Energy-Harvesting Wearables for Activity-Aware Services

Advances in energy-harvesting hardware have created an opportunity for realizing batteryless wearables for continuous and pervasive human activity recognition (HAR). Unfortunately, power consumption of accelerometers used in conventional HAR is relatively high compared to the amount of power that can be harvested practically, which limits energy harvesting's usefulness. Here, the authors present and evaluate an energy-harvesting wearable sensor architecture, HAR from Kinetic Energy (HARKE), that doesn't require using an accelerometer. Using off-the-shelf products, the authors demonstrate that a kinetic harvester's voltage exhibits distinguishable patterns to distinctly infer human activities accurately while consuming a fraction of the limited harvested energy.

Sara Khalifa, Mahbub Hassan, and Aruna Seneviratne
University of New South Wales and NICTA

Sajal K. Das
Missouri University of Science and Technology

Advancements in human activity recognition (HAR) enable a wide range of activity-aware services in various domains, including healthcare,[1] smart living,[2] military,[3] security, and indoor positioning[4]. For example, a system that can recognize various ambulation activities such as walking, sitting, standing, and jogging can enable healthcare authorities to continuously monitor a patient's status from a remote center. Similarly, a smartphone capable of detecting activities such as climbing stairs, riding an elevator, or moving up a ramp can infer a pedestrian's position in a complex indoor environment by matching the activities to the indoor map that shows the precise locations of stairs, elevators, and ramps.[5]

There are two fundamentally different approaches to HAR, using *infrastructure sensors* and *wearable sensors*. In the former, the sensors that can detect motion, pressure, temperature, and so on, are installed at specific locations and on furniture to detect human activity when the user visits these locations and interacts with the sensors. For example, a pressure sensor installed beneath a sofa could detect sitting activity from the pressure change whenever the user sits there. Cameras installed at specific locations can help detect user activities whenever the user comes within their vicinity. However, deployment and maintenance of infrastructure sensors are costly.

Wearable sensors, on the other hand, provide an alternative option. By placing various types of sensors on the human body, we can achieve accurate and pervasive HAR without the need for deploying significant infrastructure. For example, an accelerometer in a wristband can help detect different activities by simply collecting and analyzing the time-series acceleration data. Because wearables can continuously

1089-7801/15/$31.00

monitor user activities at all times and locations, they provide a more pervasive HAR solution compared to the infrastructure-based approach that requires the user to be within the sensing range for effective activity recognition. Consequently, wearable sensor-based HAR has recently become the focus of intense research and development,[6] producing a wealth of tools and algorithms to accurately detect human activities from the data collected by the wearables.

However, the major issue with wearable sensors is the battery life. To achieve sustained operation, we either need to instrument the wearables with large batteries or be prepared to manually replenish the batteries when they die. Neither of these options is desirable because large batteries make the wearables heavy and less convenient to wear, while frequent manual replacement might not be possible. This motivates us to explore a third option: energy-harvesting wearables for activity-aware services.

Energy harvesting or scavenging is a process of converting various forms of ambient energy sources, such as kinetic, thermal, radio frequency, and solar or light, into electrical energy, which we can use to power a small electronic device. In recent years, we've seen significant advancements in energy-harvesting hardware technology, leading to many off-the-shelf products available at low cost. This means that it's conceptually possible to replace the battery of a wearable sensor with an energy-harvesting unit to achieve perpetual sensing in many applications including HAR.

However, there's a caveat. Energy harvesting generally suffers from low power output, which could challenge the power requirement of the wearable sensor components, such as the accelerometer used for sampling human motion. A recent study has shown that the power requirement of the accelerometer ranges between 0.35 to 5 times the harvested kinetic power for detecting common human activities with high accuracy.[7] Given that the sensor will also have to turn on its radio for occasional communications with a nearby sink, the power generated from energy harvesting is clearly too small to simply port the existing accelerometer-based HAR techniques into an energy-harvesting wearable. How to achieve HAR using energy-harvesting wearables is indeed an extremely challenging problem that requires new solutions. A recent survey has revealed that although significant research has been carried out for battery-powered wearables, there exists limited literature on energy-harvesting wearables for HAR.[6]

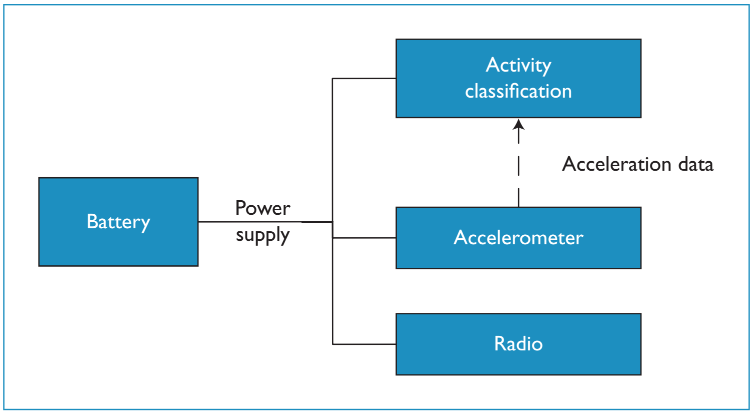

Figure 1. Block diagram of a conventional battery-based wearable sensor for human activity recognition (HAR). The battery powers three main components: an accelerometer, a classifier, and a radio.

Here, we propose and evaluate a new paradigm for energy harvesting HAR that overcomes the power limitation of energy harvesting in wearables. More specifically, our novel approach employs kinetic energy harvesting and infers human activity directly from the energy-harvesting patterns without using any accelerometer. The underlying idea lies in the fact that different human activities produce different amounts of kinetic energy that can be leveraged for activity recognition. The proposed energy-harvesting wearable sensor architecture is called Human Activity Recognition from Kinetic Energy, or HARKE.

Because we use no accelerometer in the HARKE architecture, we dedicate the harvested power to radio communication in its entirety. Using off-the-shelf products, we design a kinetic-energy-harvesting data logger, which shows that the energy-harvesting voltage switches to clearly distinguishable patterns as the user changes her activities. Our experimental results demonstrate that HARKE is as accurate as an accelerometer-based HAR system, yet it consumes only a small fraction of the limited harvested energy.

Limitations of Energy-Harvesting Wearables

Figure 1 shows a simplified block diagram of a conventional battery-powered wearable used for HAR. The battery powers three main components: an accelerometer, a classifier, and a radio. The first two make up the HAR function in which the classifier detects human activities by analyzing the features extracted from the accelerometer data. A previous detailed measurement study indicates that the average power consumption of an accelerometer running at 20 hertz (Hz) is four times as much as the average power consumption for extracting features and executing a classifier.[8] Consequently, the accelerometer is responsible for 80 percent of the total HAR power consumption.

Typically, a three-axial accelerometer is used to measure human acceleration in three dimensions. These 3D acceleration data are then used to train a classifier, which later is used to detect activities from a given sample of acceleration values. Generally, the more frequent the measurements occur, the more information is available to enable more accurate classification. The measurement frequency is called the accelerometer's sampling rate, which is measured in hertz or the number of measurements per second.

To perform a measurement, an accelerometer must be turned on for a few milliseconds. Because the accelerometer consumes power when it's active, it's turned off when it isn't measuring. Consequently, an accelerometer is continuously turned on and off, whose frequency is dictated by the sampling rate. As such, the average power consumption of an accelerometer is a linear function of the sampling rate. For example, the Harvey Weinberg's data sheet shows that an ADXL150 accelerometer consumes about 5 microwatts (μW) on average per hertz, which means that it would require 100 μW if a sampling rate of 20 Hz was required for a given activity set.[9] The required sampling rate depends on the set of activities monitored and typically ranges from 10–50 Hz.[10,11] In other words, the battery has to supply between 50–250 μW to the accelerometer. This isn't a major issue for conventional battery-powered wearables, because most batteries can supply power at a much higher rate than this.

Let's now examine the power consumption of accelerometers in the context of self-powered energy-harvesting wearables. If we simply replace the battery in Figure 1 with an energy-harvesting unit, we'll face two major problems. First, we might not be able to supply enough power to the accelerometer for accurate HAR. Using a commercial kinetic energy harvester, a study by Alberto Olivares and his colleagues shows that some activities generate only a few μW, which is much lower than what's required to sample the accelerometer at a sufficiently high rate for accurate activity classification.[12] Clearly, this will force the wearable to cut down the power to the accelerometer — that is, use a lower sampling rate and accept a lower activity classification accuracy. Second, even if the harvested power is enough to operate the accelerometer at the required sampling rate, it reduces the amount that could be accumulated in the capacitor for future radio communications. Lack of enough stored energy in the capacitor will force more aggressive duty cycling of the radio and/or more drastic reduction in the transmission power. In summary, when the power supply is limited by energy harvesting, powering the accelerometer trades off the quality of radio communication.

The HARKE Architecture

Figure 2a shows the proposed HARKE with no accelerometer. Instead, it contains a kinetic energy-harvesting (KEH) system to harvest energy from human motion. The harvested energy is then used to power the classifier, and the radio when it's turned on for communication. The KEH system typically consists of three basic components:

- a generator (transducer) to convert human motion into electrical power, typically with varying AC voltage;
- a power conditioning circuit to provide power rectification and regulation; and
- a storage element to store the harvested energy.

A storage element such as a battery or capacitor is needed to accumulate the harvested energy and supply regulated power, typically a constant DC voltage, which is suitable to power the classifier and radio communication. The regulated power isn't suitable for detecting human activities, because regulation would wipe out any potential pattern in the power signal that might be used for activity recognition. Therefore, HARKE uses the AC voltage from the transducer as an input to the classifier.

The energy savings of HARKE is directly due to the accelerometer's removal, which accounts for 80 percent of the HAR power consumption in a wearable device. Compared to the energy saved by removing the accelerometer, the power consumption for recording the AC voltage for activity classification is minimal. To continuously record AC voltage, HARKE basically needs an analog-to-digital converter (ADC) to sample the analog AC signal into digital data that we can use for feature extraction and classification. The data sheet of ADS7042 (see www.ti.com/lit/ds/symlink/ads7042.pdf), an ultra-low-power ADC from

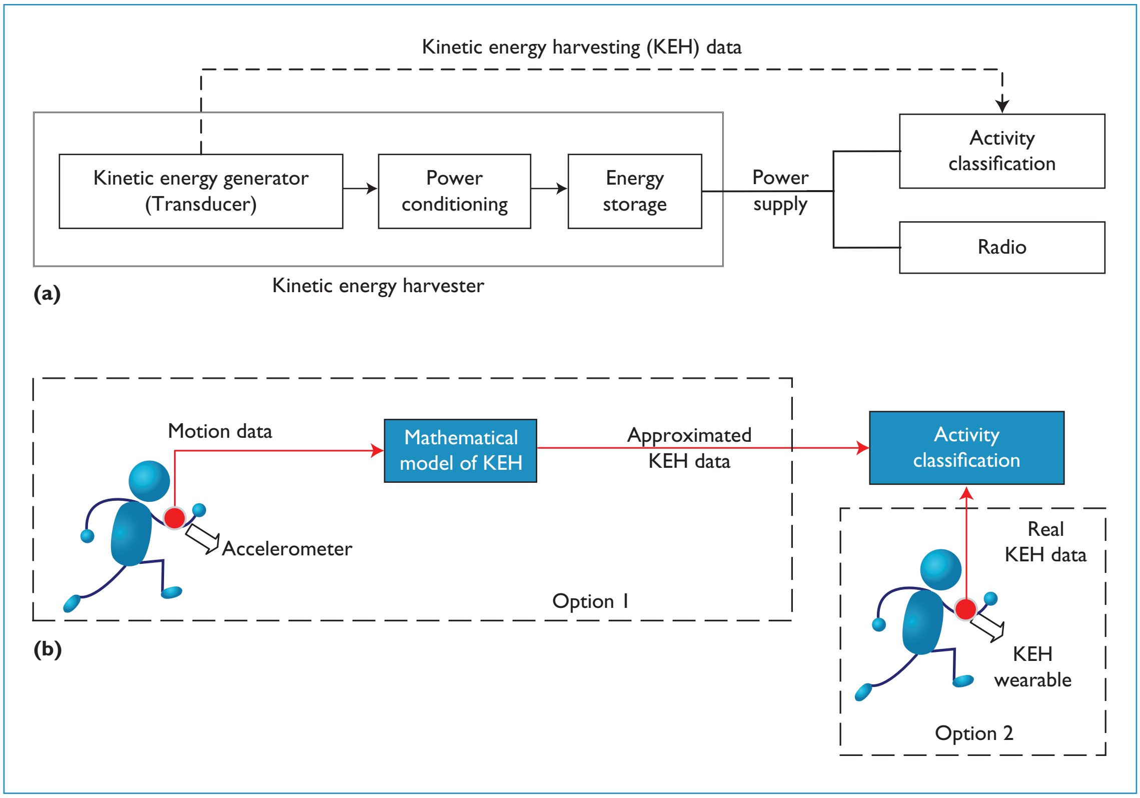

Figure 2. HAR from Kinetic Energy (HARKE) architecture. (a) HARKE block diagram. (b) Validation options.

Texas Instruments, shows that the ADC consumes approximately 1 μW per kilohertz (kHz). Comparing this to 100 μW consumed by an ADXL150 accelerometer running at 20 Hz (assuming a 5 μW/Hz power consumption), we find that 99 percent of the power that would have been consumed by the accelerometer could be saved using our proposed architecture. Given that 80 percent of HAR power consumption is due to the accelerometer, HARKE saves 72 percent of the HAR power consumption in a self-powered wearable.

It should be clear by now that the HARKE architecture can save a significant amount of the limited harvested energy in next-generation self-powered wearables. The question is how accurately can we classify human activities using the harvested energy signal?

Validation of HARKE

The basic idea of HARKE is to use the KEH data to classify common human activities. Figure 2b shows two options we can use to obtain such data. The first option is to use a mathematical model that could approximate KEH data from human motion data. We can use smartphones with incorporated accelerometers to collect human motion. This option lets us validate HARKE without the need to use special hardware. However, the generated KEH data are an approximation of the real data. The second option is to use a real device that the user could wear while performing different activities and collect real KEH data.

Validation Using KEH Data Approximation

Next, we use the first option of deriving KEH data to validate HARKE. Specifically, we use a mathematical model that relies on the most basic principle of a kinetic energy harvester,[13] namely a standard mass-spring damping system, to approximate KEH data. When the spring moves, the mechanical energy is converted into electrical energy. If the spring moves with more force, or it bounces back and forth rapidly, more energy is produced.

Given an accelerometer trace, this model lets us estimate the kinetic power signal that would have been harvested by a kinetic energy harvester. Our accelerometer traces have been collected from five

Table 1. Power-supply percentage and HAR accuracies of accelerometer-based and HARKE architectures for three different activity sets.*

Activity set	Average harvested power (μW)	Harvested power allocated to accelerometer (%)	Achievable accelerometer sampling rate (Hz)	HAR accuracy (%) Accelerometer-based	HAR accuracy (%) HARKE
Sport	104.70	100.00	20	100.00	98.15
		80.00	16	100.00	
		50.00	9	100.00	
Home monitoring	45.17	100.00	9	96.97	82.17
		80.00	7	94.81	
		50.00	3	86.58	
Indoor positioning	36.00	100.00	7	91.84	73.97
		80.00	6	89.86	
		50.00	3	77.75	

* The sport set includes standing, walking, and running activities. The home monitoring set includes standing, walking, going up or down the stairs, and vacuuming. The indoor positioning set includes standing, walking, going up or down the stairs, standing on an escalator (going up or down), and going up or down a ramp.

different subjects at 100 Hz using a Samsung Galaxy Nexus smartphone carried in the hand. The generated kinetic power traces have been used to evaluate HARKE's performance.

As is commonly used for HAR,[14] we use the K-nearest neighbor (KNN) classifier in this study. We trained the KNN classifier with 12 features (mentioned elsewhere[7]) extracted from the accelerometer traces, and only the maximum feature extracted from the approximated KEH traces. In a previous study,[7] we showed that the maximum feature provides better accuracy than simultaneously using the 12 features when KEH traces were used. In both cases, we extracted the features from 5-second windows with 50-percent overlapping between consecutive windows. Finally, we perform a 10-fold cross-validation test to obtain the accuracy.

We evaluate the performance of HARKE using three different activity sets from three different applications (sport, home monitoring, and indoor positioning), each containing between three to eight different activities to be classified. As mentioned earlier, HARKE only needs 1 μW to sample the KEH signal at 1 KHz. On the other hand, the sampling rate of the accelerometer, and hence its HAR accuracy, will depend on the power allocated to it. Table 1 compares HAR accuracies of accelerometer-based and HARKE architectures in the last two columns and shows the percentage of harvested power allocated to the accelerometer in the third column.

We can see that HAR accuracies for the accelerometer-based architecture are generally higher when higher power is allocated (for higher sampling rates). HARKE accuracies varied between 74 to 98 percent, depending on the activity set. The average accuracy over all activity sets is 85 percent, which is within 12 percent of what could be achieved with an accelerometer without any power constraints (using 100 percent of the available harvested power). However, if we consider the realistic scenario of dividing the available harvested power between the accelerometer sensor and radio communication, this difference will be reduced. For example, when 50 percent of the harvested power is allocated to the accelerometer, the average accuracy of HARKE is within only 3 percent of what could be achieved with an accelerometer with the achievable sampling rate.

Validation Using Real KEH Data

For the ultimate validation of HARKE, we built a data logger, to collect real KEH data. Our objective is to investigate whether the kinetic energy harvested by a real harvester actually contains information about human activity — and if it does, how does the accuracy of HAR from KEH data compare against conventional HAR based on accelerometers.

The prevalent commercial kinetic energy harvesters are based on the piezoelectric and electromagnetic transduction mechanisms (see Table 2), but piezoelectric transducers are the most favourable due to their simplicity and compatibility

Table 2. Commercial kinetic (vibration) energy harvesters (VEHs).

Manufacturer (country)	Product	Material	Dimensions (in) $L \times W \times H$	Weight (grams)	Output (in voltage)
Perpetuum (UK)	PMG FSH	Electromagnetic	3.4×2.6	1075	DC (5V and 8V)
Ferro Solutions (US)	VEH 460	Electromagnetic	—	430	DC (3.3V)
LORD MicroStrain (US)	PVEH	Piezoelectric	1.87×1.75	185	DC (3.2V)
	MVEH	Electromagnetic	2.25×2.56	216	DC (3.2V)
MicroGen (US)	BoLT PZEH	Piezoelectric	$1.18 \times 1.04 \times 0.6$	10	DC (3.3 V)
MIDÉ (US)	Volture V25W	Piezoelectric	$2.00 \times 1.50 \times 0.0$	8	AC
PI Ceramic GmbH (Germany)	P-876.A11 DuraAct	Piezoelectric	$2.4 \times 1.38 \times 0.02$	—	AC
Smart Material (US)	MFC M2503-P1	Piezoelectric	$1.81 \times 0.93 \times 0.0$	—	AC
OMRON and Holst Centre/imec (Belgium)	Still under testing	Electrostatic	1.96×2.36	15.4	DC

with microelectromechanical systems (MEMS).[15] Most commercial harvesters are available as a packaged system, including the transducer, power conditioning circuit, and local storage providing a regulated DC voltage to power multisensor nodes, controllers, peripherals, memory, and so on. However, we can't access the instantaneous AC voltage in such packaged products. Only a few companies, such as MIDÉ, make the AC voltage accessible by offering modular components for transduction, power conversion, and storage.

Our data logger includes a product called Volture from MIDÉ (see www.mide.com), which implements only the transducer providing AC voltage as its output, and also a three-axis accelerometer (MMA7361LC) for comparison purposes. We used an Arduino Uno as a microcontroller device for sampling the data from both the Volture and the accelerometer. We used a sampling rate of 1 kHz for data collection. We saved the sampled data on an 8-Gbyte microSD card, which we equipped to the Arduino using microSD shield. A 9V battery was used to power the Arduino. The data logger also includes two switches, one to switch on/off the device and the other to control the start and stop of data logging. Figures 3a and 3b show the hardware platform and the internal appearance of the data logger.

Ten subjects — four male and six female, with ages between 26 to 35 years — volunteered to participate in this research study. We asked the subjects to hold the data logger in either their left or right hand and perform three activities: standing, walking, and running. All subjects performed walking and running with their natural speed (there was no special speed requirement). To avoid mislabeling, we used a switch to start and stop data collection at the beginning and end of each activity. We asked subjects to stop and wait a few seconds after an activity and before starting the next activity. We labeled the data collected between the start and stop times of an activity with that activity's name. Each subject provided between 25 and 35 seconds of data for each activity. Figure 3c shows the user's output patterns of the accelerometer and Volture. We can clearly see that the AC output of Volture changes patterns whenever the user changes his or her activities. This is a clear indication that the output of kinetic energy harvesters contains information about human activity.

Next, we investigate and compare HARKE's accuracy against the conventional accelerometer-based HAR. No noise filtering has been applied on the data. We consider a technique of *window overlapping* for feature extraction. In this technique, the data traces are subdivided into smaller subsets or windows. As previously mentioned, we chose windows of 5 seconds with 50 percent of overlap between consecutive windows to reduce information loss at the edges of the window. We extracted two features — mean and standard deviation — from the consecutive windows.

In our work, we have two datasets — one for accelerometer data and one for KEH data from Volture. The accelerometer generates three time series along the *x*-, *y*-, and *z*-axes. We compute the features for each of the three axes, giving a total of six extracted features for each

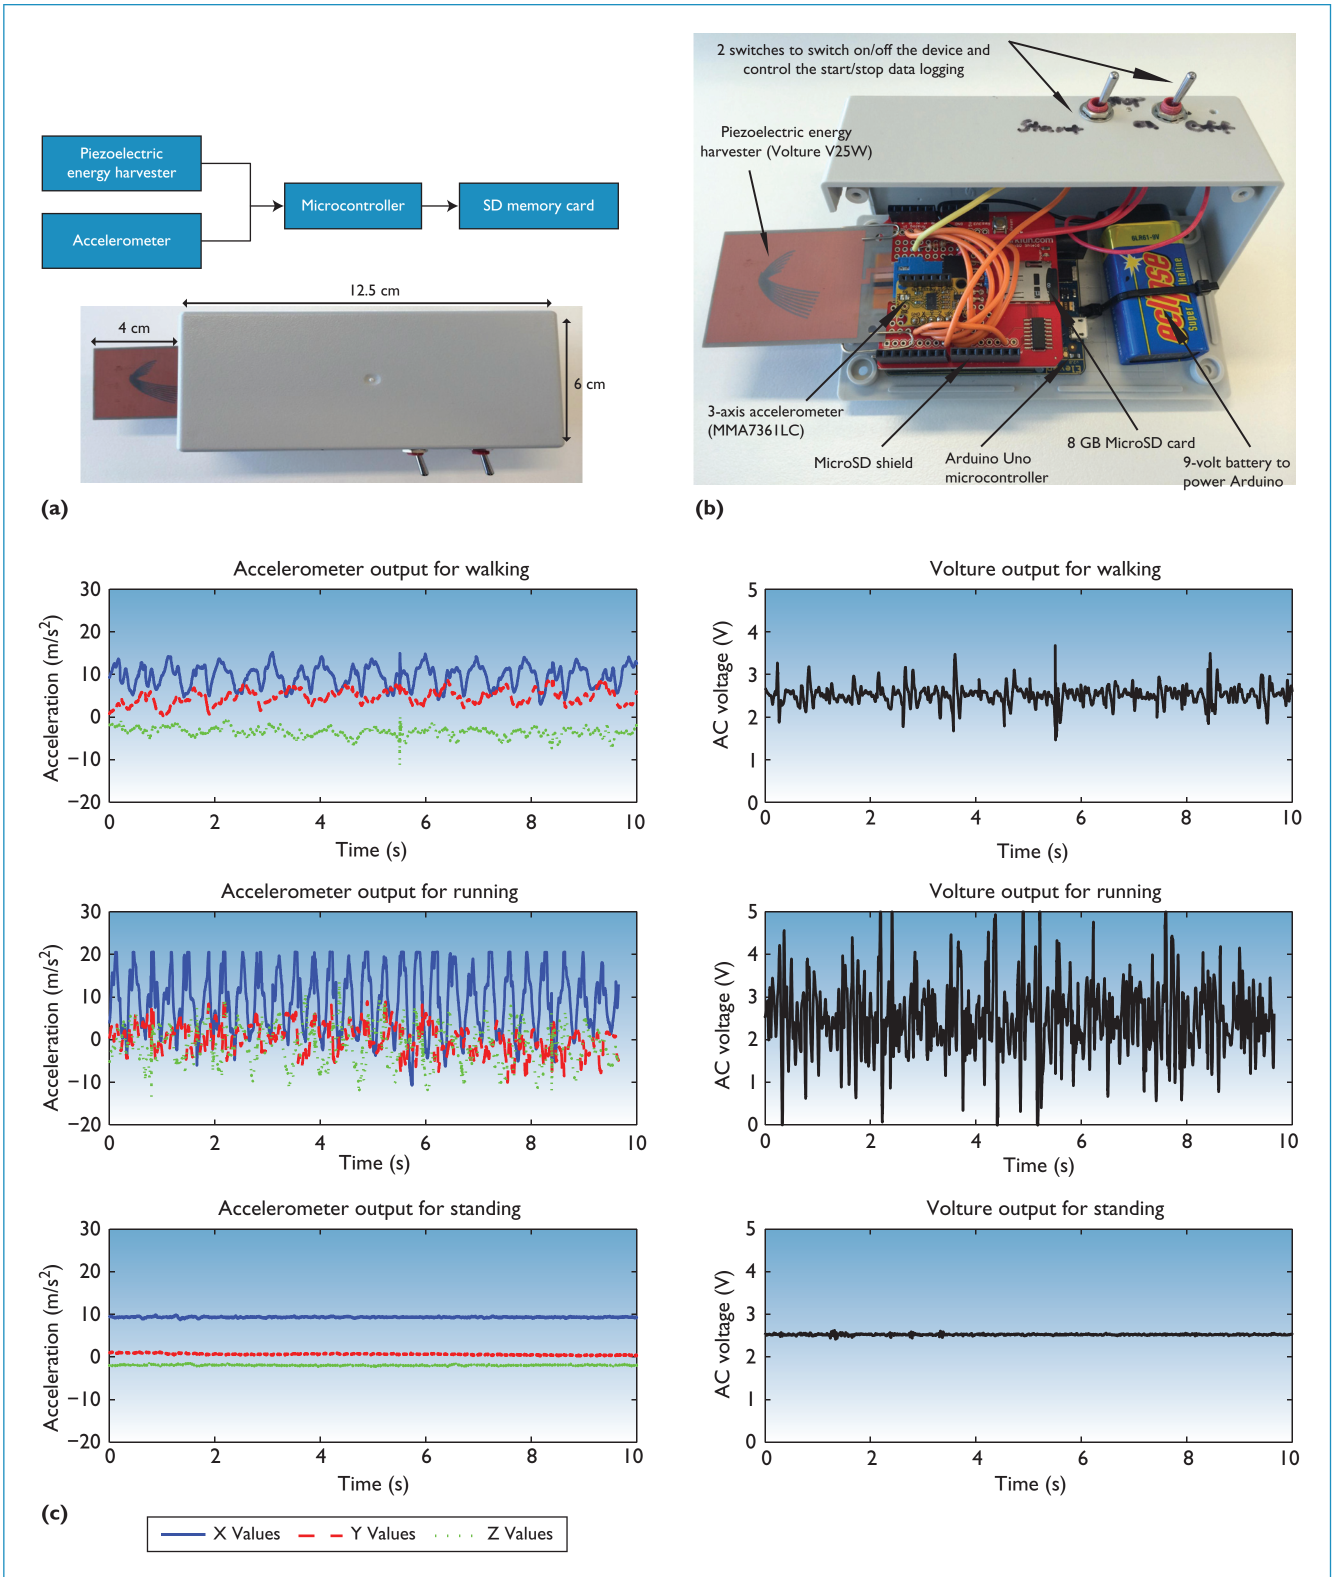

Figure 3. HARKE proof-of-concept experimentation. (a) Block diagram of kinetic energy harvesting and the accelerometer data logger (up) and the external appearance of the data logger (down). (b) The internal appearance of the data logger. (c) The accelerometer's output pattern (left) and Volture (right).

Table 3. HAR accuracies (given as percentages) of accelerometer-based HAR and HARKE architectures for six different classifiers.

Classifier	Accelerometer-based HAR (3-axes acceleration)	HARKE (1-axis AC voltage)
K-Nearest Neighbor	100.00	99.25
Multilayer Perceptron	100.00	99.25
Logistic Regression	99.25	98.50
Decision Trees (J48)	96.24	97.75
Naïve Bayes	96.24	97.74
Bayes Net	98.50	99.25

accelerometer trace. On the other hand, the Volture generates one series of AC voltage, giving only two extracted features for each Volture trace.

We chose five classifiers to evaluate the recognition accuracy of HARKE and compare it with the accelerometer-based HAR system. As previously mentioned we performed a 10-fold cross-validation test to obtain the accuracy. Table 3 shows the classification accuracies of accelerometer-based and HARKE architectures. These results confirm that although HARKE consumes 72 percent less energy compared to the conventional accelerometer-based HAR, it can classify human activities as accurately as accelerometer-based HAR.

Our experimental study shows that power consumption by the accelerometer is a major roadblock for realizing HAR in self-powered energy-harvesting wearables. By eliminating the accelerometer, however, HARKE consumes only a small fraction of energy compared to the conventional accelerometer-based HAR. Because HARKE can detect human activities as accurately as conventional accelerometer-based HAR, we believe that this finding will open the door for a new direction of research and development in realizing self-powered devices for the future Internet of Things.

A natural continuation of this current work is to quantify HARKE's performance under a wider range of activities and using a larger-scale dataset. Building and testing a complete self-powered HARKE prototype is also an important next step.

References

1. O. Chipara et al., "Reliable Clinical Monitoring Using Wireless Sensor Networks: Experiences in a Step-Down Hospital Unit," *Proc. 8th ACM Conf. Embedded Networked Sensor Systems*, 2010, pp 155–168.
2. B. Tessendorf et al., "Recognition of Hearing Needs from Body and Eye Movements to Improve Hearing Instruments," *Proc. Int'l Conf. Pervasive Computing*, 2011, pp. 314–333.
3. T. Wyss and U. Mder, "Recognition of Military-Specific Physical Activities with Body-Fixed Sensors," *Military Medicine*, vol. 175, no. 11, 2010, pp. 858–864.
4. K. Altun and B. Barshan, "Pedestrian Dead Reckoning Employing Simultaneous Activity Recognition Cues," *Measurement Science and Technology*, vol. 23, no. 2, 2012, pp. 1–20.
5. S. Khalifa, M. Hassan, and A. Seneviratne, "Adaptive Pedestrian Activity Classification for Indoor Dead Reckoning Systems," *Proc. Int'l Conf. Indoor Positioning and Indoor Navigation*, 2013; doi:10.1109/IPIN.2013.6817868.
6. O. Lara and M. Labrador, "A Survey on Human Activity Recognition Using Wearable Sensors," *IEEE Comm. Surveys and Tutorials*, vol. 15, no. 3, pp. 1192–1209, 2013.
7. S. Khalifa, M. Hassan, and A. Seneviratne, "Pervasive Self-Powered Human Activity Recognition without the Accelerometer," *Proc. Int'l Conf. Pervasive Computing and Comm.*, 2015, pp. 79–86.
8. X. Qi et al., "Adasense: Adapting Sampling Rates for Activity Recognition in Body Sensor Networks," *Proc. IEEE 19th Real-Time and Embedded Technology and Applications Symp.*, 2013, pp. 163–172.
9. H. Weinberg, *Minimizing Power Consumption of iMEMS Accelerometers*, AN-601 application note, Analog Devices, 2002; www.analog.com/media/en/technical-documentation/application-notes/5935151853362884599AN601.pdf.
10. N. Ravi et al., "Activity Recognition from Accelerometer Data," *Proc. 17th Conf. Innovative Applications of Artificial Intelligence*, 2005, pp. 1541–1546.
11. A.M. Khan, Y.-K. Lee, and T.-S. Kim, "Accelerometer Signal-Based Human Activity Recognition Using Augmented Autoregressive Model Co-Efficients and Artificial Neural Nets," *Proc. 30th Ann. Int'l Conf. IEEE Eng. Medicine and Biology Soc.*, 2008, pp. 5172–5175.
12. A. Olivares et al., "A Study of Vibration-Based Energy Harvesting in Activities of Daily Living," *Proc. 4th Int'l Conf. Pervasive Computing Technologies for Healthcare*, 2010;

http://dx.doi.org/10.4108/ICST.PERVASIVEHEALTH 2010.8808.

13. M. Gorlatova et al., “Movers and Shakers: Kinetic Energy Harvesting for the Internet of Things,” *ACM Sigmetrics Performance Evaluation Rev.*, vol. 42, no. 1, 2014, pp. 407–419.
14. K. Jia and D.-Y. Yeung, “Human Action Recognition Using Local Spatio-Temporal Discriminant Embedding,” *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, 2008; doi:10.1109/CVPR.2008.4587732.
15. E. Lefeuvre et al., “A Comparison between Several Vibration-Powered Piezoelectric Generators for Standalone Systems,” *Sensors and Actuators*, vol. 126, no. 2, 2006, pp. 405–416.

Sara Khalifa is a PhD student in the School of Computer Science and Engineering at the University of New South Wales, Australia. Her research interests include pattern recognition and energy-harvesting fields, with a focus on human activity recognition. Khalifa has an MSc in computer science from Zagazig University, Egypt. She’s a recipient of the Australian Postgraduate Award, NICTA Research Project Award, and Engineering Research Excellence Award. She’s also an IEEE student member. Contact her at sarak@cse.unsw.edu.au.

Mahbub Hassan is a full professor in the School of Computer Science and Engineering at the University of New South Wales, Australia. His research interests include nanoscale wireless networking, self-powered communication, novel uses of energy harvesting in small devices, and opportunistic networking. Hassan has PhD in computer science from Monash University, Australia. He’s a Distinguished Lecturer of the IEEE Communications Society, and an editor for IEEE Communications Surveys and Tutorials. More information about him is available at www.cse.unsw.edu.au/~mahbub. Contact him at mahbub@cse.unsw.edu.au.

Aruna Seneviratne is the director of the Australian Technology Park Laboratory and the leader of the Networks Research Group at NICTA. He holds the Mahanakorn Chair of Telecommunication in the School of Electrical Engineering and Telecommunications at the University of New South Wales. His research interests include mobile data communication systems with an emphasis on the development of mechanisms for managing the quality of service in wireless Internet environments. Seneviratne has a PhD in electrical engineering from the University of Bath. Contact him at a.seneviratne@unsw.edu.au.

Sajal K. Das is the Daniel St. Clair Endowed Chair Professor of Computer Science at the Missouri University of Science and Technology. His research interests include wireless sensor networks, smart healthcare, cyber-physical systems, mobile and pervasive computing, security and privacy, and social networks. Das has a PhD in computer science from the University of Central Florida. Contact him at sdas@mst.edu.