CONTEXT-AWARE NETWORKING AND COMMUNICATIONS

On Systems Generating Context Triggers through Energy Harvesting

Vijay S. Rao, S. N. Akshay Uttama Nambi, R. Venkatesha Prasad, and I. Niemegeers

ABSTRACT

Context awareness is one of the building blockS of smart applications that constitutes smart spaces. With the emergence of cyber physical systems, it is now possible to create spaces that are truly adaptable and smart. In these spaces, contextual parameters are captured by many wireless sensor nodes. This collected data is processed to understand the context in real time. Since a large number of sensors are deployed, processing all the data is a big task. Moreover, since the sensors are powered by batteries, they have a limited lifetime. Sensing only when there is a context event can save energy as well as reduce data processing. To make sensor nodes operate perpetually, ambient energy harvesters can be used. Typically, the energy harvesting source of a sensor is related to the physical parameters the sensor is measuring. We propose exploiting this property to develop a context-event triggering mechanism in this article. We also adapt the context-aware application framework for incorporating context-event triggers by harvesters. A Smart-M3-based architecture is also proposed. Through a real-world use case, we illustrate significant energy savings and reduced data processing in our proposed approach.

INTRODUCTION

With rapid advancements in embedded systems and wireless technologies, the vision of Mark Weiser is becoming a reality. In his momentous work [1], he envisioned "ubiquitous computing," that is, *personal* computers integrate seamlessly into a person's environment and enrich his/her everyday life by automating many routine tasks and providing information relevant to the *context*. Context is any information that can characterize the situation of a user or an entity in general [2]. Employing contextual information in applications to enrich user experience has led to powerful ideas such as smart spaces [3]. The smart space is a paradigm based on ubiquitous computing, where environments are impregnated with embedded devices to capture the context and adapt the ambiance around the user accordingly to improve his/her experience. Cyber physical systems enable connecting "things" (physical objects) and controlling them [4]. Hence, smart spaces move on to become smart cyber physical (SCy-Phy) spaces. SCy-Phy spaces not only adapt to the environment and/or respond according to the context of the user, but also enhance safety and quality of experience for the user. Obviously, smart spaces are a subset of SCy-Phy spaces. An overview of SCy-Phy spaces is shown in Fig. 1, where both information and communications technology (ICT) objects such as sensors and non-ICT objects such as appliances can be monitored and controlled based on context. SCy-Phy spaces integrate smartness across various domains such as homes, offices, buildings, transportation, logistics, and cities.

The core idea of SCy-Phy spaces[1] is to gather various contextual information (location of the user, user movements, temperature of the room, etc.) and then *act* appropriately based on the derived context. The most popular approach to collect these data is by using wireless sensor networks (WSNs). Nodes in WSNs are low-power, battery-operated, tiny embedded devices that have a sensor(s) and a radio transceiver. Typical WSNs are ad hoc networks where a multihop approach is used for communication between the nodes and a sink to conserve energy. The nodes report the sensor data periodically to the sink. The sink then processes the data received from all the nodes to determine the context and/or change in the situation.

This classical approach of periodic sampling has some problems:
- Sensor nodes drain their batteries quickly.
- A sink needs to process huge amounts of sensor data generated by all the nodes.

Many research efforts have targeted energy conservation on sensor nodes and lifetime extension of the network. Recently, there has been tremendous growth in energy harvesting technologies targeting perennial lifetime for WSNs [5, 6]. In energy harvesting WSNs (EH-WSNs),[2] nodes scavenge energy from ambient sources like light, heat, water flow, and vibrations. While harvesting alleviates the first problem of the classical approach to an extent, it does not solve the second one. To address the second problem, the

Vijay S. Rao, S. N. Akshay Uttama Nambi, R. Venkatesha Prasad, and I. G. M. M. Niemegeers are with the Embedded Software group in the Faculty of Electrical Engineering, Mathematics and Computer Science in Delft.

[1] In this work, we confine SCy-Phy spaces to indoor environments such as home or office.

[2] We assume that the nodes run only on the harvested energies and do not have secondary energy sources like batteries.

0163-6804/14/$25.00 © 2014 IEEE

sink should be notified only if a context changes (i.e., a context-event triggering mechanism is the need of the hour).

In EH-WSNs, if a node begins to harvest energy, it indicates that there is a change in the environment. In many cases, the energy harvesting source for a node will be related to the physical parameter that the sensor is measuring. For instance, a light sensor should have a photovoltaic harvester that harvests energy, which is directly proportional to the intensity of light. Therefore, in this article, we propose to exploit this property of energy harvesters to detect change in the context. We further design a context-event triggering framework and architecture based on Smart-M3 [7] for SCy-Phy spaces. This article demonstrates the above concepts and the framework with a real-world use case.

The remainder of this article is organized as follows. First, we briefly describe context-aware systems. We then describe a context-event triggering mechanism through energy harvesting. A novel framework for a context-event triggering system is discussed. Next, we demonstrate our proposal with a use case along with energy savings compared to other approaches. We describe the challenges to be addressed for a robust system. Finally, we conclude the article.

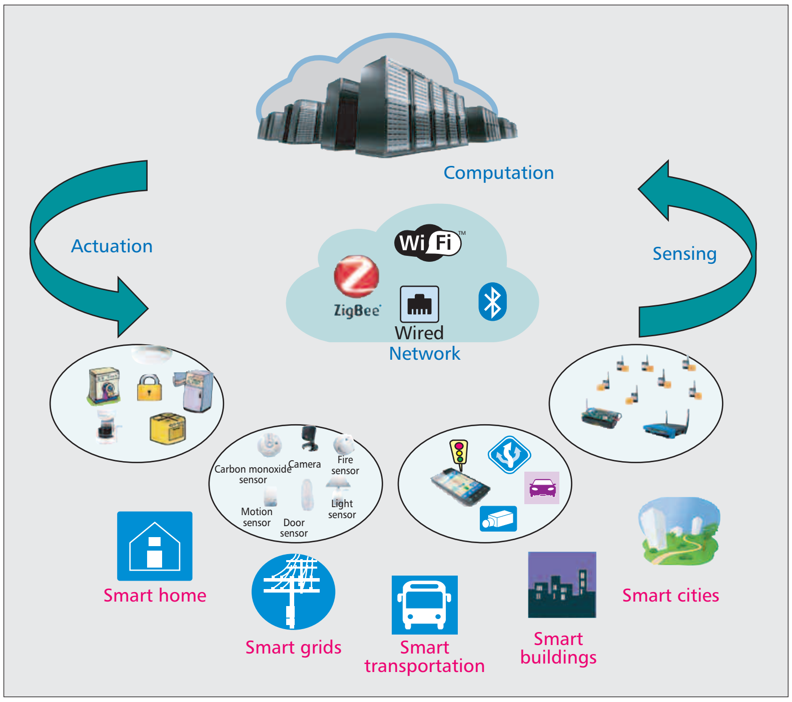

Figure 1. An overview of a smart cyber physical (SCy-Phy) space.

Context and Context Awareness in SCy-Phy Spaces

The widely accepted definition of context is "any information that can be used to characterize the situation of an entity. An entity can be a person, place or an object that is considered relevant to the interaction between the user and the application." [2] Since any information can be contextually relevant according to the definition, it is a huge task to process all data in real time. Hence, to reduce this complexity, context sources[3] are grouped into dimensions [8], which are then weighted to pick the most relevant sources. The major dimensions of context are listed below.

Ambient dimension: The set of sources that are in the proximity of the user. This also includes real-time raw sensor data about the user and his/her ambiance. Examples include room temperature and user location.

User and social dimension: The set of sources that characterize the user's preferences. Here the user's social graph members and their preferences may also play a role in determining the context. Examples include a user's shopping information and recommendations.

Derived dimension: The information is obtained from external sources — the web, calendars, weather, traffic information, and so on — and may play a role in identifying the current situation of the user.

Of the different dimensions, the ambient dimension plays the most significant role since it is close to the user and can capture the user's action and environment in indoor SCy-Phy spaces. In SCy-Phy spaces, a large number of sensors are deployed to capture the user and ambiance information.

A context-aware system [9] determines *why* the situation is occurring based on the contextual information. The system can then either adapt the environment or react to the situation. For example, in an indoor SCy-Phy environment, the temperature of a room in which the user is present can be adapted according to his/her preference. If s/he moves to another room, the context-aware system reacts to this action by adjusting the temperature in the new location.

A generalized framework for context-aware systems [10] in SCy-Phy space is shown in Fig. 2. The architecture consists of three layers: the sensing, modeling, and application layers. Each module in the architecture is briefly explained below.

Context providers furnish data about the contextual parameters. These can be sensors, user preferences, or external entities that provide data such as temperature, humidity, light, RFID, location, shopping preferences, social preferences, calendars, and weather and traffic information.

A context interpreter harmonizes the data given by heterogeneous context providers. This is required since the data will be of diverse data types, formats, and values.

A context reasoner infers the big picture from the context information. A context reasoner considers the relevance and quality of the contextual information gathered. A reasoner employs inference and rule-based mechanisms. Thus, context reasoners can derive new high-level contextual information by employing rules based on the current context information. Context reasoners can use additional resources (location, time, user information, etc.) for deriving this new high-level contextual information.

A context modeler and storage is used to represent contextual information in a machine understandable format. Context information can be modeled using a variety of approaches such

[3] *Any source that captures information about a user and can provide this information to the sink for deriving context is referred to as a context source.*

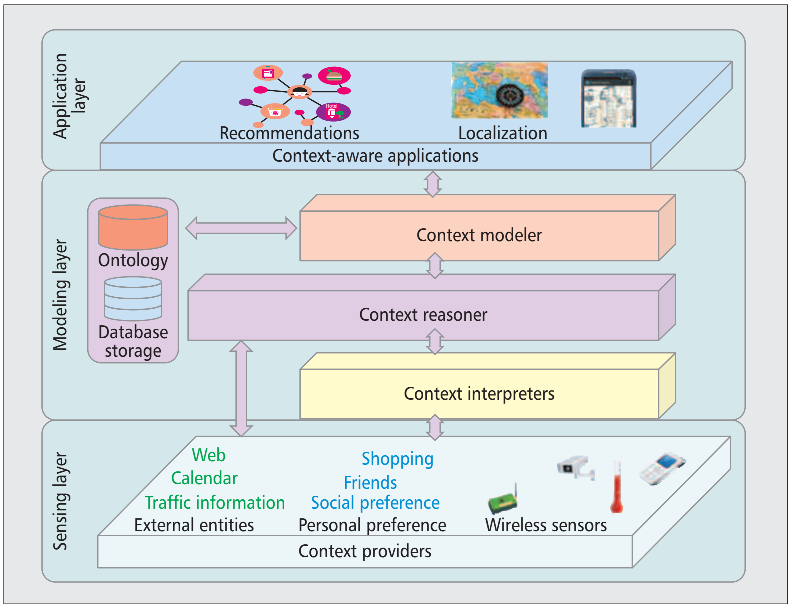

Figure 2. Generalized framework for context-aware systems.

as key-value, object-oriented, logic-based, and ontology-based models. In a key-value pair model, the value of context information is provided to the application as an environment variable. In object-oriented models, context is represented as an atomic value from a flat information model. Ontology-based models support heterogeneity and interoperability in context representation. Ontology-based context modeling has several advantages: flexibility, extensability, being generic, and supporting easy querying and retrieval of data that are crucial for context-aware systems.

Finally, the context model is used by the applications to adapt and/or respond to the user's situation. It is apparent that a huge amount of data is gathered and processed, mostly because of periodic data collection. Therefore, in the following section, we propose a method that triggers a context change event to aid in spotting only the moments when data gathering should be done.

Context-Triggered Sensing

Energy Harvesting Wireless Sensor Networks

Of context providers, WSNs play a crucial role [11]. The ambiance dimension of the context is typically determined by analyzing periodic data from these sensors. Hence, a large number of sensors are deployed to monitor various parameters.

The sensors are usually powered by batteries and are expected to run for long periods of time. To conserve energy, several strategies such as data aggregation and event detection at the node and network levels are adopted, as described in [12]. However, even with sophisticated energy conservation techniques in these sensors, their batteries burn out quickly. Frequent battery replacement is labor intensive in some cases; in many other situations, battery replacement is impractical due to harsh physical or deployment conditions. The other associated problems of batteries, such as increased size for increased capacity and higher leakage, make them unattractive.

A promising approach for perpetual network operations is to harvest energy from ambient sources, such as light, radio frequency, thermal, wind, water, salinity gradients, and motion/mechanical movements [5]. The sensor nodes are completely dependent on the harvested energy. The characteristics of nodes that use energy harvesting differ from those with conventional power sources. Ambient energy sources do not provide constant power, and the harvested energy from these sources varies drastically over location and time. The harvested energy is sometimes very low and sometimes in excess of the storage capacity of the nodes. Table 1 summarizes the sources and power that can be harvested from these sources.

While harvesting energy alleviates the lifetime problem, it can also detect a change in context. This is described in the following section.

Context-Event Triggers through Harvesters

We propose to use harvesters as context-event triggers since the energy harvesting source for a sensor will most often be related to the physical parameters that the sensor measures. Here the harvester behaves like a sensor. For example, a shoe insert measures the distance walked/run and thereby the calories burned by the user. A shoe insert energy harvester gets triggered only when there is user movement, and the total energy harvested is directly related to the calories burned. Therefore, instead of periodically sensing for user activity even when s/he is stationary, the sensor can now only be activated when the harvester generates energy. At this point, a context event is triggered and reported to the sink.

Another example is when a luminosity measuring sensor has a PV harvester. When there is change in light intensity, the energy generated by the PV harvester changes proportionally. This change can be used to trigger context-event notification and initiation of other sensors for further monitoring. Many such examples can be easily envisioned in the SCy-Phy spaces, including smart homes, smart transportation, smart cities, advanced energy metering infrastructure [14], and so on. Table 2 summarizes the possible contexts generated by different harvesters in indoor SCy-Phy spaces.

Therefore, the harvester can be used to derive initial high-level contextual information. Unlike the classical method of periodic sensing to detect context changes, now we can rely on harvesters to detect context changes and use sensors to precisely monitor all the changes occurring from then on. If need be, the sensors can operate on a low sensing frequency until a context change is detected. Consequently, all the sensor nodes deployed can save energy.[4] Thus, energy harvesting sensor nodes not only support perpetual operation of nodes, but also assist in deriving contextual information.

[4] *Energy saving is also required in EH-WSNs to mitigate the variability in harvested energy and enable operation under low-energy conditions.*

Energy source	Harvesting technology	Amount of energy harvested	Example usage in indoor SCy-Phy space
Sunlight	Solar (photo-voltaic) cells	100 mW/cm^2 (direct sun)	Sensors near windows
Ambient light	Solar (photo-voltaic) cells	100 $\mu W/cm^2$ (illuminated office)	Sensors in rooms
Wind	Anemometer	1200 mWh/day (@5 m/s wind)	In AC ducts
Water	Water (hydro) turbines	1 W/ltr	In kitchens, toilets, and showers
Thermo-electric	Temperature gradient (Seebeck effect)	60 $\mu W/cm^2$ (@40°C gradient)	In heaters and on human bodies
Vibrations	Piezoelectric	0.2 mW/cm^2	In wheelchairs, appliances like washing machines, refrigerators
Push buttons	Magnetic coils	50 µJ/N	Wireless switches, remote controls
Shoe inserts	Microgenerators	300 $\mu W/cm^3$	In shoes

Table 1. Characterization of energy sources [5].

Energy harvester	Events detected	Contextual information derived
Photo-voltaic (indoor)	Lights turned ON/OFF	Occupancy in the room; life sign
Hydro	User presence; usage of water	User location (in kitchen, shower etc.,); water usage patterns and activities, advanced metering
Thermo-electric (on body sensor)	Change in body parameters	Indicates probable change in other body parameters, advanced metering
Piezoelectric	Appliance usage; movement on wheelchairs	Indicates life sign for elderly using wheelchairs; also, appliances with this type of harvester detect usage patterns
Wireless switches	Lights ON/OFF	User location (range)
Shoe inserts	Movement	Life sign; indicates energy spent and change in body parameters

Table 2. Events detected and contextual information from harvesters.

A Novel Framework for Context-Event Triggered Systems

We propose a novel framework for context-event triggered systems using energy harvesters. We then describe the architecture using Smart-M3 adapted for our framework.

Framework

With our proposed context-event triggering system through energy harvesters, we notice that energy harvesters can directly provide high-level contextual information. At a basic level, the energy harvesters replace the batteries in WSNs. Therefore, they act as any other sensor node in collecting contextual data. For example, a light sensor can be powered through a PV harvester. The context reasoner has the responsibility of inferring high-level context from the sensor data. For example, consider the case of a user walking into his/her bedroom and turning on the light. In the classical approach, WSNs periodically send all the luminosity values to the sink. At the point when luminosity values jump, the reasoner determines the high-level context: the user walked into the bedroom at that instant. In our proposed method, when the PV harvester generates more energy, the sensor node concludes that a light is "on" in the bedroom. The sensor can then directly report the reasoned context to the sink. Thus, the reasoner is not required to collect and process huge amounts of data. The context information generated by the harvester can then be sent to the context modeler. This calls for a cross-layer design as compared to the traditional framework. Figure 3a shows the framework adapted for context-event trigger systems.

Adapted Smart-M3 Architecture

In our framework, we adapt the Smart-M3 [7] architecture that was proposed for semantic information interoperability. Smart-M3 (multipart, multidevice, and multivendor) consists of two important entities: knowledge processors (KPs) and semantic information brokers (SIBs). KPs are entities that either produce or consume information. An SIB is an entity that maintains high-level contextual information. There can be

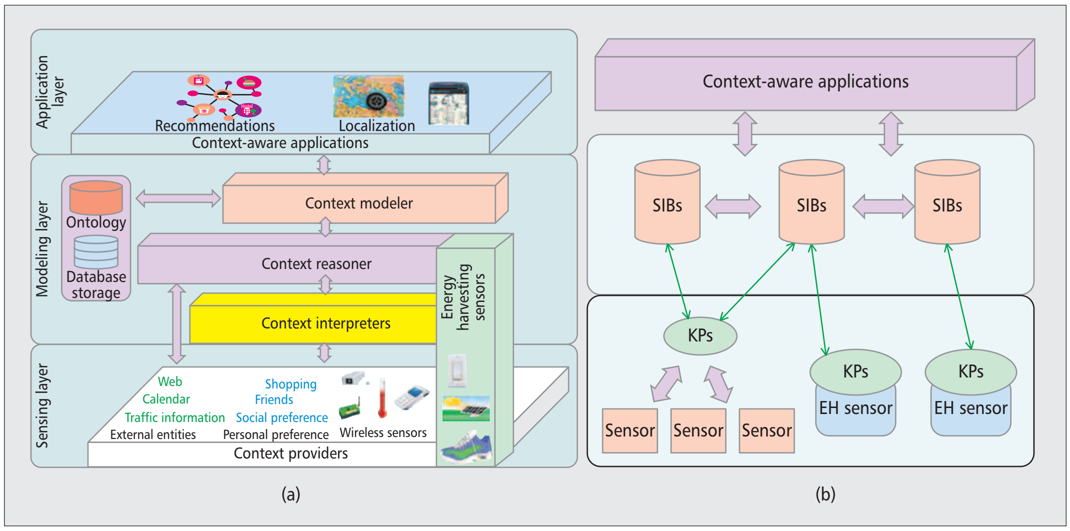

Figure 3. Proposed framework and architecture: a) framework for context-event trigger systems; b) adapted Smart-M3 architecture with energy harvesting sensors.

more than one SIB, where information is distributed among various SIBs. KPs communicate with the SIBs using access methods that are defined by the Smart Space Access Protocol (SSAP). SSAP includes access methods for various communication technologies such as WiFi, Bluetooth, Zigbee, and NFC. SSAP provides a set of primitives that enable KPs to *join* an SIB, *leave* an SIB, and *access* the information (insert, retrieve, or query) in the SIB.

Typically with Smart-M3 architecture for context-aware applications, the KPs reside in the sink node. The KPs perform the interpretation and reasoning based on the aggregated sensor data. The SIBs store the high-level contextual information. In our adapted Smart-M3 architecture, nodes with energy harvesters that directly generate high-level contextual information can themselves act as KPs. Therefore, this information is directly sent to SIBs, as shown in Fig. 3b. Applications/services subscribe to KPs to be notified when there is a change in data produced. Through this subscription, other sensors in the ambience can be triggered to monitor the changes in context. Consequently, SIBs are updated and applications in that space adapt based on this new information.

Use Case: Indoor SCy-Phy Spaces

In this section, we describe a simple use case where we demonstrate how our proposed method and framework can be employed in real-life scenarios. We then show the benefits of our proposal in terms of energy savings and data processing.

Bob is an early adopter of SCy-Phy spaces. In his house, he has a fitness room for physical exercise activities. He has converted his fitness room into a SCy-Phy space. This room is equipped with wireless switches, motion, temperature, humidity, and light sensors, and an Internet access point that also acts as a sink to the wireless sensors. He bought a wireless pulse rate monitoring device that can be worn on his wrist or shoe. All the wireless devices, except the pulse rate monitoring device, are equipped with suitable energy harvesters. The pulse rate monitor is battery operated. The shoe has a shoe insert energy harvester, and the wireless switch has a linear motion harvester. The motion, temperature, humidity, and light sensors have PV harvesters equipped as the energy source.

Consider the following scenario: Bob enters his fitness room and presses the light switch. Immediately, the wireless switch harvests energy and sends a notification to the sink, which then turns the light on. The sink also notes that Bob is in the room, and the motion sensor is turned on. The change in light conditions triggers the temperature, humidity, and light sensors' PV harvester. They conclude a person's presence in the ambience and start recording values. Bob then wears his pulse rate monitor and the shoe. Since he likes springboard jumping, he begins his routine exercise of jumping. The moment he starts jumping, the shoe insert detects a significant change in harvested energy. It immediately notes that Bob is performing a physical activity and sends this contextual information to the sink. The shoe needs to trigger body sensors to monitor the changes. Thus, it broadcasts a notification to body sensors indicating that they should start monitoring at a higher sensing frequency. In this case, the pulse rate monitor picks up this message and then begins noting his heart rate. At the end of the exercise, the shoe again detects the change in Bob's activity and thus notifies the sink. It also notifies the body sensors, here the pulse rate monitor, to switch off.

We compare our method to the classical approach and constant sensing approach for this scenario. In our approach, the motion sensor, after being turned on, samples every 30 s. In the classical approach, the sensor periodically measures the data and sends them to the sink regardless of any activity performed. For this approach, we assume that the motion sensor measures every 5 s to detect entering the room. In the constant sensing approach, the sensing frequency

Applications/services subscribe to KPs to get notified when there is a change in data produced. Through this subscription, other sensors in the ambiance can be triggered to monitor the changes in context. Consequently, SIBs are updated and applications in that space adapt based on this new information.

of all the sensors is 1 s, and this data is sent to the sink. The pulse rate monitor senses at the rate of 1 s for all the approaches considered.

We classify Bob's activities into two parts:
- Bob's presence in the fitness room is detected.
- His pulse rate is monitored.

We performed a simulation by considering the described use case. We assumed that Bob performed the above activities at random times during 24 h. For the sake of simplicity, we neglected the time of day implications in the scenario. We use energy consumption values for various sensors in this use case from [11, 12].

Figure 4a shows the energy consumption of the sensors. As evident from the figure, our proposed method consumes less energy for detecting the contexts. The localization event detection in our method takes 61.73 percent less energy than the classical approach. In this use case both the classical and constant approaches consume the same amount of energy for heart rate monitoring. Compared to those approaches, the proposed method consumes 68.95 percent less energy. Note that the energy harvesting is not considered in the above comparisons of energy consumption.

Figure 4b shows the energy consumption for pulse rate monitoring in the classical and proposed approaches. The plot is for a random 60 minutes of Bob's activities. In the plot, we see flat portions and positive slope portions in the energy consumption of our approach. The flat portions indicate that Bob is stationary. The positive slope portions indicate that Bob is exercising and his pulse rate monitor is on. As expected, energy consumption by our approach is much lower than the classical approach, since the pulse rate monitor is turned off during Bob's inactivity.

The number of data packets generated in all approaches is shown in Fig. 4c. As expected, the number of packets in our approach is significantly less than the other approaches for both activities.

CHALLENGES

As demonstrated in the previous section, our proposed method can detect and report the context instead of raw sensor data in the classical approach. This saves significant amounts of energy. Moreover, a distributed context-aware application can be designed by employing our method. However, there are a few challenges that need to be addressed.

Quality of contextual data: In the classical approach, sensors can provide fine-grained data on a user's activity. However, in our approach we need to design smart applications [13] that try to match the requirements. In applications where fine-grained activity is critically required (e.g., in ambient assisted living for the elderly), context-triggered sensing needs to be used in conjunction with the classical approach.

Detection of events with harvesters: The detection of events in our approach is highly dependent on the sensitivity of the harvesters to the energy sources. In current harvesting systems, the sensitivity is quite low. For example,

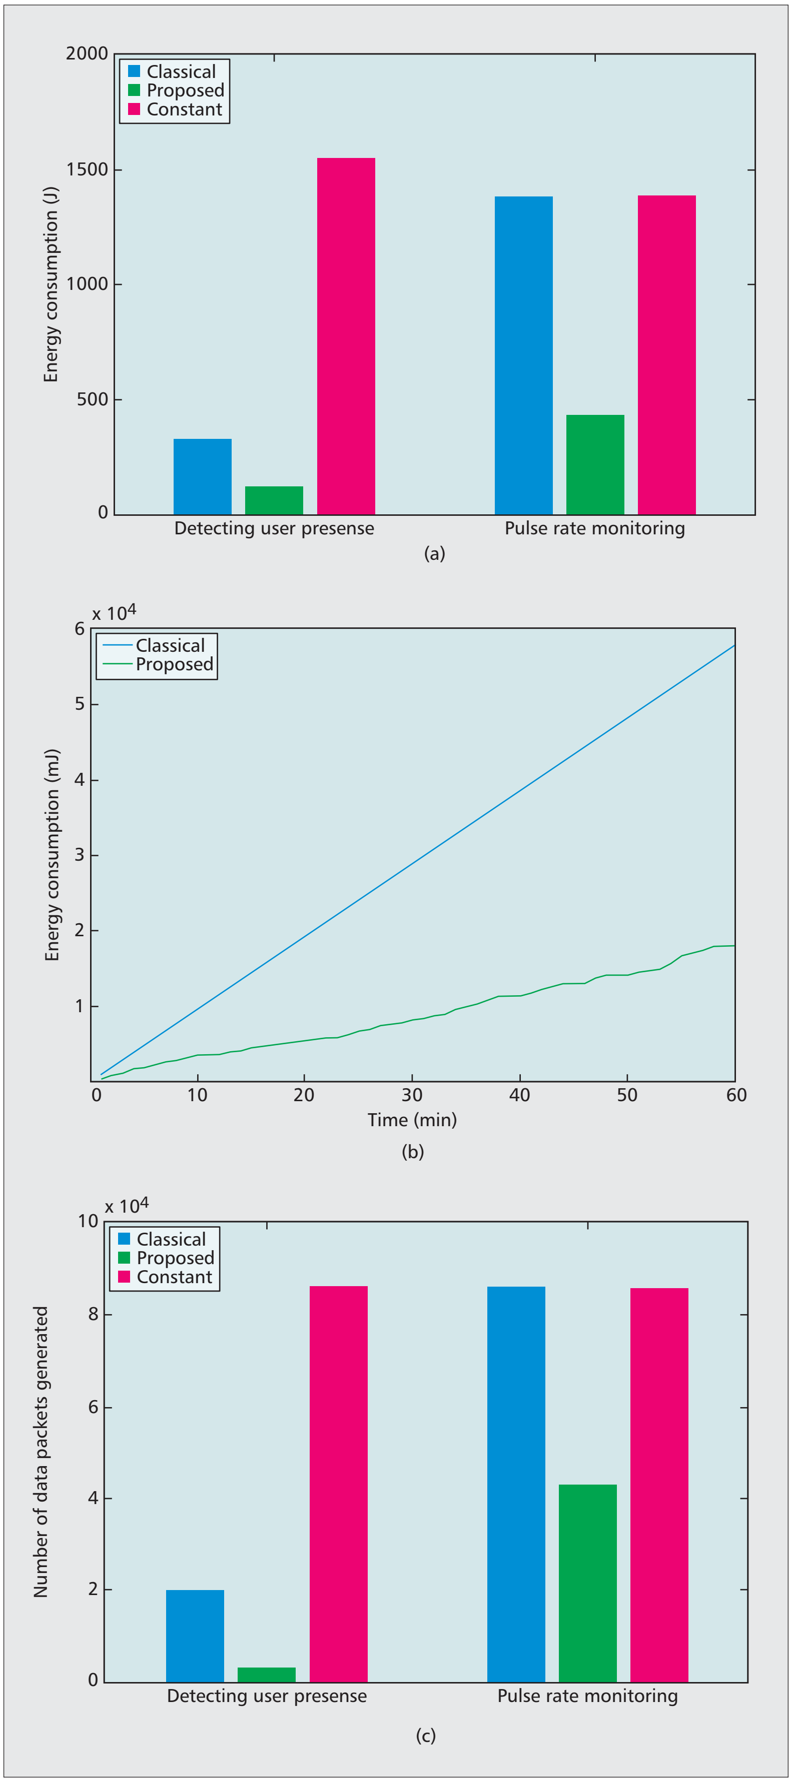

Figure 4. Comparison of different approaches: a) energy consumption for Bob's activities; b) energy consumption by pulse rate monitor; c) data packets generated by sensors.

With advancements in harvesting technologies, event detection can be made robust. We believe that using harvesters as context-event triggers is the approach for designing a truly distributed context-aware application in smart cyber physical systems.

solar panels cannot detect minor changes in illumination, especially when the panel is exposed to sunlight. Similarly, other harvesters also have a tipping point before which no energy is harvested. Making harvesters detect events robustly is an important challenge. With rapid development in harvesting technology, this problem can be addressed.

Reliability of detection in real time: Since all the nodes are run by harvested energy, some nodes may be in sleep state. Hence, making the system reliable enough to transfer the context event in real time to the sink is a major challenge in EH-WSNs.

Density of sensors: From the above challenges, we note that reliability needs to be guaranteed in such systems. Since energy harvesting is not guaranteed to always provide "enough" power, the density of such nodes need to be estimated correctly to support reliable context-event detection. Generalizing density estimation for many scenarios is a hard problem.

Distributed architecture: Having a distributed architecture facilitates the nodes, especially with our proposed method, taking distributed actions for context-aware adaptation of the environment. This makes our system scalable.

Learning over time: It is possible that a context change occurs due to more than one reason. For example, a PV harvester next to a window will trigger when sufficient sunlight falls on it. The same PV harvester will also trigger a context event when a user switches on the light in the room. Typically, when there is such ambiguity in context, the nodes are configured to report the event to the sink. The sink queries other sensors and derives the actual context. However, the nodes are required to "learn" over time to resolve such ambiguities in order to reduce latency and help realize a real-time system.

There is one drawback in the proposed approach. Not every physical parameter change can be context triggered through a harvester. For example, there is no harvester that can trigger an event due to a change in CO_2 concentration levels. Another example is that an accelerometer sensor can be used to get the orientation of the device. Unfortunately, this change cannot be used to harvest energy, so this event cannot be captured through a harvester.

Conclusions

Context awareness in applications is one of the most sought-after technologies with the growth of the Internet of Things and SCy-Phy systems. With the current technology for context-event detection, sensors need to report their measured values periodically. However, such a system is not scalable due to the facts that sensors generate huge data, and periodic data transmissions in sensors drain their battery. Ambient energy harvesting mechanisms can be exploited to address both of these issues together. We propose to exploit the harvesters to detect contextual changes in SCy-Phy spaces, and act as context-event triggers. Therefore, both data processing and energy consumption are reduced. We also adapt the context-aware framework and Smart-M3 architecture for our proposed method. We demonstrate the usefulness of our system with a real-world use case and compare our proposal with the current approach. With advancements in harvesting technologies, event detection can be made robust. We believe that using harvesters as context-event triggers is the approach for designing a truly distributed context-aware application of Scy-Phy systems.

In this space there are abundant opportunities to innovate in many ways. We have listed some important challenges, which are only the tip of the iceberg. This article provides only the beginning of highly ambitious new inter-disciplinary research.

References

[1] M. Weiser, Ed., "The Computer for the 21st Century," *Scientific American*, 1991, pp. 94–104.

[2] A. K. Dey, "Understanding and Using Context," *Personal Ubiquitous Comp.*, vol. 5, no. 1, Jan. 2001, pp. 4–7.

[3] X. Wang *et al.*, "Semantic Space: an Infrastructure for Smart Spaces," *IEEE Pervasive Computing*, vol. 3, no. 3, 2004, pp. 32–39.

[4] K.-D. Kim and P. Kumar, "Cyber Physical Systems: A Perspective at the Centennial," *Proc. IEEE*, vol. 100, 2012, pp. 1287–1308.

[5] R. V. Prasad *et al.*, "Reincarnation in the Ambiance: Devices and Networks with Energy Harvesting," *IEEE Commun. Surveys & Tutorials*, vol. PP, no. 99, 2013, pp. 1–19.

[6] S. Sudevalayam and P. Kulkarni, "Energy Harvesting Sensor Nodes: Survey and Implications," *IEEE Commun. Surveys & Tutorials*, vol. 13, no. 3, 2011, pp. 443–61.

[7] J. Honkola *et al.*, "Smart-m3 Information Sharing Platform," *IEEE Symp. Computers and Commun.*, 2010, pp. 1041–46.

[8] V. Penela, C. Ruiz, and J. Gómez-Pérez, "What Context Matters? Towards Multidimensional Context Awareness," *Ambient Intelligence and Future Trends — Int'l. Symp. Ambient Intelligence*, Springer, vol. 720, 2010, pp. 113–20.

[9] S. Pantsar-Syvaniemi, K. Simula, and E. Ovaska, "Context-Awareness in Smart Spaces," *IEEE Symp. Computers and Commun.*, 2010, pp. 1023–28.

[10] N. Rodriguez, "A Framework for Context-Aware Applications for Smart Spaces," *IEEE/IPSJ 11th Int'l. Symp. Applications and the Internet*, 2011, pp. 218–21.

[11] E. Reetz, R. Tonjes, and N. Baker, "Towards Global Smart Spaces: Merge Wireless Sensor Networks into Context-Aware Systems," *5th IEEE Int'l. Symp. Wireless Pervasive Computing*, 2010, pp. 337–42.

[12] P. Gyorke and B. Pataki, "Energy-Aware Measurement Scheduling in WSNs Used in AAL Applications," *IEEE Trans. Instrumentation and Measurement*, vol. 62, no. 5, 2013, pp. 1318–25.

[13] T. Prabhakar *et al.*, "A Distributed Smart Application for Solar Powered WSN," *Networking*, LNCS, Springer, 2012.

[14] M. D. Vithanage *et al.*, "Medium Access Control for Thermal Energy Harvesting in Advanced Metering Infrastructures," *IEEE EUROCON*, 2013, pp. 291–99.

Biographies

Vijay S. Rao (V.Rao@tudelft.nl) is currently pursuing his Ph.D. in the Embedded Software group of Delft University of Technology, The Netherlands. He received his M.S.in telecommunications from Delft University of Technology in 2009 and his B.Eng. in information science and technology from Visvesvaraya Technological University, India, in 2004, both with distinction. He worked for Trans research in the Netherlands in 2010. During his Master's research, he worked part-time on the EU project MAGNET-Beyond and the Dutch project PNP2008, both in the area of personal networks. Between 2004 and 2007, he worked as a senior software engineer at ESQUBE Communication Solutions, Bangalore, India, where he designed and developed VoIP products. His main research interests lie in communication networks, especially energy harvesting networks, cognitive radio networks, and the Internet of Things.

S. N. Akshay Uttama Nambi (Akshay.Narashiman@tudelft.nl) is a Ph.D. student in the Embedded Software group of

Delft University of Technology. Prior to this, he worked as a research assistant at École Polytechnique Fédérale de Lausanne (EPFL), Switzerland, and the Indian Institute of Science (IISc), Bangalore. He received a B.Eng. in computer science and engineering from Visvesvaraya Technological University with distinction. His research interests are broadly in the area of cyber physical systems, smart grids, energy harvesting systems, and the Internet of Things.

R VENKATESHA PRASAD (R.R.VenkateshaPrasad@tudelft.nl) received his Bachelor's degree in electronics and communication engineering and M.Tech degree in industrial electronics from the University of Mysore, India, in 1991 and 1994. He received a Ph.D. degree in 2003 from IISc. During 1996 he worked as a consultant and project associate for THE ERNET Lab of ECE at IISc. While pursuing his Ph.D. degree, from 1999 to 2003 he also worked as a consultant for CEDT, IISc, Bangalore for VoIP application development as part of Nortel Networks sponsored project. In 2003 he headed a team of engineers at ESQUBE Communication Solutions for the development of various real-time networking applications. From 2005 to 2012, he was a senior researcher at the Wireless and Mobile Communications group, Delft University of Technology, working on the EU funded projects MAGNET/MAGNET Beyond and PNP-2008, and guiding graduate students. From 2012 onward, he has been an assistant professor at the Embedded Software group at Delft University of Technology. He is an active member of TCCN and IEEE SCC41, and a reviewer of many IEEE transactions and Elsevier journals. He is on the Technical Program Committees of many conferences, including IEEE ICC, IEEE GLOBECOM, ACM MM, ACM SIGCHI, and others. He is TPC Co-Chair of the CogNet Workshop in 2007, 2008, and 2009, and TPC Chair for E2Nets at IEEE ICC '10. He is also running PerNets workshop from 2006 with IEEE CCNC. He was the tutorial Co-Chair of CCNC 2009 and 2011, and Demo Chair of IEEE CCNC 2010. He is Secretary of the IEEE ComSoc Standards Development Board and Associate Editor of *Transactions on Emerging Telecommunications Technologies*.

IGNAS NIEMEGEERS (I.G.M.M.Niemegeers@tudelft.nl) received a degree in electrical engineering from the University of Ghent, Belgium, in 1970. In 1972 he received an M.Sc.E. degree in computer engineering and in 1978 a Ph.D. degree from Purdue University in West Lafayette, Indiana. From 1978 to 1981 he was a designer of packet switching networks at Bell Telephone Mfg. Cy, Antwerp, Belgium. From 1981 to 2002 he was a professor at the Computer Science and Electrical Engineering faculties of the University of Twente, Enschede, The Netherlands. From 1995 to 2001 he was scientific director of the Centre for Telematics and Information Technology (CTIT) of the University of Twente, a multi-disciplinary research institute on ICT and applications. From May 2002 until his retirement in 2012, he held the chair Wireless and Mobile Communications at Delft University of Technology, where he headed the Telecommunications Department. He is currently a Professor Emeritus of Delft University of Technology. He was involved in many European research projects, including the EU projects MAGNET and MAGNET Beyond on personal networks, EUROPCOM on UWB emergency networks, and eSENSE and CRUISE on sensor networks. His present research interests are 5G wireless infrastructures, future home networks, ad hoc networks, personal networks, and cognitive networks. He has (co)authored close to 300 scientific publications and has coauthored a book on personal networks.