectifier and high efficiency buck converter which provides an output voltage V_{out} regulated at 2.5 V. Finally, a 0.1-F supercapacitor is used as energy storage device and both the supercapacitor voltage V_{cap} and V_{out} are recorded using an oscilloscope.

Results

The supercapacitor is initially charged with the harvester. When the voltage is larger than $V_{on} = 4.04$ V, the buck converter is enabled and regulation is effective. The sensor node is then connected to the circuit output and both V_{cap} and V_{out} are recorded for 3 min.

Over this period of time, the energy balance is positive. In fact, the average harvested power = $541 \mu W$ and that is higher than the average power consumed by the node during the observed period, which resulted to be $10.8 \mu W$. Figure 17 shows that the supercapacitor is slightly charged over the observed period of time, as V_{cap} increases from 4.3 V to 4.5 V over the duration of the experiment. A negative voltage peak is shown every 16 s, corresponding to data transmission. The figure also traces V_{out} , which is regulated at 2.5 V by the bulk converter to power the sensor node.

9 Conclusion

We have reported of a novel comprehensive study on vibration energy-harvesting sensor networks for bridge applications, which relies on a sensor node powered by a harvester having the ability to be driven into both direct resonance and autoparametric resonance. A power profiling methodology, based on real traffic data collected from Tamar Bridge during a preliminary acquisition campaign, drove the development of a framework that allows to conduct realistic experiments before on-site deployments using the COOJA/MSPSim simulator. The framework also includes an energy-harvesting software for WSN, relying on an enhancement of the RI-MAC protocol, which ensures multihop communication under the challenging energy restrictions of the limited vibration power. We achieved energy neutral operation of a WSN in simulation and also validated the system through laboratory experiments

on real hardware.

Our study provides a methodology and potential solution to deploy a WSN powered solely by vibration harvesters, which does not entirely depend on the radio technology in use, and where multihop may still be necessary. However, we plan to explore long-range radios as they might throw the multihop out of the picture in some cases. The study also highlights the limitations of the harvested power generated by traffic-induced vibrations, which clearly shrink the applicability of vibration energy-harvesting technology for WSN to low data rate applications. We believe this is the first attempt at building a real bridge monitoring framework using a vibration energy-harvesting WSN. Finally, the design concepts of the energy-harvesting software, and the way it leverages the results of a separate power profiling methodology, lead to a framework straightforward to employ and adapt to a wide range of settings, which extends beyond bridge applications and vibration-based energy-harvesting systems.

10 Acknowledgments

This research has been funded by the EPSRC Innovation and Knowledge Centre for Smart Infrastructure and Construction project (EP/K000314/1). We thank the Tamar Bridge and Torpoint Ferry Joint Committee for allowing access to Tamar Bridge. We also thank Nicholas de Battista, Xiaomin Xu, Ki Young Koo for their assistance with sensor deployment.

11 References

- [1] Eurocode 1. Actions on structures. General actions. Thermal actions (BS EN 1991-1-5:2003). CEN, European Committee for Standardization. Bruxelles, Belgium, 2004.
- [2] Source code and relevant data, (accessed December 20, 2017). Web page—<http://www.cl.cam.ac.uk/research/srg/mobsys/IKC/>.
- [3] Network rail infrastructure limited – annual return 2016, (accessed March 28, 2017). Web page—<https://www.networkrail.co.uk/who-we-are/publications-resources/regulatory-and-licensing/annual-return/>.
- [4] M. Buettnner, G. V. Yee, E. Anderson, and R. Han. X-MAC: a short preamble MAC protocol for duty-cycled wireless sensor networks. In *SenSys*, 2006.
- [5] Y. Cao, J. Yim, Y. Zhao, and M. L. Wang. Temperature effects on cable stayed bridge using health monitoring system: a case study. *Structural Health Monitoring*, 10(5):523–537, 2010.
- [6] M. Ceriotti, L. Mottola, G. P. Picco, A. L. Murphy, S. Guna, M. Corra, M. Pozzi, D. Zonta, and P. Zanon. Monitoring heritage buildings with wireless sensor networks: The Torre Aquila deployment. In *IPSN*, 2009.
- [7] N. de Battista, J. M. Brownjohn, H. P. Tan, and K.-Y. Koo. Measuring and modelling the thermal performance of the tamar suspension bridge using a wireless sensor network. *Structure and Infrastructure Engineering*, 11(2):176–193, 2015.
- [8] A. H. Dehwah, M. Mousa, and C. G. Claudel. Lessons learned on solar powered wireless sensor network deployments in urban, desert environments. *Ad Hoc Networks*, 28:52–67, 2015.
- [9] A. Dunkels, B. Grönvall, and T. Voigt. Contiki - a lightweight and flexible operating system for tiny networked sensors. In *LCN*, 2004.
- [10] A. Dunkels, F. Österlind, N. Tsiftes, and Z. He. Software-based on-line energy estimation for sensor nodes. In *EmNets*, 2007.
- [11] S. Duquenooy, B. Al Nahas, O. Landsiedel, and T. Watteyne. Orchestra: Robust mesh networks through autonomously scheduled tsch. In *SenSys*, 2015.
- [12] J. Eriksson, F. Österlind, N. Finne, N. Tsiftes, A. Dunkels, T. Voigt, R. Sauter, and P. J. Marrón. Cooja/mspsim: Interoperability testing for wireless sensor networks. In *SIMUTools*, 2009.
- [13] X. Fafoutis, A. Di Mauro, M. D. Vithanage, and N. Dragoni. Receiver-initiated medium access control protocols for wireless sensor networks. *Computer Networks*, 76:55–74, 2015.
- [14] F. Flammini, A. Gaglione, F. Ottello, A. Pappalardo, C. Pragliola, and A. Tedesco. Towards wireless sensor networks for railway infrastructure monitoring. In *ESARS*, 2010.
- [15] O. Gnawali, R. Fonseca, K. Jamieson, D. Moss, and P. Levis. Collection tree protocol. In *SenSys*, 2009.
- [16] N. A. Hoult, P. R. A. Fidler, P. G. Hill, and C. R. Middleton. Long-term wireless structural health monitoring of the ferry road bridge. *Journal of Bridge Engineering*, 15(2):153–159, 2010.
- [17] Y. Jia, S. Du, and A. A. Seshia. Twenty-Eight Orders of Parametric Resonance in a Microelectromechanical Device for Multi-band Vibration Energy Harvesting. *Scientific Reports*, 6:1–8, 2016.
- [18] Y. Jia, J. Yan, S. Du, T. Feng, P. Fidler, C. Middleton, K. Soga, and A. A. Seshia. Real world assessment of an auto-parametric electromagnetic vibration energy harvester. *Journal of Intelligent Material Systems and Structures*, 2017.
- [19] X. Jiang, J. Polastre, and D. Culler. Perpetual environmentally powered sensor networks. In *IPSN*, 2005.
- [20] R. Karoumi. *Response of cable-stayed and suspension bridges to moving vehicles. Analysis methods and practical modelling techniques*. Phd thesis, KTH Royal Institute of Technology, 1998.
- [21] S. Kim, S. Pakzad, D. Culler, J. Demmel, G. Fenves, S. Glaser, and M. Turon. Health monitoring of civil infrastructures using wireless sensor networks. In *IPSN*, 2007.
- [22] J. P. Lynch and K. J. Loh. A summary review of wireless sensors and sensor networks for structural health monitoring. *Shock and Vibration Digest*, 38(2):91–130, 2006.
- [23] L. Mateu, C. Codrea, N. Lucas, M. Pollak, and P. Spies. Energy harvesting for wireless communication systems using thermogenerators. In *DCIS*, 2006.
- [24] J. J. McCullagh, T. V. Galchev, R. L. Peterson, R. J. M. Gordenker, Y. Zhang, J. Lynch, and K. Najafi. Long-term testing of a vibration harvesting system for the structural health monitoring of bridges. *Sensors and Actuators A: Physical*, 217:139–150, 2014.
- [25] C. R. Middleton, P. R. A. Fidler, and P. J. Vardanega. *Bridge Monitoring – A practical guide*. ICE Publishing, 2016.
- [26] S. Nabavi and L. Zhang. Portable wind energy harvesters for low-power applications: A survey. *Sensors*, 16(7), 2016.
- [27] M. Peigney and D. Siegert. Piezoelectric energy harvesting from traffic-induced bridge vibrations. *Smart Materials and Structures*, 22(9):095019:1–11, 2013.
- [28] V. Raghunathan, A. Kansal, J. Hsu, J. Friedman, and M. Srivastava. Design considerations for solar energy harvesting wireless embedded systems. In *IPSN*, 2005.
- [29] C. Renner, S. Unterschütz, V. Turau, and K. Römer. Perpetual Data Collection with Energy-Harvesting Sensor Networks. *ACM Transactions on Sensor Networks*, 11(1):12:1–12:45, 2014.
- [30] S. Roundy, P. K. Wright, and J. M. Rabaey. *Energy Scavenging for Wireless Sensor Networks - With Special Focus on Vibrations*. Kluwer Academic Publishers, 2003.
- [31] E. Sazonov, H. L. H. Li, D. Curry, and P. Pillay. Self-Powered Sensors for Monitoring of Highway Bridges. *IEEE Sensors Journal*, 9(11):1422–1429, 2009.
- [32] B. F. Spencer, Jr. and S. Cho. Wireless Smart Sensor Technology for Monitoring Civil Infrastructure: Technological Developments and Full-scale Applications. In *ASEM*, 2011.
- [33] Y. Sun, O. Gurewitz, and D. B. Johnson. RI-MAC: a receiver-initiated asynchronous duty cycle MAC protocol for dynamic traffic loads in wireless sensor networks. In *SenSys*, 2008.
- [34] Y. K. Tan and S. K. Panda. Review of energy harvesting technologies for sustainable wireless sensor network. *Sustainable wireless sensor networks*, Dec:15–43, 2010.
- [35] S. Unterschütz, C. Renner, and V. Turau. Opportunistic, receiver-initiated data-collection protocol. In *EWSN*, 2012.
- [36] T. Voigt, H. Ritter, and J. Schiller. Utilizing solar power in wireless sensor networks. In *LCN*, 2003.
- [37] N. Xu, S. Rangwala, K. K. Chintalapudi, D. Ganesan, A. Broad, R. Govindan, and D. Estrin. A Wireless Sensor Network For Structural Monitoring. In *SenSys*, 2004.
- [38] L. Yerva, B. Campbell, A. Bansal, T. Schmid, and P. Dutta. Grafting energy-harvesting leaves onto the sensornet tree. In *IPSN*, 2012.