Proceedings of the IDEA League Doctoral School on Transiently Powered Computing

Muhammad Hamad Alizai
Jan Beutel
Jó Ágila Bitsch
Olaf Landsiedel
Luca Mottola
Przemysław Pawełczak
Klaus Wehrle
Kasım Sinan Yıldırım

Department of Computer Science

Technical Report

Aachener Informatik-Berichte (AIB) | ISSN 0935-3232 | AIB-2017-09
RWTH Aachen University | Department of Computer Science | November 2017

The publications of the Department of Computer Science of *RWTH Aachen University* are in general accessible through the World Wide Web.

http://aib.informatik.rwth-aachen.de/

Table of Contents

Poster: Network Bootstrapping and Leader Election Utilizing the Capture Effect in Low-power Wireless Networks

Beshr Al Nahas
beshr@chalmers.se
Chalmers University of Technology
Gothenburg, Sweden

Simon Duquennoy
simon.duquennoy@ri.se
RISE SICS, Sweden
Inria Lille - Nord Europe, France

Olaf Landsiedel
olafl@chalmers.se
Chalmers University of Technology
Gothenburg, Sweden

ABSTRACT

Many protocols in low-power wireless networks require a *leader* to bootstrap and maintain their operation. For example, Chaos and Glossy networks need an initiator to synchronize and initiate the communication rounds. Commonly, these protocols use a fixed, compile-time defined node as the leader. In this work, we tackle the challenge of dynamically bootstrapping the network and electing a leader in low-power wireless scenarios.

CCS CONCEPTS

• **Networks** → **Network protocol design**; **Sensor networks**;

KEYWORDS

Synchronous transmissions, Capture effect, WSN, IoT

1 INTRODUCTION

Context and Challenge. Many protocols in low-power wireless networks require an entity to bootstrap and maintain the operation, which we denote a *leader*. For example, RPL networks need a network root to build the routing tree and Glossy/Chaos networks need an initiator to synchronize and initiate the communications rounds. In applications that build their operation on consensus; *e.g.*, two-phase commit, the leader is responsible for proposing and committing transactions. In recent work [1, 3], the common solution was to use a fixed, compile-time defined node as the leader.

The use of a statically defined leader exhibit the following weaknesses; (a) it assumes a known network deployment; thus, it does not suit random deployments; *e.g.*, throwing nodes from the air; (b) it assumes a static network; thus, mobility is limited, and, (c) initiator failure means a network failure and might require manual intervention to restart the network operation. While the problem of clustering and leader election is not new as it was tackled by Heinzelman *et al.* in LEACH [2] and subsequent work, there is a need for an approach that both suits and benefits from the low latency of recent approaches to synchronous transmissions, such as Glossy and Chaos.

Approach. In this paper, we tackle the challenge of dynamically electing a leader in low-power wireless networks. We propose mechanisms that achieve (a) network bootstrapping; *i.e.*, network synchronization, and clustering; (b) leader election and ensuring the convergence toward one leader and (c) leader failure recovery. We build on top of our work A^2 [4] and its lower layer, Synchrotron, the synchronous transmission protocol that is inspired by Chaos.

Outline. We provide the required background on A^2 and synchronous transmission in §*2*. Then we explain the design in §*3* and conclude with preliminary results in §*4*.

2 BACKGROUND: A^2 AND SYNCHROTRON

A^2 builds on top a synchronous transmissions kernel, Synchrotron, and utilizes in-network processing to provide primitives for network-wide, all-to-all dissemination, collection, aggregation, voting, consensus (two- and three-phase commit) and membership services. A^2 operates in *rounds* where nodes send packets synchronously and receive data thanks to the *capture effect*.

Synchrotron: Synchronous transmissions and capture effect. Synchrotron roots in approaches to synchronous transmissions, such as Chaos, where multiple nodes synchronously transmit the data they want to share. Nodes overhearing the concurrent transmissions receive one of them with high probability, due to the capture effect. For example, to achieve capture with IEEE 802.15.4 radios, nodes need to start transmitting within the duration of the preamble of $160\mu s$ [3].

Synchrotron operates as a time-slotted protocol. The minimum time unit is a *slot*, which fits one packet transmission/reception and processing. Slots are grouped in *rounds*, where a designated function, such as collect or disseminate is run network-wide. Within each slot, a node transmits, receives or sleeps according to the transmission policy of the application.

In-network aggregation. In A^2, each packet contains *progress flags*, with one bit assigned to each node in the network. The initiator starts a *round* by sending a packet with its own flag set. Upon successful reception, a node sets its flag and merges the received packet with own. It transmits in the next time slot when it receives new information, *i.e.*, new flags, or when it hears a node transmitting messages with fewer flags set, *i.e.*, a neighbor knows less than the node. The process continues until all nodes have set their flag.

Similar to Chaos, the rules for merging are application specific. With the *Max* operation, for example, A^2 identifies the maximum value: Next to the flags, the only payload is the maximum value collected so far. Upon reception, nodes compute the max between their local value and the payload, write it to the packet, merge the flags, and set their flag before transmitting in the next time-slot.

3 DESIGN

To start a proper network operation in A^2, there shall be an initiator node or a *leader* that (a) ensures network-wide *synchronization*; (b) *joins* the nodes that wish to participate in the network and assigns each of them a unique ID; and (c) initiate the communication rounds and ensures the application objectives are met.

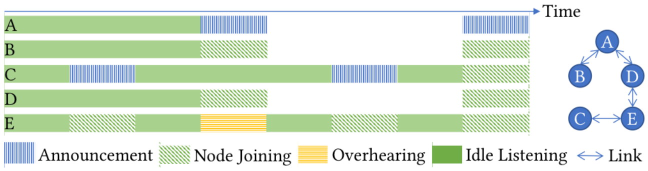

Figure 1: An example run with five nodes. *Nodes A and C propose as leaders. Due to the network topology, two clusters form. E and C keep listening between rounds since their cluster has less than half of the nodes. E overhears A's cluster and notifies its cluster members. The network converges to select A as the leader.*

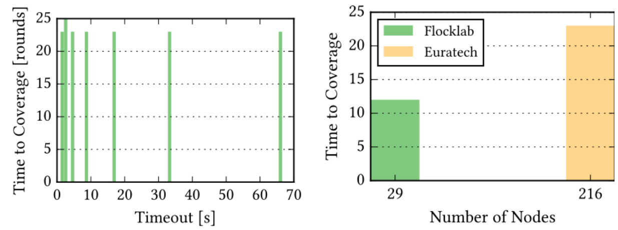

(a) Election timeout does not affect convergence time. *The Euratech network converges within 23 rounds = 41 seconds on average.*

(b) Leader election on a sparse and a dense testbed; *repectively, Flocklab and Euratech.*

Figure 2: The time it takes the leader election procedure to converge to one cluster and join all the nodes in the network.

To be able to communicate, we need to synchronize the nodes. We start by forming clusters that ensure neighborhood synchronization. Then, we merge the clusters gradually to form one network-wide cluster with a single leader.

3.1 Bootstrapping and Clustering

We start by having two assumptions (a) the nodes are homogeneous and any of them could be a leader. This assumption is by no means compulsory, but it simplifies the discussion; and (b) the maximum number of nodes is known before hand.

Every node start by listening to the radio and generates a random timeout. It keeps listening until it hears a valid A^2 message to synchronize on. If it times out without hearing, then it proposes itself a leader and starts sending *join* announcements such that other nodes hear them and join it to form a cluster. At this phase, multiple clusters could form as different nodes might not hear each other. The next step is to converge towards one leader in one cluster.

3.2 Leader Election

To ensure the convergence toward one leader, we put a quorum stability condition: A stable cluster is the one that has more than half of the nodes. Given that the nodes cannot join more than one cluster, we can ensure convergence.

Until a cluster is stable, it keeps running join rounds, and its members keep sampling the medium between the communication rounds looking for bigger clusters to join. When a node hears another cluster, it saves the synchronization information of the largest foreign cluster it heard. Each node shares this information in the next join round with its cluster, and use the *Max* primitive to find the information about the largest cluster. At the end of the join round, nodes drop their cluster, their IDs and jump to join the new cluster if it is bigger than their current cluster. With time, only the largest cluster survives, and only one leader exists.

3.3 Failure Recovery

Upon leader failure, nodes no longer hear packets, and the random timeout mechanism kicks in. This restarts the whole process and elects a new leader as illustrated in §*3.1* and §*3.2*.

4 PRELIMINARY RESULTS AND CONCLUSION

Implementation. We implement the algorithm in §*3* in C for the Contiki OS targeting wireless nodes equipped with a low-power radio such as the TelosB platform which features a 16bit MSP430 CPU @ 4 MHz, 10 kB of RAM, 48 kB of firmware storage and CC2420 radio compatible with 802.15.4.

Figure *1* illustrates an example run on a network of 5 nodes. Figure *2* summarizes the results of running on the testbeds FIT-IoTLAB Euratech and Flocklab. First, we vary the maximum timeout period and run on Euratech. We find that the choice of the timeout has minimal effect on the time to converge to one cluster that inlcudes all the nodes. Second, we compare the performance when running on sparse and dense testbeds. The network converges to one leader and all the nodes join the leader within 41 seconds for 216 nodes (on Euratech) and 21 seconds for 29 nodes (on Flocklab).

ACKNOWLEDGMENTS

This work was supported by the Swedish Research Council (VR) through the project ChaosNet, the Swedish Foundation for Strategic Research (SSF) through the project LoWi, CPER Nord-Pas-de-Calais/FEDER DATA and Sweden's innovation agency (VINNOVA).

REFERENCES

[1] Federico Ferrari, Marco Zimmerling, Luca Mottola, and Lothar Thiele. 2012. Low-Power Wireless Bus. In *Proceedings of the Conference on Embedded Networked Sensor Systems (ACM SenSys).*

[2] W. R. Heinzelman, A. Chandrakasan, and H. Balakrishnan. 2000. Energy-efficient communication protocol for wireless microsensor networks. In *Proceedings of the Annual Hawaii International Conference on System Sciences.*

[3] Olaf Landsiedel, Federico Ferrari, and Marco Zimmerling. 2013. Chaos: Versatile and Efficient All-to-All Data Sharing and In-Network Processing at Scale. In *Proceedings of the Conference on Embedded Networked Sensor Systems (ACM SenSys).*

[4] Beshr Al Nahas, Simon Duquennoy, and Olaf Landsiedel. 2017. Network-wide Consensus Utilizing the Capture Effect in Low-power Wireless Networks. In *Proceedings of the Conference on Embedded Networked Sensor Systems (ACM SenSys).*

Design Support for Energy Harvesting Driven IoT Devices

Lukas Sigrist
ETH Zurich
Computer Engineering and Networks Laboratory
Zurich, Switzerland
sigristl@ethz.ch

Lothar Thiele
ETH Zurich
Computer Engineering and Networks Laboratory
Zurich, Switzerland
thiele@ethz.ch

ABSTRACT

With the emerging Internet of Things, wireless sensor applications are increasingly being supplied from energy harvesting. While this shift away from batteries provides many advantages, it also increases the complexity of designing these highly energy constrained systems. Due to environmental dependencies, novel tools are necessary to support their design process. With the RocketLogger we introduce a measurement tool that addresses the specific needs of energy harvesting systems. It provides a highly increased dynamic range for current measurements, accommodating both ultra-low sleep currents of few nanoamperes as well as wireless communication currents in the range of hundreds of milliamperes. The portable design allows for in-situ measurements for characterizing and validating the system performance, while an extensible sensor bus provides flexible recording of the application specific environmental variables from which the energy is extracted. While the RocketLogger is an important first step, additional tools are still required to provide the necessary support for designing efficient and reliable harvesting-driven systems.

KEYWORDS

Energy Harvesting, Internet of Things, Wireless Sensor Networks.

1 INTRODUCTION

Energy harvesting is seen as a key long-term technology to power the billions of devices of the emerging Internet of Things (IoT). Advances in low power system design have enabled battery powered wireless sensor networks with increasing lifetimes that are today deployed in various application areas. However, with the vast amount of IoT devices and their deployment in hardly accessible locations, batteries are not a practical option, since their limited lifetimes would require expensive maintenance. Extracting energy from the surrounding environment is therefore seen as a key solution to the energy supply problem [1]. This lead to an increased integration of energy harvesting into new wireless sensor node designs and the new research of transiently powered computing systems, focusing on the design of highly energy constrained, batteryless systems that progress only as a function of their environment.

Unfortunately, the use of energy harvesting comes with additional complexity in the design process of these systems. Where previously a constant supply was guaranteed, the variable and application-specific environmental conditions have a direct impact on parameters like e.g. harvested energy and power conversion efficiencies [2]. To efficiently use the small amount of energy harvested by tiny IoT devices, these systems need to adapt to the ever changing supply conditions, which vary over a large range from nW to mW.

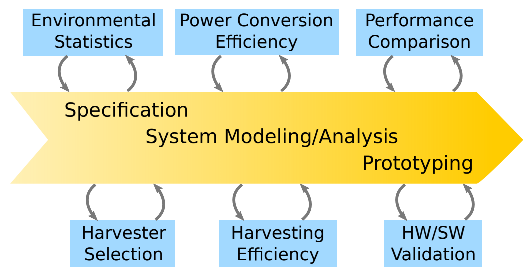

Figure 1: Sample design aspects of harvesting-based systems

Only application-specific solutions could provide the desired performance given the environmental, size and cost restrictions. While these solutions are highly specific, they all demand for a precise characterization of the environment-dependent energy budget, characterization and optimization of the application's active and sleep energy consumption during their design, as illustrated in Figure 1. The measurement of these metrics, however, include widely ranging currents: quiescent/sleep currents in the order nA-μA, and active currents of 100's mA. Because of the application and environment specific design, tools must be deployed with the energy harvesting system to record its behavior the production environment. Although measuring power and environmental properties are well known problems, portable tools that accurately and reliably measure harvesting-based systems do not exist.

Building these tools provides several challenges: they are required to measure a wide dynamic power range and environmental conditions in the field and for long periods of time. In addition, these tools must run independently from the system being measured, while minimizing the impact on them. Otherwise, the measured systems will not work in adverse power conditions or suffer a significantly degraded harvesting or power conversion efficiency.

While the RocketLogger presented in this work provides the required measurement capabilities, we also motivate the need for novel tools allowing to rapidly explore various system and harvester configurations under the same environmental conditions to further improve and accelerate the design of harvesting driven systems.

2 PORTABLE ACCURACY MEASUREMENTS

Currently available measurement equipment does not fulfill all the requirements of these different, but related design aspects outlined above. Portable or embedded solutions designed for battery operated devices lack ultra-low current measurement in the nanoampere range [3, 6], do not feature the desired dynamic range from below micro- up to hundreds of milliamperes [4, 5, 9], or do not

Table 1: RocketLogger performance overview

Characteristic	Performance
Voltage Range/Accuracy	±6 µV – 5.5 V / 0.02 % + 13 µV
Current Range/Accuracy	±1 nA – 500 mA / 0.09 % + 4 nA
Sampling Rate	up to 64 kSPS
Voltage Channel Impact	input impedance ~ 1 TΩ
Current Channel Impact	max. voltage drop ≤ 53 mV

target portable measurements with remote control [5, 6]. Benchtop devices with the desired performance exist, but focus on single measurement tasks with high accuracy and are bulky. They are therefore infeasible for mixed-signal in-situ measurements.

To address the specific measurement challenges in the designing and evaluation of energy harvesting systems we introduce the RocketLogger [7]. The novel mixed-signal data logger features high-accuracy power measurements and logging of environmental conditions like temperature or illuminance in a portable design of 103 mm × 68 mm. Two seamlessly switched current channels combined with up to four voltage channels, allow uninterrupted power measurements with a large dynamic range from 40 pW at 10 mV up to 2.75 W at 5.5 V with minimal impact on the device being measured. A remote web interface facilitates control and observation of long-term in-situ measurements. An brief summary of the loggers measurement performance is given in Table 1.

To show how the RocketLogger's unique features enhance the design process of harvesting driven systems, we consider a wearable multi-source harvesting circuit. The harvesting power and environmental conditions shown in Figure 2 were recorded in a scenario where the user walks outside during a warm, sunny day with high illuminance levels and then enters a colder, darker indoor space allowing for higher body to ambient temperature gradients. For the harvester measurements, it should be noted that the harvested TEG power is in the order of 100's µW, while solar power is in the mW range. This result shows that the solar harvesting power dominates outdoors, while the TEG generates more power indoor, although at a lower power level. This data is very valuable for subsequent iterations of system modeling and analysis to optimize important system parameters like harvesting efficiency.

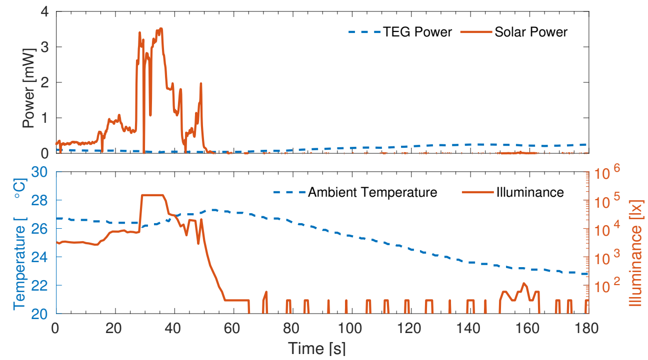

Figure 2: RocketLogger case study: in-situ measurement of a multi-source wearable harvesting system [7]

3 CONCLUSION AND FUTURE WORK

In this work we motivated the need for novel tools to support the design process of energy harvesting driven applications. The presented RocketLogger device provides a combination of high-accuracy power measurements with large dynamic range, environmental logging, a mobile form factor, and an easy-to use remote interface. This unique set of features make it a versatile measurement instrument satisfying essential needs of the system design process for energy harvesting driven applications. While this solves important measurement needs in the design process of harvesting-driven systems, these measurements still require the deployment of the devices in the actual application environment.

We therefore see the need of emulating the physical environment in the lab to provide repeatable and consistent environmental conditions to accelerate the design process. State-of-the-art harvesting source emulators [3] allow recording and emulation of transducers like solar panels or TEGs in the application environment with reasonable accuracy. On other hand these solutions are limited to the emulation of the behavior of one specific transducer under specific environmental conditions. However, for the design of an optimized harvesting application, careful selection and precise characterization of the transducers is crucial to tune the harvesting circuitry and other system parameters [8]. For this reason we will focus on the emulation of the physical processes from which the energy is being extracted rather than the harvested power of a specific transducer. We are convinced that accurately reproducible environmental conditions are an important step that enables the comparison of different implementations of harvesting-driven applications and brings us closer to fully optimizing the design of purely energy harvesting-driven, transient computing systems.

ACKNOWLEDGMENTS

This research was funded by the Swiss National Science Foundation under grant 157048: Transient Computing Systems.

REFERENCES

[1] Naveed Anwar Bhatti, Muhammad Hamad Alizai, Affan A Syed, and Luca Mottola. 2016. Energy Harvesting and Wireless Transfer in Sensor Network Applications: Concepts and Experiences. *ACM Transactions on Sensor Networks* 12, 3 (2016).
[2] Andres Gomez, Lukas Sigrist, Thomas Schalch, Luca Benini, and Lothar Thiele. 2017. Efficient, Long-Term Logging of Rich Data Sensors using Transient Sensor Nodes. *ACM Transactions on Embedded Computing Systems* (2017).
[3] Josiah Hester, Timothy Scott, and Jacob Sorber. 2014. Ekho: Realistic and Repeatable Experimentation for Tiny Energy-harvesting Sensors. In *Proc.SenSys'14*.
[4] Xiaofan Jiang, Prabal Dutta, David Culler, and Ion Stoica. 2007. Micro Power Meter for Energy Monitoring of Wireless Sensor Networks at Scale. In *Proc. IPSN'07*.
[5] Saman Naderiparizi, Aaron N Parks, Farshid Salemi Parizi, and Joshua R Smith. 2016. µMonitor: In-situ Energy Monitoring with Microwatt Power Consumption. In *Proc. RFID'16*.
[6] A Pötsch, A Berger, C Leitner, and A Springer. 2014. A power measurement system for accurate energy profiling of embedded wireless systems. In *Proc. ETFA'14*.
[7] Lukas Sigrist, Andres Gomez, Roman Lim, Stefan Lippuner, Matthias Leubin, and Lothar Thiele. 2017. Measurement and Validation of Energy Harvesting IoT Devices. In *Proc. DATE'17*.
[8] Moritz Thielen, Lukas Sigrist, Michele Magno, Christofer Hierold, and Luca Benini. 2017. Human Body Heat for Powering Wearable Devices: From Thermal Energy to Application. *Energy Conversion and Management* 131 (2017), 44–54.
[9] R Zhou and G Xing. 2013. Nemo: A high-fidelity noninvasive power meter system for wireless sensor networks. In *Proc. IPSN'13*.

Adaptive Wireless Sensor Networks: Robust but Efficient

Extended Abstract

Robert Hartung
Institute of Operating Systems and Computer Networks
hartung@ibr.cs.tu-bs.de

Ulf Kulau
Institute of Operating Systems and Computer Networks
kulau@ibr.cs.tu-bs.de

Lars Wolf
Institute of Operating Systems and Computer Networks
wolf@ibr.cs.tu-bs.de

ABSTRACT

Future Wireless Sensor Network deployments need to be highly adaptive to changing conditions, while being both robust and efficient at the same time. This paper introduces challenges and possible solutions to this new class of adaptive Wireless Sensor Networks. We present our research project *REAP* and present past, current and future research directions and preliminary results.

1 INTRODUCTION

Wireless Sensor Networks (WSNs) can be used in a variety of applications, which sometimes are challenging. Especially in outdoor environments such as precision farming or smart cities, the demands to WSNs is to be highly reliable and robust. However, the most crucial points to such deployments is the lifetime of the overall network. Lifetime is traditionally dependent on the size of batteries and the duty cycling of the hardware including communication devices. Due to the recent development of energy harvesting platforms, micro controllers and complete sensors nodes can be powered from other sources than batteries. A challenge to these new class of applications is the unreliability and unpredictability of energy harvested from different sources.

Traditional WSNs have been deployed and operated until their end of lifetime (e.g. when the battery died). A key goal of research for these WSNs is to extend lifetime as much as possible, e.g. due to using sleep modes and duty cycling the radio transceiver. In contrast, more recent WSNs including those in the context of the Internet of Things (IoT), are likely to be a lot more dynamic and need to adapt to dynamic scenarios, including available energy [6]. Due to the high variability of energy available, the mutual influence of efficiency and dependability is likely to increase. Traditionally, the energy budget was limited and known (battery capacity) and frequency of sensor readings was fixed. This resulted in a good estimation of the lifetime of a sensor node and the whole network. Hence, with the introduction of energy harvesting and the unreliability and unpredictability of these energy sources, an imbalance between efficiency and dependability is likely to arise. In addition the changing environmental conditions have a huge impact on the baseline-reliability of WSNs per se. Especially the temperature can prune the performance of wireless communication which leads to another unpredictable energy demand to compensate such issues [7].

To overcome this challenge, we introduce the term *adaptability* for this new class of WSNs, where both hardware and software need to be highly adaptable to its environment. Figure 1 illustrates this adaptability in terms of low-power and transiently-powered devices. Even with low energy budgets, we aim to balance efficiency and dependability.

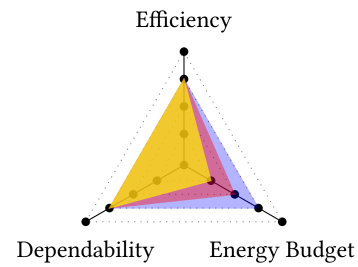

Figure 1: Illustration of the dynamics of WSNs

The next section will briefly describe the research directions of the *REAP* project towards **R**obust, **E**nergy Efficient Wireless Sensor Networks for Outdoor Scenarios by **A**daption of Operation **P**arameters. Finally, section 3 will explain adaptability in relation to our project.

2 REAP PROJECT

The *REAP* project is a 3-year project funded by the German Research Council (DFG) under the grant no. BU 3282/2-1. It is focused on improving robustness and efficiency of WSNs under changing environmental conditions. Improving hardware is our first task.

The improvement of the application's software, including routing protocols is the second focus of the *REAP* project.

2.1 Adaptability

As previously shortly introduced, we use the term *adaption* or *adaptabiliy* to describe a new class of WSNs that has to deal with a highly dynamic context. The presented concepts can be used to achieve this goal. An energy harvester has not only a limited power output, but may also be unreliable or unpredictable. Thus, a constant execution of tasks (e.g. measurements) might not be guaranteed, due to a low energy budget. Additionally, a dynamic adaption of the sensor sample rate can be used to increase network's lifetime [6]. The following sections will describe how the *REAP* project faces this challenge.

2.2 Past research

In the past, some prior research have already been done to complete the overall tasks of the project.

We developed IdealVolting [4], that allows operating the sensor node below its specified voltage levels. It was proven that this is safe and does not affect the WSN functionality, while allowing to operate at up to 40% less energy consumption. Within REAP IdealVolting will further be coupled with energy harvesting and used in an adaptive manner to allow steady QoS despite changing

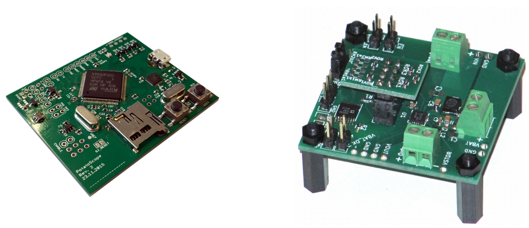

(a) PotatoScope **(b) REAPer**

Figure 2: Pictures of different components used in the REAP project

energy budgets. Our dynamic sensor sample rate adaption [6] can be used to react to the change in sensed data to increase or decrease sampling rate dynamically. An example is the measurement of the air temperature, where changes are small during night time. A dynamic sample rate can help to reduce the amount of data and energy consumption during night time. A helpful tool has been accomplished within the project to allow a reliable online energy measurement – the *PotatoScope* [2] (cf.Figure 2a). It is used for a distributed, temperature-invariant measurement of devices such as sensors nodes. We use this to measure the actual power consumption of our software and other devices within the network (such as sensors etc.) even for rough environmental conditions.

2.3 Current Research

The current research focuses on energy harvesting solutions and deriving temperature-dependencies of sensor nodes. The first research direction focuses on building a modular energy harvesting platform that can be used to harvest energy from different sources and store energy accordingly. A prototype called *REAPer* [3] of this platform is shown in Figure 2b. It can use different types of storage, including super capacitors. The second research direction aims to find out more about temperature dependency of sensor nodes. While this is a well-researched objective for ChipCon-based radio transceivers, a lack of literature exists for others. We are the first to make intensive experiments for other radio transceivers to find generic solutions for classifications. The third part we are currently working on is a robust version of the *RPL* protocol. We have researched the impact of transient node resets [5] which results in a big energy loss even for short resets of a single node.

2.4 Future Research

Future research aims to improve existing protocols such as the RPL protocol [1]. With the knowledge about temperature dependencies, as well as charging and discharging behavior of super capacitors and other energy stores, we will be able to predict the environmental conditions to react accordingly, e.g. by local scheduling or topology changes. In combination with IdealVolting a plethora of research is to be expected.

A simple example of this idea is shown in Figure 3 (solely focused on the energy budget). Figure 3a shows the network before an adaption. Node A wants to send data to the sink and takes the route via B and C. Nodes C and D are sort of a bottle neck in the network, because they have to route data for a disproportionately

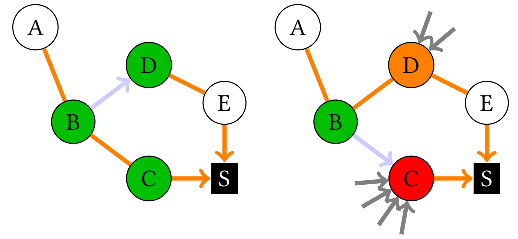

(a) Routes before adaption **(b) Routes after adaption**

Figure 3: RPL: Example of adaptive routes due to changes in available energy

number of nodes, compared to the network average, indicated by arrows pointing towards them. This results in a low energy state, as indicated by their color (orange and respectively red). The adaption within *RPL* will automatically choose a different route to increase lifetime of the overall network, because C is already low on energy and B should not route data along C. However, the goal is to include more metrics (e.g. reliability or efficiency) to allow robust but energy efficient WSNs. Furthermore, scheduling tasks based on available and predicted energy will be part of this research.

3 CONCLUSION

In conclusion we will make two major contributions to the field. First, extensive experiments will provide information about the temperature dependency of sensor nodes and related energy storage devices such as batteries and super capacitors commonly used in energy harvesters. Second, existing protocols and task scheduling will be adapted with information about current and predicted available energy. Finally, this should improve the network's lifetime and efficiency while being robust at the same time.

REFERENCES

[1] Roger Alexander, Anders Brandt, JP Vasseur, Jonathan Hui, Kris Pister, Pascal Thubert, P Levis, Rene Struik, Richard Kelsey, and Tim Winter. 2012. RPL: IPv6 Routing Protocol for Low-Power and Lossy Networks. RFC 6550. (March 2012). https://doi.org/10.17487/RFC6550

[2] R. Hartung, U. Kulau, and L. Wolf. 2016. Distributed Energy Measurement in WSNs for Outdoor Applications. In *2016 13th Annual IEEE International Conference on Sensing, Communication, and Networking (SECON)*. 1–9.

[3] Ulf Kulau, Daniel Bräckelmann, Felix Büsching, Sebastian Schildt, and Lars Wolf. 2017. REAPer – Adaptive Micro-Source Energy-Harvester for Wireless Sensor Nodes. In *Twelfth IEEE Workshop on Practical Issues in Building Sensor Network Applications 2017 (IEEE SenseApp 2017)*. Singapore. accepted for pubication.

[4] Ulf Kulau, Felix Büsching, and Lars Wolf. 2016. IdealVolting: Reliable Undervolting on Wireless Sensor Nodes. *ACM Trans. Sen. Netw.* 12, 2, Article 11 (April 2016), 38 pages. https://doi.org/10.1145/2885500

[5] Ulf Kulau, Silas Müller, Sebastian Schildt, Arthur Martens, Felix Büsching, and Lars Wolf. 2017. Energy Efficiency Impact of Transient Node Failures when using RPL. In *Proceedings of the 18th IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM 2017)*. Macau, China. accepted for publication.

[6] Ulf Kulau, Johannes van Balen, Sebastian Schildt, Felix Büsching, and Lars Wolf. 2016. Dynamic Sample Rate Adaptation for Long-Term IoT Sensing Applications. In *IEEE World Forum on Internet of Things 2016 (WF-IoT)*. Reston, USA.

[7] Hjalmar Wennerström, Frederik Hermans, Olof Rensfelt, Christian Rohner, and Lars-Åke Nordén. 2013. A long-term study of correlations between meteorological conditions and 802.15. 4 link performance. In *Sensor, Mesh and Ad Hoc Communications and Networks (SECON), 2013 10th Annual IEEE Communications Society Conference on*. IEEE, 221–229.

Energy Efficient Multi-Connectivity in Ultra-Dense 5G Networks

[Extended Abstract] *

Valentin Poirot
Chalmers University of Technology, Sweden
poirotv@chalmers.se

Mårten Ericson
Ericsson Research
Luleå, Sweden

Karl Andersson
Luleå University of Technology, Sweden

ABSTRACT

In 5G systems, two radio air interfaces, evolved LTE and New Radio (NR), will coexist, the former providing coverage and the latter high capacity. A feature for 5G called multi-connectivity has therefore been proposed, allowing multiple simultaneous links from one user to different access nodes. In this work, we investigate how multi-connectivity can improve user reliability and the system's energy efficiency. Multiple algorithms for secondary cell association are presented and evaluated. We show a decrease of the radio link failure rate of up to 50% at high speeds and an improvement of the energy efficiency of up to 20%.

CCS Concepts

•**Networks → Network performance evaluation;** Network simulations;

Keywords

Multi Connectivity, Energy Efficiency, Ultra Dense Network, 5G, Multi RAT

1. INTRODUCTION

Mobile networks are being more and more used throughout the world. It is expected that the total mobile traffic will increase by a factor of ten in 2022 [2], while new usages, such as connected vehicles or the industry 4.0, will require stronger constraints on the network, such as ultra-low latency and ultra-high reliability. The next generation of mobile communications, 5G, is therefore being designed as a way to solve these challenges.

Amongst the numerous technologies investigated, a new air interface using millimeter waves, called *New Radio* (NR), has been chosen to provide higher throughputs, in addition to an evolution of the traditional LTE air interface. However, using high frequency bands will also lead to higher losses and poor coverage. Ultra dense networks, a paradigm in which the number of access nodes exceeds the number of active users, is seen as a cornerstone of 5G and one solution to counter the propagation issues. In addition, a feature allowing users to connect to multiple access nodes at the same time, Multi-Connectivity (MC), can be used to benefit from the advantages of both air interfaces.

In this work, we investigate multi-connectivity's effects on the network performance within an ultra-dense deployment. A special interest is given to energy efficiency, and how to ensure that multi-connectivity is more energy efficient than single connectivity. A condition on the power consumption is given and several secondary cell association algorithms are proposed and tested by simulation.

2. SECONDARY CELL ASSOCIATION

The cell association comprises the cell selection algorithm and the connection procedure. The focus here is on the selection logic as the procedure is standardized.

2.1 Condition on Power Consumption

One possibility to ensure energy efficiency for multi connectivity is to envision a metric related to the power consumption of target cells. In this work, we propose an inequality comparing the hypothetical power consumption if a downlink transmission for a specific user is made using the traditional master cell or associating a secondarycell, such as:

$$P_{C|M}(t) \geq P_{C|S}(t) \tag{1}$$

Where $P_{C|M}(t)$ is the power consumed by the system knowing that the transmission has been done by the master cell, while $P_{C|S}(t)$ represents the system consumption when transmitting with a secondary cell. By using a model for power consumption, such as the EARTH E^3 framework [1], and by assuming that the non-related traffic is constant, we can transform the equation (1) as:

$$\begin{aligned} &\Delta P_{\text{C,master}}(t) + P_{\text{TX,u|M}}(t) * \Delta_{\text{P,master}} \\ &\geq \Delta P_{\text{C,sec}}(t) + P_{\text{TX,u|S}}(t) * \Delta_{\text{P,sec}} \end{aligned} \tag{2}$$

Where

$$\Delta P_{\text{C,BS}}(t) = \begin{cases} P_0 - P_{\text{sleep}}, & \text{BS was sleeping} \\ 0, & \text{otherwise} \end{cases}$$

Where $\Delta P_{\text{C,BS}}(t)$ is the evolution of the power consumption between $t-1$ and t, and represents the cost of turning on a sleeping cell, $P_{\text{TX,u}}$ the emitted power for transmitting data related to the user u, $\Delta_{\text{P,BS}}$ a coefficient defined in the power model and P_{sleep} the static power consumed when a cell is sleeping.

*The research reported here was supported and funded by the PERCCOM Erasmus Mundus Program of the European Union [3]. It was conducted at Ericsson Research, Sweden, and supervised by Karl Andersson (Luleå University of Technology), Mårten Ericson and Mats Nordberg (Ericsson Research).

Table 1: Multi-Connectivity Schemes Parameters

Schemes	Perf.	Robst.	Engy.	Uti.
Max Bitrate	x			
Max SINR		x		
Max Bitrate-EE	x		x	
Max Clustered-Bitrate	x		x	x
AHP	x	x	x	x

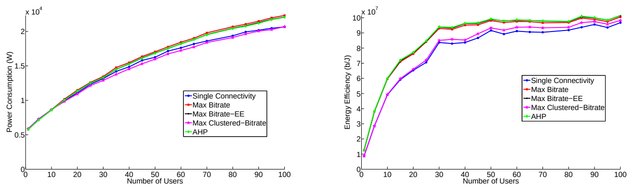

Figure 1: Power Consumption (a) and Energy Efficiency (b)

2.2 Multi-Connectivity Algorithms

A framework for secondary cell association is proposed. As for handovers, ping pong effects must be avoided at all costs. To do so, we implement a counter mechanism, called *Time To Trigger* (TTT), to counter the noisiness of the measurements. Furthermore, we decouple the procedure for connection and disconnection and use a hysteresis on the measured metrics.

Numerous metrics can be taken as input for an algorithm. We classified them in four categories:

- **Reliability**: this class regroups the most common metrics, such as RSRP, RSSI or the SINR, and are used to estimate the channel quality and optimize the reliability.
- **Performance**: If estimating the achievable bitrate is feasible, the expected performance can be quantified.
- **Energy**: the power consumption can be seen as one metric when designing energy efficient algorithms. Good power consumption models must however be implemented within the scheme to offer the best results.
- **Cell utilization**: Metrics such as the cell load, the number of connected devices or the BS state can also be used as inputs.

We propose five algorithms, from which two can be used for optimizing user reliability and four for performance optimization. Table 1 summarizes in which classification their inputs are taken from.

3. SIMULATION RESULTS

The performances of our proposals are investigated through simulations. The network is modelled as an ultra-dense deployment of 3 tri-sectors LTE base stations and 61 NR small cells, while users are modelled as a linear mob with a velocity of 3 km/h. The traffic is a full buffer, FTP-type communication and is modelled following a Poisson process.

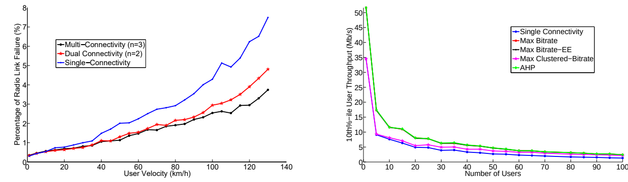

Figure 2: User Reliability (a) and 10^{th} Percentile User Throughput (b)

3.1 System Level

Figure 1 represents the power consumption in [W] and the energy efficiency in [bit/J] of the network. In Fig. 1.a, we can see that multi connectivity induces a increase of consumed power. Max-Bitrate, which does not implement the energy condition, see an increase of up to 4%, while Max-Bitrate-EE increases of up to 2% only. We can also see that one algorithm even reduces the power consumption compared to a single connectivity scenario by around 1%.

Regarding energy efficiency, we can first observe in Fig. 1.b that MC is always more energy efficient than the single connectivity scenario. This is due to the disconnection parameter chosen in our algorithms. An increase of around 15 to 20% is noted for the best three algorithms.

3.2 User Level

In Figure 2, the reliability is presented as the rate of radio link failure, obtained through simulations with different user velocities. We can observe that aggregating multiple links does indeed reduce the RLF rate, or increase the reliability, by up to 50% at high speeds, while there is little to no improvement at low speeds.

As expected, we can also observe that the user throughput is increased with MC, since a user can request more resources. We see an improvement of around 80% except for Max Clustered-Bitrate.

4. CONCLUSIONS

We showed that multi connectivity is a important feature of 5G as it will increase the system's energy efficiency given right parameters in ultra-dense networks. Furthermore, the results demonstrated that the increase in power consumption can be tackled by implementing our proposed energy condition. The interest of MC for the users was also proven for both their reliability and throughput. However, future work should assess the impact of MC over the terminal user's power consumption. It is expected that a more punctual usage of MC would benefit more users than maintaining multiple links, especially for low network load.

5. REFERENCES

[1] G. Auer and V. Giannini et al. How much energy is needed to run a wireless network? volume 18, pages 40–49, oct 2011.

[2] Ericsson AB. *Ericsson Mobility Report*, nov 2016.

[3] A. Klimova, E. Rondeau, K. Andersson, J. Porras, A. Rybin, and A. Zaslavsky. An international masterś program in green ICT as a contribution to sustainable development. *Journal of Cleaner Production*, 135:223–239, nov 2016.

Incremental Checkpointing Techniques for Transiently Powered Computers

Extended Abstract

Saad Ahmed
RWTH Aachen University, Germany
Lahore University of Management Sciences (LUMS), Pakistan

ABSTRACT

We propose three different incremental checkpointing techniques for transiently powered computers, to minimize the size of checkpoint updates in the secondary storage. These approaches lie in different regions of the cost-benefit spectrum, thus offering application-specific alternatives to transiently powered systems for successful completion of tasks under intermittent power.

1 INTRODUCTION

Miniaturization of transiently powered computers has made it difficult for a conventional battery to fit in such devices. This limitation forces these devices to harvest their energy from environment such as wind, solar and vibration. Harvested energy is generally erratic and has a lot of variance[1, 2]. With the passage of time, these devices are performing complex tasks which demands more time and energy. With erratic energy supply, they will not be able to complete their task in a single activation cycle. Checkpointing computational state (registers, global variables, and call stack etc.) before power blackout and restoring it at the start of next activation cycle (aka. *intermittent computing*) is thus essential to allow these transiently powered devices to resume, and not restart, the previously running computations.

Major challenge associated with checkpointing solutions is the size of checkpoint. A very large checkpoint size means the device would spend more time saving the state than doing the actual task. Ideally, it should be as small as possible. Recent state-retention solutions for transiently powered computers [2, 3] are suboptimal; they checkpoint complete program state, either the whole memory [3] or at least its occupied portion [2], each time a call to a checkpointing system is made due to depleting energy buffer. These approaches are unable to track program state that did not change from the previous checkpoint. Due to this reason, these approaches copy unmodified state to the secondary storage wasting time and energy.

Goal of transiently powered computing solution must be to "smartly" copy program state to secondary storage. Energy should not be wasted in copying the state that has not changed from the previous checkpoint or is never modified by the code. TPC solutions should reduce the amount of state to be saved as much as possible without excluding any important segments. This will save energy which can be used in actual computations.

To this end, we propose three different, platform independent *incremental checkpointing* approaches that can proactively track changes in the computational state.

2 INCREMENTAL CHECKPOINTING

Our first approach is near-optimal, as it accurately tracks and records modifications in the main memory except for processor registers. The second approach avoids such computational overhead by binding variables to program paths, only updating the relevant variables in the secondary storage if the corresponding path has been executed. The third approach does not require program path to variable binding rather it efficiently identifies modified memory locations using an approach called Hash of Hashes(HoH).

2.1 Inch: Tracking Changes in State

This approach is based on the fact that there are very few, well defined statements in the program, such as assignment, increment, shift operations, and function calls and returns, which modify program state. By tracking all those instructions, we can keep track of the changes made in program state from one checkpoint to the next. These statements are tracked by instrumenting the code to insert special function calls before these statements. So whenever state-modifying statements will be executed, they will get recorded and our approach will precisely update only those memory locations which were modified from the previous checkpoint.

Figure 1a shows the checkpoint sizes for different intervals of checkpoint. It is clear from the figure that, with exponential increase in checkpointing interval, the checkpoint size increases very slowly highlighting the fact that this approach offers light weight checkpoints. However, these light weight checkpoints come at the cost of computational overhead due to continuous tracking of program state. It increases the execution time of a task which implies more energy consumption for the same task. This increase in energy consumption is compensated by the energy saved by reduction in checkpoint size which is significantly larger than the energy consumed by functional calls to record changes.

2.2 EVM: Event to Variable Mapping

Our second approach avoids any computational overhead by binding variables to events through offline, static program analysis. For each event occurred, this approach identifies all variables modified by the program path of that event. This event can be an interrupt, such as radio, sensor, a timer or any other period or aperiodic task. This approach logs all the events occurred from one checkpoint to the next one. At each checkpoint call, size of variables modified by each event make up the checkpoint size. This approach removes the computational overhead created by the previous approach. Figure 1b shows the behavior of checkpoint sizes for different intervals. One can see that checkpoint size for this approach is greater than that for Inch. It is because of its inability to locate, at which point

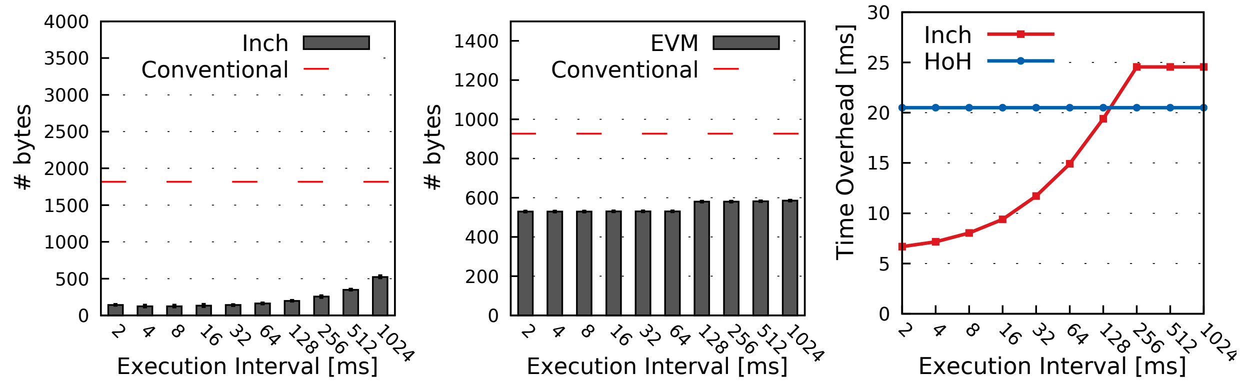

(a) RSA using Inch **(b) Sense using EVM** **(c) RSA using HoH**

Figure 1: Checkpoint sizes for different approaches

in path, did the checkpoint call occurred. It has to checkpoint all the variables being modified in the path where number of variables actually modified can be less than the ones checkpointed.

2.3 HoH: Hash of Hashes

Naive approach to identify modified memory locations requires reading all memory locations and comparing each byte of a memory to that in the checkpoint. Clearly, it is a time consuming process when memory is sufficiently large. To avoid this overhead, we propose our third approach to divide the memory into chunks of equal size. Each chunk is called a block. For each memory block, we compute a hash value.These hash values serve as the leaf nodes of the tree. For two (or more depending upon the type of tree) consecutive memory blocks, we again compute the hash value. This processes is repeated to construct a tree with root node representing the hash of the entire main memory. If nothing changed from the previous checkpoint, this approach will have to do only one comparison with the hash of the root node thus avoiding the need to read all blocks.

This approach has a similar behavior in terms of checkpoint size reduction to Inch. However, it reduces the time to checkpoint the program state by efficiently identifying what has changed from the previous checkpoint. While it can efficiently calculate the modified program-state, limitation of HoH is the regeneration of the entire tree before comparing it with the existing checkpoint. It incurs a constant overhead as shown in the figure 1c. With increase in interval length, computational overhead of inch increases as more number of state tracking function calls get executed between two checkpoint calls.

First two approaches are implemented through platform independent, pre-compiler extensions that automatically instrument the source code with relevant incremental checkpointing functionality while the third one uses simple SHA-256 hash.

3 CONCLUSIONS AND FUTURE WORK

We have already implemented Inch, EVM in TinyOS and HoH in mbed and have performed preliminary evaluations. In future, we plan to perform rigorous evaluations of these approach to demonstrate their true potential. Furthermore, we would also like to perform hardware and software optimizations which includes the use of byte addressable FRAM and relocation of variables in memory.

In case of Inch, we would like to reduce the checkpoint size even further by monitoring call stack to copy only the changed portion. Another challenge is the pointer that points to a global variable. Local variables are created on stack and get deleted when an activation record instance of the function is removed. However, if local variable is a pointer that points to global variable, it should also be tracked to correctly identify of state modifications.

In case of EVM, we have to test our approach on Contiki as well and compare it with the results of TinyOS. For HoH, we would like to compare the computational overhead of all three approaches with this approach. By making these comparisons, we would like to answer the question, "Which approach is best among all?" or "Is there any fit-for-all solution?". These question can only be answered by extensive analysis of these approaches which makes our future work.

REFERENCES

[1] N.A.Bhatti et al. 2014. Sensors with Lasers: Building a WSN Power Grid. In *Proc. 13th Int. Symp. Information Processing in Sensor Networks (IPSN '14)*. 261–272.
[2] N.A.Bhatti et al. 2016. Efficient State Retention for Transiently-powered Embedded Sensing. In *Proc. 2016 Int. Conf. Embedded Wireless Systems and Networks (EWSN '16)*. 137–148.
[3] Ransford et al. 2011. Mementos: System Support for Long-running Computation on RFID-scale Devices. In *Proc. 16th Int. Conf. Architectural Support for Programming Languages and Operating Systems (ASPLOS XVI)*. 159–170.

Intermittently-Powered Executor

Amjad Yousef Majid
Delft University of Technology
Delft, The Netherlands
a.y.majid@tudelft.nl

ABSTRACT

Enabling battery-free devices is a mandatory step towards an environment friendly Internet of Things (IoT). However, removing the batteries requires IoT to operate on an harvested power supply. This make sustaining long computation very challenging. Correspondingly, there are two approaches to enable long-running computations on intermittently-powered IoT: (i) checkpointing, where the volatile state of a program is frequently saved to the non-volatile memory, and (ii) task-based approach, where a programmer splits the code into small idempotent modules. Results show that task-based approach performs better than checkpointing. However, the use of tasks requires that the energy to execute any task must not exceed the maximum limit of the energy buffer. In other words, the task-based model is static approach that does not take advantage of variation in the amount of the harvested energy. Therefore, we introduce Intermittently-Powered Executor which pushes the boundaries of the task-based execution model by enabling an application to enlarge its task size, on the fly, when there is redundant energy.

CCS CONCEPTS

• **Computer systems organization → Embedded software**; *Availability*; • **Software and its engineering** → *Scheduling*; *Embedded software*;

KEYWORDS

Energy harvesting, transient operation, operating system

1 INTRODUCTION

Intermittently powered devices (IPDs) have unique features and pose significant challenges. They are battery-free devices which makes them light, cheap, require less maintenance and harvest energy from the environment. For example WISP (Wireless Identification and Sensing platform) [4] uses the power of RF signal to drive its computation and communication. However, reliance of IPDs on the real-time harvested power, makes sustaining long computations very difficult. Therefore, applications that run on IPDs require a software execution model that complies with the nature of a discontinuous power supply.

The intermittent (discontinuous) execution model defines a program execution as a cumulative discrete process. The main difference between the intermittent and the conventional (or continuous) execution models is that, a power failure is seen by the continuous model as an *exception* that may reset the progress of a program to its beginning. Whereas, in intermittent execution a power failure is regarded as a temporary *pause* to the execution that may result in some progress degradation. Generally we can classify the intermittent execution models into:

- *Sequential Execution Model*: Under this model a program is seen as one big idempotent task that has one common execution context. Generally, the progress of the program is saved and updated by means of checkpointing—where all the program context (e.g. CPU registers, the stack and the global variables) is saved to a non-volatile memory. Normally, the sequential model relies on a hardware assistant to measure the voltage level in the energy reservoir to place a checkpoint [1, 3]. The benefit of this model is that it does not require code modification by a programmer. However, it has a number of drawbacks: (i) it suffers from significant overhead [2]; (ii) the programmer should not access the non-volatile memory to guarantee the consistency of the memory [2]; and (iii) it restricts the IPDs to run only a single application.
- *Modular Intermittent Execution Model*: At the heart of this model is the concept of an idempotent task. The idempotent task is a function that does not have arguments and does not return a value. This task uses a well defined interface to interact with the non-volatile memory. Therefore, it tolerates arbitrary number of power interrupts. This model, generally, produces less overhead [2] and allows multiple applications to run on an IPD by interleaving their tasks. However, it obviously requires code modification—for example, if an algorithm is written according to the continuous execution model it has to be split by a programmer into small tasks to run under the modular intermittent execution model.

Despite the superiority of the Modular Intermittent Execution model, it is still a static approach that completely depends on a programmer's estimation which is mostly result in a sub-optimal code division. Moreover, this model can only consider a single hardware configuration and it does not take environment changes into considerations.

2 INTERMITTENTLY-POWERED EXECUTOR

Intermittently-Powered Executor (IPE) is runtime library that facilitates tasks navigation and preserves data/memory consistency of IPDs. IPE aims at reducing energy consumption and applications execution time. It optimizes the commit rate (saving tasks contexts into the non-volatile memory), subject to the number of power interrupts and completed tasks. Furthermore, IPE provides a number of services to facilitate intermittently powered applications development:

(1) *Power Interrupt Immune Scheduler*: Any scheduler that facilitates intermittent execution must make a firm transition from one task to another, and resume to the same task when

the supplied power is interrupted. The scheduler is responsible for estimating least commit rate given that forward progressing is preserved. IPE introduces a *persistent round-robin scheduler* to schedule task execution and commit operations. The benefit of round-robin scheduler in intermittent execution relies on the fact that it does not require protecting a large number of state variables to be a power interrupt immune scheduler. Additionally, round-robin scheduler makes adapting the commit rate relatively easy. Furthermore, IPE protects tasks' state information (task address and accessing permissions) in a (*persistent linked list*.

(2) *Persistent-/Non-persistent Semaphore*: Semaphore normally defines the relationship between producers and consumers tasks. IPE identifies two different types of relationships between the tasks:
   - time-dependent synchronization where the relationship is only defined within the same power interval. In other words, a power failure breaks the synchronization of these tasks;
   - content-dependent tasks synchronization; An example of such a relationship can be defined as *consume the data only once*. In such a scenario a power loss does not corrupt the data, but the data should not be consumed twice.

(3) *Communication between Tasks*: IPE provides two types of mailbox for communication between tasks, namely, a persistent and non-persistent mailboxes. Since the value of the communicated information may depend on the time or the content, IPE has to provides these two different communication services.

(4) *Memory Management*: In order for the IPE to change the commit rate, it requires an interface that controls the non-volatile memory access. IPE memory interface includes a volatile buffer (a buffer in the volatile memory) to hold temporarily all the outputs of the tasks (see Algorithm 1 line 4). Moreover, these tasks must first attempt to read variables from the buffer before the non-volatile memory (see Algorithm 1 lines 5–8) to ensure a correct execution and the consistency of the memory. Once the scheduler decides to commit the volatile data, it first copies all the values from the volatile buffer to a persistent buffer (see Algorithm 1 lines 11-16). If the power is interrupted (before IPE finishes copying all the data from the volatile buffer to the persistent one) the persistent buffer must be discarded and the execution must start again from the location of the last commit. The second phase commit is a power failure immune operation and up on finishing it the non-volatile memory will be in a consistent state (see Algorithm 1 lines 17-21).

Algorithm 1 Virtualized Operational Buffer

```
1:  var ∈ {global variables}
2:  TASKS()
3:  while executing do                          ▷ Execution stage
4:      var → volatile buffer
5:      if var in volatile buffer then
6:          var ← volatile buffer
7:      else
8:          var ← FRAM
9:      if power interrupts then
10:         back to 2
11: while volatile buffer ≠ ∅ do                ▷ First phase commit
12:     volatile buffer → persistent buffer
13:     if power interrupts then
14:         discard persistent buffer
15:         back to 2
16:
17: while persistent buffer ≠ ∅ do              ▷ Second phase commit
18:     persistent buffer → FRAM
19:     if power interrupts then
20:         Continue
21:
22: return
```

3 SUMMARY

This poster introduces IPE: a power interrupt immune library that facilitates a dynamic modular execution model. It enables execution optimization on the run by adapting the commit rate. It uses the history of execution as a metric for the optimization process. Moreover, it provides a number of fundamental services to simplify the intermittent applications development cycle. Finally, IPE left the door open for further development specially when the I/O operation is considered.

REFERENCES

[1] Domenico Balsamo, Alex S. Weddell, Geoff V. Merrett, Bashir M. Al-Hashimi, Davide Brunelli, and Luca Benini. 2015. Hibernus: Sustaining computation during intermittent supply for energy-harvesting systems. *IEEE Embedded Syst. Lett.* 7, 1 (March 2015), 15–18.

[2] Alexei Colin and Brandon Lucia. 2016. Chain: tasks and channels for reliable intermittent programs. In *Proc. OOPSLA*. ACM, Amsterdam, The Netherlands, 514–530.

[3] Benjamin Ransford, Jacob Sorber, and Kevin Fu. 2012. Mementos: System support for long-running computation on RFID-scale devices. In *Proc. ASPLOS*. ACM, Newport Beach, CA, USA, 159–170.

[4] Alanson P. Sample, Daniel J. Yeager, Pauline S. Powledge, Alexander V. Mamishev, and Joshua R. Smith. 2008. Design of an RFID-based battery-free programmable sensing platform. *IEEE Trans. Instrum. Meas.* 57, 11 (Nov. 2008), 2608–2615.

Verified Boot in IoT Devices with Low Power Consumption

Christos Profentzas
Chalmers University of Technology
Gothenburg, Sweden
chrpro@chalmers.se

ABSTRACT

In this paper, we describe our ongoing research regarding the security of operating systems for IoT devices. We try to highlight energy consumption issues posed by security measures. We start by securing the device boot-up process to provide the necessary dependency towards a trustful operating system. Lastly, our focus is a holistic view of the security model, which combines security measures and energy consumption in IoT devices.

CCS CONCEPTS

• **Security and privacy → Embedded systems security**; *Trusted computing*;

KEYWORDS

IoT, Verified Boot, Trusted Boot, Secure Boot, Power Consumption

1 INTRODUCTION

The modern field of *Internet of Things* (IoT) has emerged from the evolution of Wireless Sensor Networks (WSNs). The explosion of technologies and protocols has made applications beyond the simple data gathering available. For instance, an adaptive lighting in road tunnels [1] is possible using WSN as part of a closed-loop control system.

To extend and use such applications in the wider area of IoT, we need to add more capabilities to the embedded devices. Therefore, we have seen several Operating Systems (OS) emerge to manage the growing number of resources. In the survey of Padmini Gaur & Mohit P. Tahiliani [3], we find an extensive comparison among the most recent operating systems for IoT devices. Although the performance of the device is an essential factor of success in the field of IoT, the security of the device is another factor equivalent in importance.

Tock-OS [4], an embedded operating system, has already made a significant effort to address operating system security issues. The developers wrote the kernel of this particular OS in a new system programming language called RUST [8], where the compiler checks the safety of the memory during the compiling phase. Therefore, it avoids numerous security flaws. However, the user can only trust an IoT device if the device has booted an authenticated operating system in the first place. In this way, we must also secure the boot method of the device.

Even though, the security aspect of the boot process has already been highlighted in embedded devices for safe-critical applications such as Avionics Wireless Networks (AWN) [7], securing the boot process for IoT device has not received as much attention yet. The main reasons are extreme limitations regarding the power consumption and latency requirements for IoT devices.

Thus, we investigate the overall power consumption caused by security measures during the boot process. From our experience, we have seen a significant compromise between dealing with real-time constraints and implementing security measures in the boot process. Therefore, our hypothesis estimates a similar compromise between the power consumption and the security level we require to achieve.

Finally, it is within our research priorities to investigate a holistic security model that ensures all parts of the device from the boot process to operating system and finally the application level. Our research goals aim towards achieving a reliable platform for IoT devices that take energy consumption into consideration.

2 SECURITY THREATS

The main reason for a user to not trust the Operating System is the risk of malicious modification to the Kernel. The term Rootkit [6] covers these modifications, which is a set of tools designed to maintain privileged access to a compromised Operating System. How an attacker can gain privileged access to the system is outside of our scope, and we do not analyze it further. Therefore, we assume that the attacker has already found a way to compromise the system. Additionally, attackers try to hide their malicious software in deeper operating system structures, ultimately targeting the boot process and the startup code of the device. These advanced tools are covered by the term Bootkit [5].

Finally, the tools mentioned above are not the only methods for modifying the operating system for malicious purposes. For example, Cloaker [2] is a dynamic way to undermine the normal execution of the operating system running on ARM embedded systems. Therefore, it is within our research scope to construct security models to capture any limitations of security measures.

3 VERIFIED BOOT OVERVIEW

In general, we can say that a Verified Boot mainly provides a report (verification) about the authenticity of the boot-up code and the OS kernel regarding unauthorized modifications. However, we can find different terms used in various context.

Secure Boot [7] is another term used to describe such a verification method, which requires the device to verify before loading the operating system. However, we must not confuse the extent of the guarantees provided by the validation technique. The validation results only report modifications of the operating system; this states nothing about the trustworthiness of the verification process.

In this way, we need to introduce a trust between the verification process and the device. We can issue trust by measuring certain device configuration properties before and after the verification process. Also, the measurements could be validated by a third party upon a request.

The above architecture requires external hardware to provide the necessary trust. This entity already exists in the form of Trusted Platform Module (TPM) [7]. This module is a passive reporting device, and the enforcement policy is open to suggestions. Moreover, TPM can play the role of a root of trust, where we maintain the validation data. The aim is to create a chain of trust from the boot of device to the operating system and eventually the application layer. Figure 1 illustrates the boot process using TPM as a root of trust.

One of the challenges that our research is trying to address is the minimal hardware requirements. It is important to modify to a minimum the existing architecture of IoT platforms as cost and energy consumption should be kept low. Another challenge is the limitations of introducing a verification method into the low-level boot code, due to the small and compound memory space that such code occupies. Lastly, we have set the primary objective of minimizing power consumption that may conflict with the verification process. In this way, we should take into consideration the possibility of a compromise between security measures and power consumption.

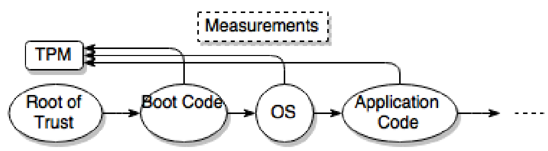

Figure 1: A Trusted Boot Process using TPM

4 CONCLUSIONS

In this paper, we introduce our ongoing research and interests regarding the security of IoT devices. Moreover, we argue that security is an essential factor in the success of IoT devices; as the ever-growing number of security threats make the importance of the security more urgent than ever. Any security measure should take into consideration the power consumption of the device, which is also another significant factor in the sustainability of IoT devices and their success. Our research focuses on addressing the challenges of a fully trusted IoT platform with low-power consumption.

REFERENCES

[1] M. Ceriotti, M. CorrÃă, L. D'Orazio, R. Doriguzzi, D. Facchin, S. T. GunĂČ, G. P. Jesi, R. L. Cigno, L. Mottola, A. L. Murphy, M. Pescalli, G. P. Picco, D. Pregnolato, and C. Torghele. 2011. Is there light at the ends of the tunnel? Wireless sensor networks for adaptive lighting in road tunnels. In *Proceedings of the 10th ACM/IEEE International Conference on Information Processing in Sensor Networks.* 187–198.

[2] Francis M. David, Ellick M. Chan, Jeffrey C. Carlyle, and Roy H. Campbell. 2008. Cloaker: Hardware Supported Rootkit Concealment. In *Proceedings of the 2008 IEEE Symposium on Security and Privacy (SP '08).* IEEE Computer Society, Washington, DC, USA, 296–310. https://doi.org/10.1109/SP.2008.8

[3] Padmini Gaur and Mohit P. Tahiliani. 2015. Operating Systems for IoT Devices: A Critical Survey. In *Proceedings of the 2015 IEEE Region 10 Symposium (TENSYMP '15).* IEEE Computer Society, Washington, DC, USA, 33–36. https://doi.org/10.1109/TENSYMP.2015.17

[4] Amit Levy, Michael P. Andersen, Bradford Campbell, David Culler, Prabal Dutta, Branden Ghena, Philip Levis, and Pat Pannuto. 2015. Ownership is Theft: Experiences Building an Embedded OS in Rust. In *Proceedings of the 8th Workshop on Programming Languages and Operating Systems (PLOS '15).* ACM, New York, NY, USA, 21–26. https://doi.org/10.1145/2818302.2818306

[5] X. Li, Y. Wen, M. H. Huang, and Q. Liu. 2011. An Overview of Bootkit Attacking Approaches. In *2011 Seventh International Conference on Mobile Ad-hoc and Sensor Networks.* 428–431. https://doi.org/10.1109/MSN.2011.19

[6] XiangYu Li, Yi Zhang, and Yong Tang. 2015. Kernel Malware Core Implementation: A Survey. In *2015 International Conference on Cyber-Enabled Distributed Computing and Knowledge Discovery, CyberC 2015, Xi'an, China, September 17-19, 2015.* 9–15. https://doi.org/10.1109/CyberC.2015.26

[7] K. Markantonakis, R. N. Akram, and R. Holloway. 2016. A secure and trusted boot process for Avionics Wireless Networks. In *2016 Integrated Communications Navigation and Surveillance (ICNS).* 1C3–1–1C3–9. https://doi.org/10.1109/ICNSURV.2016.7486322

[8] Nicholas D. Matsakis and Felix S. Klock, II. 2014. The Rust Language. In *Proceedings of the 2014 ACM SIGAda Annual Conference on High Integrity Language Technology (HILT '14).* ACM, New York, NY, USA, 103–104. https://doi.org/10.1145/2663171.2663188

Modulation Schemes in Ambient Backscatter Communication

[Extended Abstract]

Oliver Harms
Chalmers University of Technology
Gothenburg, Sweden
mail@oliverharms.eu

ABSTRACT

This paper presents a low-cost backscatter system using off-the-shelf components. It uses simple modulation schemes and is usable with a constant carrier as well as with ambient television signals as carrier.

The evaluation of the system shows that ranges of up to 225 m are coverable in line-of-sight environments using a constant carrier. Furthermore, it is usable with ambient TV signals with bit error rates as low as 10^{-3}. These results show that this backscatter system outperforms state-of-art backscatter systems.

KEYWORDS

Backscatter communication, Internet of Things, Wireless

1 INTRODUCTION

In the world of Internet of Things (IoT) low power communication is more and more important. One way to reduce power needs to a bare minimum is backscatter communication. However, many backscatter systems use expensive hardware.

This project presents a backscatter system with sensor tags using different modulation schemes to transmit data receivable by off-the-shelf radio chips. Figure 1 gives an overview over such a system. For studying the best possible performance of this setup, a constant carrier is needed generated by an Software Defined Radio (SDR). But the ultimate goal is to even remove this carrier and use ambient TV signals for communication. Those signals are available almost everywhere and would therefore be the optimal carrier for a cheap backscatter solution.

2 BACKGROUND

Backscatter communication. Backscatter communication is a passive communication method not requiring to actively generate RF signals. Instead the radar cross-section (RCS) of the antenna of the backscatter tag is modulated to either absorb or reflect a carrier signal. This modulation of the RCS is performed by changing the impedance of the antenna circuit between two states.

To communicate on a frequency different to the frequency of the carrier signal, a shifting operation in addition to the RCS modulation is performed. This is achieved by multiplying the carrier signal with a square wave generated by the circuit of the backscatter tag. The frequency of this square wave is the required offset Δf, between the carrier frequency f_c and the desired frequency. The backscattered signal has the form

$$2 \sin ft \sin \Delta f_c t = \cos(f_c - \Delta f)t - \cos(f_c + \Delta f)t.$$

This means that the backscattered signal is shifted to the frequencies $f_c + \Delta f$ and $f_c - \Delta f$ [3, 6].

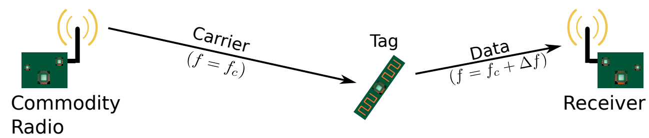

Figure 1: Backscatter System. Shift of carrier signal to avoid self interference. Data encoding on shifted signal.

TV signals. Terrestrial television signals are available in the most places and transmitted continuously. Therefore they are theoretically usable as a carrier signal. The terrestrial TV signals available in Europe and many other places worldwide are DVB-T signals using the signal format OFDM with a maximum bandwidth of 8 MHz [2, 4].

3 DESIGN

The design presented in Figure 1 consists of three parts. The carrier generator, a SDR in this project, a backscatter tag and a receiver.

The receiver used is a CC2500 transceiver of Texas Instruments. Its task is to receive and decode the backscattered data. It supports different modulation schemes like Frequency-Shift-Keying (FSK) and On-Off-Keying (OOK) with bitrates between 1.2 and 500 kbps [5].

The backscatter tag consists of the Analog Devices RF switch HMC190BMS8 as its front-end and a BeagleBone Black as its back-end. The BeagleBone has a programmable real-time unit (PRU) and is therefore able to execute time-critical software. This makes it possible to backscatter data using FSK and OOK with a frequency offset of over 2 MHz and a deviation for FSK of ∼95 kHz without using an FPGA [1].

4 EVALUATION

The evaluation begins with experiments in a line-of-sight outdoor environment using a constant carrier. From the first experiments on FSK showed much better performance than

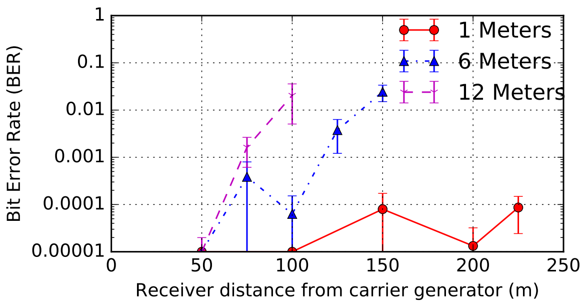

Figure 2: Low bitrate (2.9 kbps) with tag close to the carrier generator (outdoors). High range achievable by placing the backscatter tag close to the carrier generator.

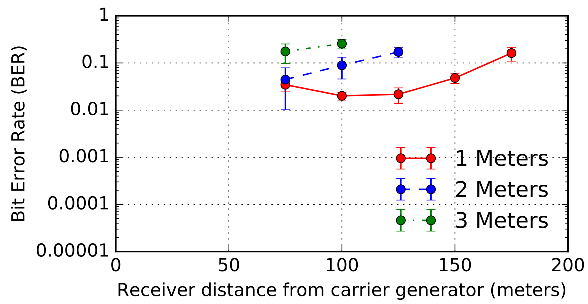

Figure 3: High bitrate (197 kbps) with tag close to the carrier generator (outdoors). The achievable range is smaller than for low bitrates and more errors are introduced.

OOK which was expected. Therefore experiments evaluating the maximum possible performance of the system were performed for FSK only.

Figure 2 shows the results for short distances between the carrier generator and the backscatter tag using a bitrate of 2.9 kbps. It shows that distances of 225 m, 150 m or 100 m between the carrier generator and the receiver are reachable depending on the tag placement. In the most cases the BER is below 10^{-2}. Figure 3 shows results for a high bitrate of 197 kbps. Even for this bitrate distances of up to 175 m are coverable with small distances between the carrier generator and backscatter tag. The better performance for a placement of the tag close to the carrier generator is common for backscatter systems and an explanation of it can be found e.g. in [3].

To study the performance of TV signals as carrier, a DVB-T signal was recorded and replayed in the 2.4 GHz ISM band. This was necessary to compare both carrier systems using the same receiver. For this comparison both modulation schemes OOK and FSK were used. Due to the low signal strength of maximally −40dBm of the replayed signal, only short distances were coverable and therefore the comparison was done based on the signal-to-noise ratio. Table 1 shows the bit error rate for different signal to noise ratios using FSK or OOK with a constant carrier or a TV signal as carrier. The results are that OOK with a constant carrier and both modulation schemes with a TV signal as carrier show overall comparable performance whereas FSK with a constant carrier shows much better performance. Generally speaking, a TV signal as carrier is usable especially for SNR values of 16 or more.

SNR	0	8	16
FSK (constant)	0.01	0.0001	0.00001
OOK (constant)	0.1	0.01	0.001
FSK (TV)	0.1	0.08	0.003
OOK (TV)	0.2	0.1	0.001

Table 1: Bit error rate of OOK and FSK for different carrier signals and signal-to-noise ratios.

5 CONCLUSION & FUTURE WORK

This paper presents a backscatter system which is capable of achieving ranges beyond 200 m with off-the-shelf components using a constant carrier. Moreover, it shows that using TV signals as carrier is possible with bit error rates below 10^{-2}.

In future work, a sub-GHz receiver should be used to build a backscatter system working directly with TV signals and evaluating possible ranges coverable with such a system.

REFERENCES

[1] beagleboard.org. 2015. PRU-ICSS Resources. (2015). Accessed 18-07-2017. http://beagleboard.org/pru.

[2] electronicsnotes. [n. d.]. What is DVB-T: Digital Video Broadcasting. ([n. d.]). Accessed 18-07-2017. https://www.electronics-notes.com/articles/audio-video/broadcast-tv-television/what-is-dvb-t-basics-tutorial.php.

[3] Bryce Kellogg, Vamsi Talla, Shyamnath Gollakota, and Joshua R. Smith. 2016. Passive Wi-Fi: Bringing Low Power to Wi-Fi Transmissions. In *NSDI 16*. 151–164. https://www.usenix.org/conference/nsdi16/technical-sessions/presentation/kellogg

[4] Vincent Liu, Aaron Parks, Vamsi Talla, Shyamnath Gollakota, David Wetherall, and Joshua R. Smith. 2013. Ambient Backscatter: Wireless Communication out of Thin Air. In *Proceedings of the ACM SIGCOMM 2013 Conference on SIGCOMM (SIGCOMM '13)*. 39–50. https://doi.org/10.1145/2486001.2486015

[5] Texas Instruments. [n. d.]. CC2500. ([n. d.]). Accessed 18-07-2017. http://www.ti.com/lit/ds/symlink/cc2500.pdf.

[6] Pengyu Zhang, Dinesh Bharadia, Kiran Joshi, and Sachin Katti. 2016. HitchHike: Practical Backscatter Using Commodity WiFi. In *Proceedings of the 14th ACM Conference on Embedded Network Sensor Systems CD-ROM (SenSys '16)*. 259–271. https://doi.org/10.1145/2994551.2994565

Collaborative and Environmentally-Powered Sensors and Actuators for Smart Environments

Lars Hanschke
Hamburg University of Technology
Research Group smartPORT
lars.hanschke@tuhh.de

ABSTRACT

Smart cities and smart production plants pose novel demands on sensor networks than years ago such as latency or heterogeneity of sensor nodes. These issues are tackled by fog computing. However, a joint energy consumption consideration of collaborative harvesting devices as part of fog computing adds a new layer of complexity. Furthermore, the number of gateways to ensure interoperability between vendors and radio communication protocols increases. Thus, we present a suitable WiFi energy harvesting platform, which can be used without additional gateways, and show a solar testbed for future evaluations. Upcoming work introduces a mathematical model aiming at distributing the energy consumption within the network. This leads to decreased drop-out rates of single nodes while fulfilling a joint task.

1 INTRODUCTION

Connected devices are now being integrated in every 'thing' in our life. Smart home, smart production plants, even smart cities or smart wearables are just a few examples for gathering data and improving processes. Especially for smart cities, application scenarios are various: increasing maintenance efficiency by monitoring infrastructure or making living healthier by monitoring pollution indicators. The state ministry of the city of Hamburg agreed on paving the ground for new smart transportation systems making use of smart roads or even autonomous public transport vehicles [2]. Usually, a variety of gateways is needed to connect smart devices from different vendors or with different wireless transmission protocols to the Internet [8]. Within their digitalization plans, most of the cities offer or plan to offer public WiFi networks. This allows smart devices equipped with WiFi modules to be connected directly, without the need for additional infrastructure.

Traditional sensor networks mainly consist of devices with equal capabilities in computing power and energy resources. They are configured once, potentially adjusting their duty cycle but primarily following the same task: sensing data and transmitting it to the sink and potentially to the cloud. With new fields of deployment evolving, demands on sensor networks change: for real-time applications, the communication path to the cloud and back again might be too long to make time-critical decisions. For example, autonomous vehicles need to react instantly to accidents and cannot rely on data processing in the cloud. This rises the need for a local decision or controlling instance, which is commonly treated as fog computing [1]. Furthermore, fog computing considers heterogeneous devices: differentiated by their computing or sensing capabilities and also energy resources.

The field of energy harvesting adds a new layer of complexity for joint working in fog computing. The available resources might differ highly between devices: e.g. local shadowing of solar harvesting devices changes their overall and timely distributed energy budget completely [7]. Consequently, their ability to perform tasks varies, since they need to adjust their duty cycle. To ensure the whole sensor network is able to fulfill the envisaged task, a coordination instance is needed. This instance, e.g. a smart edge node, can coordinate the activity in the fog network depending on the on-line capabilities of each sensor also considering its location. Therefore, an efficient exchange of status information, e.g. battery level, and a resulting re-allocation of computing or sensing resources is needed. E.g. the fine particle pollution in a production plant is continuously monitored only by a sub-set of sensor nodes to decrease the energy consumption of the network. In case of an emergency accident, previously idle sensor nodes nearby the accident can be instructed to help gaining information about the accident: either by increasing spatial resolution or by switching on additional sensors to detect different physical phenomena.

We argue, that WiFi-equipped energy harvesting sensor nodes play a great role in improving smart environments, which we explain in the following.

2 WORK-IN-PROGRESS

On our way to energy harvesting WiFi sensor nodes, we made first steps in two major compartments. First, we show the potential of WiFi sensor nodes and second, we introduce improvements of our energy harvesting platform.

2.1 Why WiFi Helps

In fog networks, the role of smart edge nodes gets more important. To reduce the amount of data being transferred to the cloud or to decrease reaction time on events, more computing power is needed than regular gateways can offer. Modern WiFi routers offer high computational power, e.g. a ASUS AC2400 router is equipped with a 1.2 GHz dual-core CPU. Consequently, software updates can enable the already deployed infrastructure to pre-process data or decide instantly. Even small WiFi devices like the Espressif ESP 8266 offer a 32 bit MCU with 160 MHz and can dynamically switch roles between gateway and regular WiFi station. We illustrated challenges faced while working with low-cost WiFi-enabled sensors like Arduino Nano in [3]. Now, we are able to show that operation of ESP 8266 is feasible with a solar-driven energy harvester. Our harvester uses a solar panel producing up to 250 mA at 2.7 V and stores the energy inside a super-capacitor with 400 F at 2.7 V. We display our observations in Fig. 1. The algorithm presented in [7] predicts the amount of harvested energy based on previous days and adjusts the duty cycle of the node.

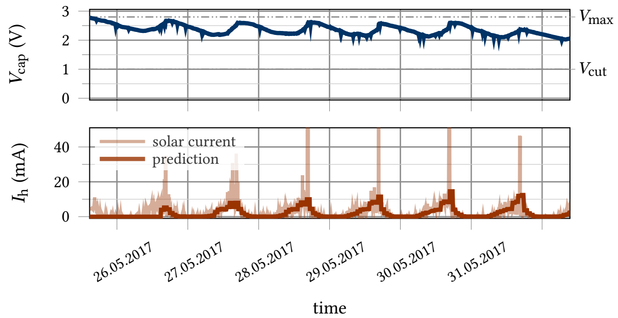

Figure 1: V_{cap}, harvested solar current I_{h} and predicted current of 16-day long-term test (extract).

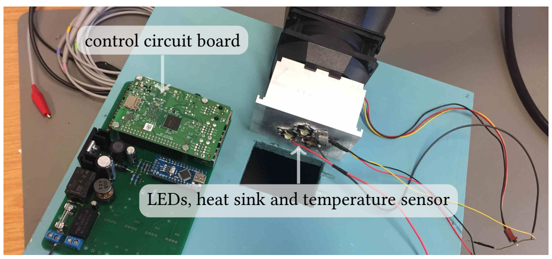

Figure 2: Light box with high-power LEDs; generates 75 mA **with our harvester, expandable to** 250 mA.

2.2 Energy Harvesting

For our future investigations on energy harvesting algorithms, we need reproducible solar current conditions. The authors in [5] also recognized that energy harvesting algorithms lack reproducibility because of changing and not repeatable environmental conditions. Therefore, we developed a light box, which is displayed in Fig. 2 and explained extensively in [4].

We use high-power LEDs to replay solar current traces emulating a current of up to 75 mA with our solar cell. Additionally, the box is expandable for generating up to 250 mA with a 75 cm^2 panel. With costs below € 110, we plan to produce more boxes to build a testbed for nodes experiencing different energetic conditions.

To improve our duty cycle and harvest prediction algorithm, we implemented the strategy presented in [6] for our WiFi platform. Using conventional on-line weather forecasts to adjust the radio duty cycle of the node. We combine local knowledge of past harvesting conditions, e.g. shadowing due to buildings, with a prediction of future harvesting conditions relying on cloud cover of the sky.

3 CONTRIBUTION & FUTURE WORK

With our previous investigations, we provide a platform for performing hardware tests of our future steps. Our next step is to develop a mathematical model to optimize the overall energy consumption in the network based on current and future energy harvesting capabilities. Given is a set of sensor nodes with unequal sensing and computing capabilities and energy harvesting conditions. The smart edge node has knowledge about the position of each sensor node and is instructed to offer a service to a cloud user. Figure 3 illustrates this scenario. The edge node decides which

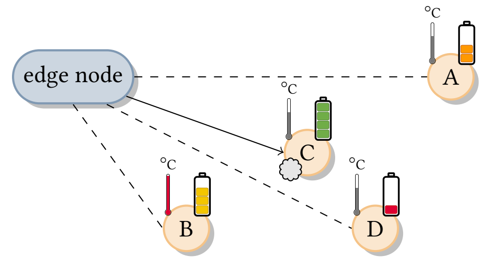

Figure 3: Heterogeneous network; node B **detects temperature** $T > T_{\text{thres}}$**; spatial resolution increasable by activating** C **or** D**.** D **has low battery, so** C **is activated; smoke sensor of** C **also activated to detect fires.**

additional nodes to activate for reacting on the temperature threshold exceeding. Up to now, the consideration of energy harvesting into routing or joint task scheduling is based on past knowledge of harvested energy. However, a consideration based on prediction of future energy income is still missing. The problem is as follows: due to the fluctuating nature of environmental resources, the harvesting capabilities might change quickly. Routing metrics taking past energy harvesting capabilities into account risk depletion of energy storage of nodes with good prior harvesting capabilities. Integrating prediction of future energy income into these metric can lead to a substantially improved lifetime in the network.

We plan to formulate this optimization problem which is mainly determined by the set of nodes to activate without loosing the availability of the whole network. This involves efficient exchange of status information and control commands as well as fulfilling latency requirements. Further steps include validation of this model, simulations and hardware tests using our solar harvesting platform and light box. Additionally, real-life test in conjunction with the Hamburg Port Authority are planned.

REFERENCES

[1] Flavio Bonomi, Rodolfo Milito, Jiang Zhu, and Sateesh Addepalli. 2012. Fog Computing and Its Role in the Internet of Things. In *Proc. of the First Edition of the MCC Workshop on Mobile Cloud Computing (MCC).*

[2] Free and Hanseatic City of Hamburg - State Ministry of Economic Affairs, Transport and Innovation. 2016. Verkehr 4.0 - ITS-Strategie für Hamburg. http://www.its2021.hamburg/downloads/ITS%20Strategie%20Hamburg.pdf. (May 2016).

[3] Lars Hanschke, Jan Heitmann, and Christian Renner. 2016. Challenges of WiFi-Enabled and Solar-Powered Sensors for Smart Ports . In *Proc. of the 4th International Workshop on Energy Harvesting and Energy-Neutral Sensing Systems (ENSsys).*

[4] Lars Hanschke, Christian Renner, Jannick Brockmann, Tobias Hamann, Jannes Peschel, Alexander Schell, and Alexander Sowarka. 2017. Light Insight — Emulation of Radiation Traces for Analysis and Evaluation of Solar-Harvesting Algorithms. In *Proc. of the 5th International Workshop on Energy Neutral Sensing Systems (ENSsys).* Accepted for publication.

[5] Josiah Hester, Timothy Scott, and Jacob Sorber. 2014. Ekho: Realistic and Repeatable Experimentation for Tiny Energy-Harvesting Sensors. In *Proc. of the 12th ACM Conference on Embedded Network Sensor Systems (SenSys).*

[6] Christian Renner. 2013. Solar Harvest Prediction Supported by Cloud Cover Forecasts. In *Proc. of the 1st International Workshop on Energy Neutral Sensing Systems (ENSsys).*

[7] Christian Renner, Stefan Unterschütz, Volker Turau, and Kay Römer. 2014. Perpetual Data Collection with Energy-Harvesting Sensor Networks. *Transactions on Sensor Networks* 11 (2014).

[8] Thomas Zachariah, Noah Klugman, Bradford Campbell, Joshua Adkins, Neal Jackson, and Prabal Dutta. 2015. The Internet of Things has a Gateway Problem. In *Proc. of the 16th International Workshop on Mobile Computing Systems and Applications (HotMobile).*

Designing Reliable Transient Applications

Andres Gomez, Luca Benini, Lothar Thiele
ETH Zurich
Zurich, Switzerland
{gomeza,lbenini,thiele}@ethz.ch

ABSTRACT

State-of-the-art wireless sensor systems are typically performance-constrained, battery-based devices which can, at most, reach self-sustainability using energy harvesting and aggressive duty-cycling. In certain scenarios, however, the energy availability is such that no large scale storage device is necessary to fulfill application requirements. Transient systems, which operate in an energy proportional way, are good candidates for such applications because they can reduce both the transducer area and storage device required for functional correctness, thus minimizing the systems cost and form factor. Designing transient systems which guarantee optimal program progress and electrical efficiency requires novel hardware-software design technique to address specific challenges. By using an Energy Management Unit (EMU), designers can abstract away transducer specifics and still guarantee maximum power transfer between the source/EMU and EMU/load.

KEYWORDS

Energy Harvesting, Transient Systems, Wireless Sensor Networks.

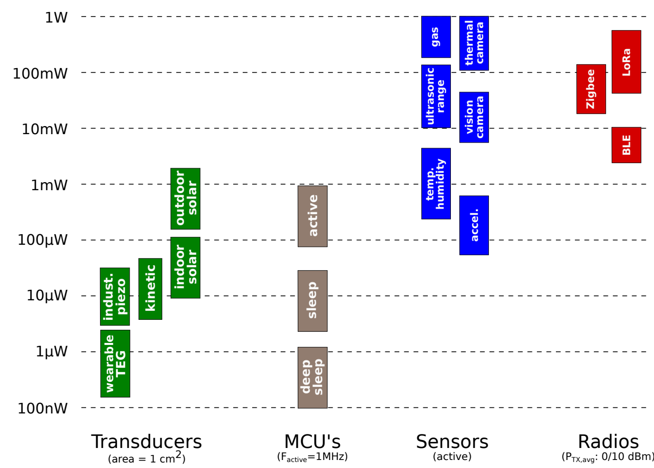

Figure 1: Power density ranges for common transducers, sensors, microcontrollers and transceivers.

1 INTRODUCTION

The advances in ultra-low power design over the past decades have significantly extended the lifetime of battery-powered devices. However, the billions of devices in the emerging Internet of Things (IoT) will demand for deploy-and-forget installations with virtually unlimited lifetimes. Battery-based designs can conceivably have long lifetimes, but they can be very expensive, bulky, energetically inefficient, and significantly hinder large scale deployments. Energy harvesting, though a mature technology, has only been successfully deployed in large scale systems where size is not a limiting factor. In many scenarios, such as wearable systems, a source's physical dimensions are just as important as the power levels [1, 8].

Transiently powered systems operate efficiently in adverse harvesting conditions, requiring only limited storage capacity and input power to reliably execute power-hungry applications. The recently-proposed Energy Management Unit (EMU) [3] allows a transient system to operate in an energy-proportional manner. As opposed to traditional duty-cycling, an EMU-based system is energy driven, meaning that as the available energy increases, so does application's execution rate. This would allow devices to operate reliably and efficiently in the wide input power range found in typical wearable sources: solar panels, thermo-electrical generators (TEG's) and piezo-electric.

Different energy sources can have widely ranging power levels depending on the environmental conditions, but they are almost always below the levels required for sensing and transmitting information, as shown in Fig. 1. Even when a transducer is large enough to directly power a load in the correct voltage and current range, there is no guarantee that it will harvest at its maximum power point [10]. If the load adjusts its operating point to extract the maximum power from the load, it will most likely not be the load's optimal operating point, which depends only on the application and the peripherals it uses, not the environmental conditions. EMU-based designs leverage decoupling to guarantee that the maximum power is harvested and the minimum power is consumed, simultaneously.

Vision sensors, which acquire and process images, are typically power-hungry devices which require substantial computational resources. For this reason it was not until recently that it became feasible to have batteryless vision sensors in a wearable form factor, thanks to new paradigms in energy harvesting systems [4, 9]. In this work we describe the design and implementation of an energy-opportunistic, wearable vision sensor node capable of executing computationally intensive tasks with temporal dependencies [5]. In particular, we consider the example of a solar-powered vision sensors for wearable applications. These transient sensor nodes have the property of guaranteed information and energy availability, since darkness does not provide neither energy nor information and light provides both.

2 EMU-BASED APPLICATION DESIGN

Transient systems must be able to tolerate highly volatile sources and still guarantee program progress. In order for these systems to operate reliably and efficiently, they have to accumulate harvested energy until enough is available for the execution of one single atomic task, also called a burst. Afterwards, the system should be shut down completely until enough energy is accumulated for the next burst execution. The time interval between two bursts depends on the instantaneous input power. This type of operation directly leads to three challenges for the design of transiently powered systems:

- Constraint (1) **Minimum Energy Guarantee** The energy harvester cannot directly power the system. To guarantee the execution of atomic tasks, the storage device should provide this *minimum energy availability*.
- Constraint (2) **Temporal Independence** There is no control over the length of the time interval between two bursts, since this only depends on the currently available input power. The application needs therefore to be split into *separate bursts with no temporal dependencies*.
- Constraint (3) **Non-Volatility** Between two bursts, the system is shut down and peripherals are powered off. Therefore, if an application cycle requires several bursts *non-volatile memory* (NVM) technologies to retain the system's state between bursts. Even if an application cycle fits in a single burst, logging data requires NVM.

Fig. 2 shows the prototype mounted in a user's glasses. The device performs three main tasks: *Image Acquisition*, *Processing* and *Storage*. To overcome *Constraint (1)*, the application-specific requirements must be used to determine the EMU's energy burst size. Since motion estimation algorithms require at least two successive images to detect pixel displacements, each burst needs to guarantee at least the energy required to acquire two pictures, thus satisfying *Constraint(2)*. The *Storage* tasks copies the estimated velocity to an external FRAM memory for logging. Thus, the simplest solution to *Constraint (3)* is to group the entire application in a single burst. Based on experimental characterizations, it was determined that the entire applications cycle consumes around 1.3 mJ, requiring only a 150 uF capacitance for the transient system to work. The prototype was then connected to a 42 cm^2 flexible solar panel, and exposed to varying light levels. The system's energy efficiency as well as the resulting execution rate can be seen in Fig. 3. The energy efficiency is defined as the ratio between the energy consumed by the application to the harvested energy.

3 CONCLUSION AND FUTURE WORK

In this work, we have argued that batteryless systems can offer all the necessary guarantees to build reliable sensor applications, even with high-power peripherals. As opposed to battery-based

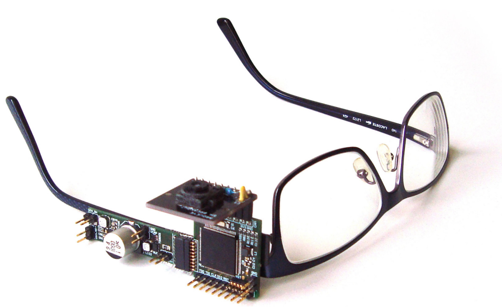

Figure 2: Wearable transient vision node [5].

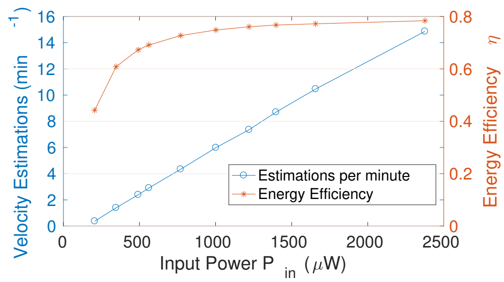

Figure 3: Measured energy efficiency and execution rate as a function of the input power [5].

devices, these systems are energy-driven and thanks to their energy-proportionality, they can operate efficiently even in very variable and adverse harvesting conditions.

Thanks to EMU-based design, we have designed an application that can reliably and efficiently acquire and process images to estimate the user's walking speed in a wide variety of harvesting scenarios. Our proposed vision sensor has an average active power consumption of 6.85 mW, but requires only 100's of μW's to begin estimating the velocity. Furthermore, it can reach up to 5.8 velocity estimations per second, and has a motion estimation error of 1.4 % of the distance traveled.

The concept of the EMU-based design is a promising approach to build reliable IoT applications with minimized energy storage and harvesting requirements. While the EMU has bridged the gap between small, volatile energy sources and power-hungry sensor nodes, there are many possibilities for improvement. In particular, cpu-based tasks have more flexibility during execution, since they can be arbitrarily paused and resumed. This approach is orthogonal to the EMU and has been explored in different works, such as [2, 6, 7] and can be used to reduce the energy required for processing.

Acknowledgements

This research was funded by the Swiss National Science Foundation under grant 157048: Transient Computing Systems.

REFERENCES

[1] NA Bhatti et al. 2016. Energy Harvesting and Wireless Transfer in Sensor Network Applications: Concepts and Experiences. *ACM Trans. Sens. Netw.* (2016).
[2] NA Bhatti et al. 2017. HarvOS: efficient code instrumentation for transiently-powered embedded sensing.. In *IPSN*. 209–219.
[3] A Gomez et al. 2016. Dynamic Energy Burst Scaling for Transiently Powered Systems. In *Proc. DATE Conf.* EDA Consortium, 349–354.
[4] A Gomez et al. 2017. Efficient, Long-Term Logging of Rich Data Sensors using Transient Sensor Nodes. *ACM Transactions on Embedded Computing Systems*.
[5] A Gomez et al. 2017. Wearable, Energy-Opportunistic Vision Sensing for Walking Speed Estimation. In *Proc. SAS Symp.* IEEE, 1–6.
[6] H. Jayakumar et al. 2014. QUICKRECALL: A Low Overhead HW/SW Approach for Enabling Computations across Power Cycles in Transiently Powered Computers. *Proc. Int. Conf. VLSI Design* (2014). https://doi.org/10.1109/VLSID.2014.63
[7] B. Lucia et al. 2017. Intermittent Computing: Challenges and Opportunities. In *LIPIcs-Leibniz International Proceedings in Informatics*, Vol. 71. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik.
[8] M. Magno et al. 2016. InfiniTime: Multi-sensor wearable bracelet with human body harvesting. *Sustainable Computing: Informatics and Systems* (2016).
[9] S. Naderiparizi et al. 2015. WISPCam : A Battery-Free RFID Camera. In *Proc. IEEE RFID*.
[10] L Sigrist et al. 2017. Measurement and validation of energy harvesting IoT devices. In *Proc. DATE Conf.* IEEE, 1159–1164.

Making LPWANs Batteryless

Vijay S Rao
Delft University of Technology, The Netherlands
v.rao@tudelft.nl

ABSTRACT

With the advent of low power miniaturized electronics supporting high-end computations and advances in wireless technologies, it is certain that Internet of Things (IoT) has become one of the key technological enablers for smart-* systems in this decade. As an enabler for the Cyber Physical Systems, IoT has made a great impact on our lifestyle and the way we interact with our surroundings, environment, and even machines. IoT infrastructure comprises of embedded devices, and network of such devices having sensors and actuators, leveraging existing wired or wireless infrastructures for communication and control of different electronic systems. The current trend is to connect these devices directly to a base-station over long ranges while adopting low-power communication methodologies as well. These are termed as Long-range Low-Power Wide Area Networks (LPWANs).

Unfortunately, merely replacing the batteries with energy harvesters does not provide the necessary alternative. Ambient energy sources do not provide constant power, and the harvested energy from the sources varies drastically over location and time. The harvested energy, sometimes, is very low and sometimes in excess of the storage capacity of the nodes. One of the challenges is that the devices die and re-enter networks due to fluctuations in harvested energy. This becomes aggravated when the devices are completely batteryless. Consequently, energy harvesting in these devices necessitates a redesign of algorithms, communication techniques, and network protocols to achieve perpetual operations while satisfying the application requirements. The idea, therefore, is to use transient powered computing methodologies in enabling LPWANs.

KEYWORDS

Internet of Things, LPWAN, Batteryless, Transient powered computing

1 INTRODUCTION

The Internet of Things (IoT) is one of the disruptive technologies in today's connected world. The idea is to connect every thing to the Internet. It holds the key to many current and future technologies that will significantly influence the quality and sustainability of life. The vision of IoT is to enable large-scale monitoring and/or control in order to either observe a phenomenon or to automate tasks. Many novel IoT applications are fueling an exponential growth in the deployment of embedded devices. These devices equipped with sensors and/or actuators with wireless communication capabilities are central in realizing the IoT infrastructure. These devices must have small form factor in order to be portable, deployable and economical. Therefore, these devices are resource constrained with respect to the available power, computing and memory.

Many early adopters of IoT used short- range communication for their sensors requiring multiple hops before reaching the Internet cloud. However, recent advances have enabled a simpler technological solution called Low Power Wide Area Networks (LPWAN) that offer low energy communication over distances of multiple kilometers, making sensors send data to the Internet backbone over only one hop.

LoRaWAN (Long Range Wide Area Network) is one of the several competing LPWAN technologies with amongst others SigFox, NBIoT and Weightless [3, 7]. LoRaWAN has been the most successful of these technologies in providing an easily accessible LPWAN [4]. This protocol is being developed by the LoRa Alliance [1, 2], which is an open alliance. In comparison with other LPWAN technologies, LoRaWAN claims an inexpensive, secure and power efficient communication method for applications with small end-devices that have to send small amounts of data over large time intervals. Therefore, we consider LoRaWAN as the specific LPWAN technology to work with. However, we wish to extend the philosophy of this work to other LPWAN technologies.

Although IoT (or sensor) devices are required to last for a long time, batteries limit the lifetime of the devices, therefore that of the network and the applications. Powering all the IoT devices through batteries is neither scalable nor environmentally sustainable. Frequent battery replacement is labor intensive in some cases; in many other situations, battery replacement is impractical due to physical or deployment conditions.

Therefore, we adopt ambient energy-harvesting techniques. By tapping into the harvesting opportunities in the ambiance, the nodes gain autonomy with respect to energy. While many existing works consider recharging batteries with the ambient energy-harvesting techniques, we look to eliminate batteries completely and replace them with more sustainable energy storage buffers such as supercapacitors.

Unfortunately, merely replacing the batteries with energy harvesters does not provide the necessary alternative. Ambient energy sources do not provide constant power, and the harvested energy from the sources varies drastically over location and time. The harvested energy, sometimes, is very low and sometimes in excess of the storage capacity of the nodes. One of the challenges is that the devices die and re-enter networks due to fluctuations in harvested energy. Consequently, energy harvesting in these devices necessitates a redesign of algorithms, communication techniques, and network protocols to achieve perpetual operations while satisfying the application requirements.

Backscatter promises to be an extremely low power, smaller and cheaper alternative to active radios [5, 6]. They are inexpensive due to the lack of active radio components, consume three to four orders of magnitude lower power and peak current than radios and hence can operate with emerging flexible, printed and ultra thin printed battery technologies. However, despite all these benefits, current backscatter designs have seen very limited adoption beyond RFID applications due to its limited range. Therefore, we wish to investigate the possibility of using backscatter for long range communications. We briefly describe the challenges in the next section.

2 CHALLENGES

We list some of the important challenges in employing ambient backscatter for LoRaWAN communications (based on [8]).

- *Operating range.* Radios including Wi-Fi, Bluetooth and ZigBee operate up to 100s of meters while wide area LoRa and SigFox deployments extend operation to kilometers. However, existing backscatter solutions such as RFID and Passive Wi-Fi are limited to tens of meters of operating distance in best-case scenarios. *The question here is how to extend the range of LoRa to at least a kilometer's range?*
- *Power consumption.* Radios including Wi-Fi, BLE, ZigBee, Lora and SigFox all consume between 10 to 500 mW, which is 3 - 4 orders of magnitude higher than the power consumption of backscatter systems including RFID and Passive Wi-Fi. As a key feature of LoRaWAN is low-power consumption, it is a challenge to reduce the power consumption further in order to survive for a long time while guaranteeing performance similar to active radio based LoRaWAN.
- *Sources of power.* We consider only ambient energy-harvesting techniques. By tapping into the harvesting opportunities in the ambiance, the nodes gain autonomy with respect to energy. While many existing works consider recharging batteries with the ambient energy-harvesting techniques, we look to eliminate batteries completely and replace them with more sustainable energy storage buffers such as supercapacitors. The sources include sun, flow-based such as water and wind, electromagnetic waves, vibrations and mechanical movements (including human movements). Ambient energy sources do not provide constant power, and the harvested energy from the sources varies drastically over location and time. Therefore, the challenge is to make the devices work with the available power no matter how low it is.
- *Transient operation.* Since the system we wish to design is batteryless and that the devices die and re-enter networks due to fluctuations in harvested energy, the devices must adopt transient powered computing techniques. The traditional method of rebooting each time may not satisfy the application requirements as the nodes may go through too many reboots due to lack of energy.
- *Communications.* The most important aspect of the LoRaWAN nodes is to communicate. Since we are looking at backscatter, it is *highly challenging to adhere to the standards of LoRa PHY layer (Chirp spread spectrum based modulation) while employing backscatter.*
- *Cost and size.* While backscatter radios are said to have small form factors due to the elimination of active RF components, it is still important to keep an eye on the size and cost of the devices.

REFERENCES

[1] LoRa Alliance. 2015. LoRaWAN: What is it? (2015). https://www.lora-alliance.org/portals/0/documents/whitepapers/LoRaWAN101.pdf [Online].
[2] LoRa Alliance. 2015. LoRaWANâĎć Specification. *LoRa Alliance* (2015).
[3] Martin Bor, John Edward Vidler, and Utz Roedig. 2016. LoRa for the Internet of Things. (2016).
[4] Marco Centenaro, Lorenzo Vangelista, Andrea Zanella, and Michele Zorzi. 2016. Long-range communications in unlicensed bands: The rising stars in the IoT and smart city scenarios. *IEEE Wireless Communications* 23, 5 (2016), 60–67.
[5] Vikram Iyer, Vamsi Talla, Bryce Kellogg, Shyamnath Gollakota, and Joshua Smith. 2016. Inter-technology backscatter: Towards internet connectivity for implanted devices. In *Proceedings of the 2016 conference on ACM SIGCOMM 2016 Conference*. ACM, 356–369.
[6] Bryce Kellogg, Aaron Parks, Shyamnath Gollakota, Joshua R Smith, and David Wetherall. 2014. Wi-Fi backscatter: Internet connectivity for RF-powered devices. In *ACM SIGCOMM Computer Communication Review*, Vol. 44. ACM, 607–618.
[7] Paul J Marcelis, Vijay S Rao, and R Venkatesha Prasad. 2017. DaRe: Data Recovery through Application Layer Coding for LoRaWAN. In *Internet-of-Things Design and Implementation (IoTDI), 2017 IEEE/ACM Second International Conference on*. IEEE, 97–108.
[8] Vamsi Talla, Mehrdad Hessar, Bryce Kellogg, Ali Najafi, Joshua R Smith, and Shyamnath Gollakota. 2017. LoRa Backscatter: Enabling The Vision of Ubiquitous Connectivity. *arXiv preprint arXiv:1705.05953* (2017).

[Extended Abstract] Backscatter Communications

Carlos Pérez-Penichet
Uppsala University
Sweden
carlos.penichet@it.uu.se

ABSTRACT

Backscatter communication is poised to become the go-to mechanism to network many embedded devices. The low power consumption of this mechanism makes it attractive as an alternative to batteries and their associated drawbacks by enabling the devices in question to be powered by harvested energy. Harvested energy, however, may be insufficient to continuously power the device and may run out at any time. As a consequence these devices must operate only during the short, irregular intervals when power is available. This situation brings a new set of challenges both from the point of view of computation and sensing as well as from the point of view of wireless communication protocols.

My research focuses on the challenges associated with communications. Specifically, my focus is on integrating this class of devices in existing standard-compliant wireless networks and to tackle the associated challenges at the protocol level.

CCS CONCEPTS

• **Computer systems organization** → **Sensor networks**; • **Hardware** → **Wireless devices**; • **Networks** → *Network architectures*;

KEYWORDS

Low-power wireless, Backscatter communication, IoT

1 INTRODUCTION

Many applications in the context of the Internet of Things are expected to run unattended for extended periods of time that can reach many months or years. Oftentimes these devices need to be installed in inaccessible places such as embedded in concrete or implanted in a living creature. This situation creates a problem for maintenance and energy supply; batteries become a liability. Furthermore, communications are responsible for the largest part of the energy consumption of traditional low-power sensor nodes. For these reasons, energy is one of the most important resources these devices have and needs to be managed very carefully.

Backscatter communication, a technique that is virtually free in term of energy consumption, has become very attractive for battery-free sensing devices. This technique enables devices to exploit different sources to harvest energy for the devices' operation and to communicate with other sensors [6]. As a consequence, nodes have the potential to operate battery free for extended periods of time.

Battery-free operation, however, comes with its own associated challenges. The rate at which energy can be collected is typically lower than the power drawn by the devices while operating and, at the same time, the supply of harvested energy is generally erratic and unpredictable. As a consequence devices are forced to operate in an intermittent manner. They execute their code during the time when there is enough energy available and fall back to a dormant state while there is not, only to resume operation when energy is available again.

This mode of operation is unprecedented in traditional computers and brings a number of challenges from the point of view of software development. The key challenge is that the state of execution must be saved reliably and efficiently before the power runs out. On the other hand, communications present their own set of particular challenges in multiple aspects across the communications stack. Furthermore, future transiently powered devices will be operating in conjunction with active wireless communications technologies such as WiFi or ZigBee. Integration of these two kinds of devices is a particularly interesting aspect that needs to be addressed.

My PhD studies center on the communication aspects of this technology. Specifically in efforts to bridge the gap between existing standard compliant wireless communication technologies and the future backscatter communication based devices.

2 BACKSCATTER COMMUNICATIONS

When a radio signal reaches an antenna, part of the signal is reflected back (scattered) into open space. The reflected power depends on the load impedance attached to the antenna. Backscatter communication exploits this effect to transfer information among devices by taking advantage of pre-existing signals such as TV or cellular stations or of purposely-generated signals such as an unmodulated carrier.

A backscatter transmitter will alternate between absorbing the received signal and reflecting it back. These two states can be considered as zeroes and ones in the information transfer, making it possible for a receiver to decode its transmissions. The facts that these devices do not need to actively generate the radio signals and that alternating between reflecting and non-reflecting states can be achieved with a single transistor, mean that they have extremely low power consumption. On the receiver side, very low energy consumption is also achieved by demodulating the received signal with very simple analog electronics.

3 RESEARCH CHALLENGES

Exploiting the backscatter phenomenon for virtually zero-power communication in wireless sensor networks is highly attractive. Further research in the area should help pave the way to successful applications of the technique. In particular, it would be worth exploring ways to improve communication range in backscatter communication and improving interoperability between backscatter devices and standard-compliant wireless communication devices.

To unleash the full potential of backscatter communication, its range should be increased so that it can surpass room level without compromising reliability and data rate. Previous research efforts [10, 13, 14] have shown good results in increasing communication range. They propose the use of different coding schemes inspired by spread spectrum techniques to increase the robustness and thereby also the range of communication. We have also explored this direction and a summary of our efforts is presented in section 4.

Future backscatter communication devices will operate in a transiently powered fashion while interconnected with other types of wireless networks such as WiFi, Bluetooth or ZigBee. Interoperation among these networks will be of great importance and there are significant research efforts showing progress in this direction [3, 5, 15]. In section 4, we describe our own efforts in this direction.

4 EARLY RESULTS AND NEXT STEPS

Inspired by previous works aiming at increasing the communication range of ambient backscatter systems [10], we have explored the trade-off between the spread spectrum code rate and the robustness of communication [12]. We asked whether increasing the code rate increases the throughput in ambient backscatter and concluded that under good channel conditions, it is more favorable to reduce the code rate (i.e., using less chips per symbol) rather than to increase the number of bits per symbol. While both approaches increase the throughput, the former provides better robustness.

We have explored the integration of backscatter devices in existing active networks. To that end we proposed a system that would allow augmenting an existing ZigBee network with new sensor devices without the need to alter the existing hardware [11]. Installing a new sensor would be as simple as placing a sticker, similar to an RFID tag, next to the existing active sensor nodes. These tags would be battery free and backscatter their sensor readings using the ZigBee protocol, thus there is no need for any modification to the pre-installed hardware.

Our future work includes looking further into medium access protocol problems associated with this kind of augmented networks.

5 RELATED WORK

Backscatter communications hold vast potential for batteryless sensing devices owing to a reduction of orders of magnitude in energy cost compared to conventional radio [2]. A limitation of these applications is the need for a high-powered reader. Nikitin et al. design a, low-cost reader for RFID applications [9] and demonstrate tag-to-tag communication at a range of a few centimeters [8].

Liu et al. present *Ambient Backscatter*, a shift in the design of backscatter communication systems [6] that leverages ambient RF, such as TV signals, both as a carrier for communication and source of energy. This strategy avoids the need for a separate reader and enables device-to-device communication. Further research efforts on ambient backscatter include the use of antenna cancellation to improve throughput and coding to improve range [10], enabling full-duplex backscatter communication [7] and using ambient WiFi signals as the carrier [1, 4, 15].

Many recent efforts demonstrate the ability to generate transmissions compatible with standard wireless protocols such as WiFi and Bluetooth [3, 5]. At the same time there have been successful efforts to increase the communication range of backscatter systems [13].

REFERENCES

[1] Dinesh Bharadia, Kiran Raj Joshi, Manikanta Kotaru, and Sachin Katti. 2015. BackFi: High Throughput WiFi Backscatter. *SIGCOMM'15* 45, 5 (Aug. 2015), 283–296. https://doi.org/10.1145/2829988.2787490

[2] Michael Buettner. 2013. *Backscatter Protocols and Energy-Efficient Computing for RF-Powered Devices.* Ph.D. Dissertation.

[3] Vikram Iyer, Vamsi Talla, Bryce Kellogg, Shyamnath Gollakota, and Joshua Smith. 2016. Inter-Technology Backscatter: Towards Internet Connectivity for Implanted Devices *(SIGCOMM '16)*. ACM, 356–369. https://doi.org/10.1145/2934872.2934894

[4] Bryce Kellogg, Aaron Parks, Shyamnath Gollakota, Joshua R. Smith, and David Wetherall. 2014. Wi-Fi Backscatter: Internet Connectivity for RF-powered Devices *(SIGCOMM '14)*. ACM, New York, NY, USA, 12. https://doi.org/10.1145/2619239.2626319

[5] Bryce Kellogg, Vamsi Talla, Shyamnath Gollakota, and Joshua R. Smith. 2016. Passive Wi-Fi: Bringing Low Power to Wi-Fi Transmissions *(NSDI'16)*. 151–164. https://www.usenix.org/conference/nsdi16/technical-sessions/presentation/kellogg

[6] Vincent Liu, Aaron Parks, Vamsi Talla, Shyamnath Gollakota, David Wetherall, and Joshua R. Smith. 2013. Ambient Backscatter: Wireless Communication out of Thin Air *(SIGCOMM '13)*. ACM, 39–50. https://doi.org/10.1145/2486001.2486015

[7] Vincent Liu, Vamsi Talla, and Shyamnath Gollakota. 2014. Enabling Instantaneous Feedback with Full-duplex Backscatter *(MobiCom '14)*. ACM, New York, NY, USA, 12. https://doi.org/10.1145/2639108.2639136

[8] P.V. Nikitin, S. Ramamurthy, R. Martinez, and K.V.S. Rao. 2012. Passive tag-to-tag communication. In *RFID*. https://doi.org/10.1109/RFID.2012.6193048

[9] Pavel V Nikitin, Shashi Ramamurthy, and Rene Martinez. 2013. Simple low cost UHF RFID reader. In *Proc. IEEE Int. Conf. RFID*. 126–127.

[10] Aaron N. Parks, Angli Liu, Shyamnath Gollakota, and Joshua R. Smith. 2014. Turbocharging Ambient Backscatter Communication *(SIGCOMM '14)*. ACM, 619–630. https://doi.org/10.1145/2619239.2626312

[11] Carlos Pérez-Penichet, Frederik Hermans, Ambuj Varshney, and Thiemo Voigt. 2016. Augmenting IoT Networks with Backscatter-enabled Passive Sensor Tags *(HotWireless '16)*. ACM, New York, NY, USA, 23–27. https://doi.org/10.1145/2980115.2980132

[12] Carlos Pérez-Penichet, Ambuj Varshney, Frederik hermans, Christian Rohner, and Thiemo Voigt. 2016. Do Multiple Bits per Symbol Increase the Throughput of Ambient Backscatter Communications? *(MadCom @ EWSN '16)*. ACM.

[13] Vamsi Talla, Mehrdad Hessar, Bryce Kellogg, Ali Najafi, Joshua R. Smith, and Shyamnath Gollakota. 2017. LoRa Backscatter: Enabling The Vision of Ubiquitous Connectivity. *arXiv:1705.05953 [cs]* (2017). http://arxiv.org/abs/1705.05953

[14] A. Varshney et al. 2017. LoRea: A Backscatter Architecture that Achieves a Long Communication Range. In *ACM SenSys 2017*.

[15] Pengyu Zhang, Dinesh Bharadia, Kiran Joshi, and Sachin Katti. 2016. HitchHike: Practical Backscatter Using Commodity WiFi *(SenSys '16)*. ACM, New York, NY, USA. https://doi.org/10.1145/2994551.2994565

[Extended abstract] Towards Backscatter Networks

Ambuj Varshney
Uppsala University
Sweden
ambuj.varshney@it.uu.se

ABSTRACT

Backscatter communication by reflecting or absorbing ambient wireless signals enables transmissions at several orders of magnitude lower energy cost when compared to conventional radios. Backscatter, till recently, has been assumed to be a low range communication mechanism with recent systems generating transmissions compatible with WiFi, ZigBee and BLE at a maximum communication range of tens of meters. Our recent system LoRea demonstrates that it is possible to achieve significantly higher communication range while consuming µWs of power. LoRea demonstrates a communication range of upto a km operating in bistatic mode with the tag colocated with the carrier source. High range is a key enabler to develop network of backscatter tags, my research tackles the challenges to enable wide-area network of backscatter tags.

(a) Backscatter tag and carrier generator are co-located (1 ft). (b) Backscatter transmissions received 1 km away from tag.

Figure 1: *Long distance backscatter communication.* We can receive backscatter transmissions even when the tag and the receiver are separated by a distance of 1 km. The tag consumes 70 µW during transmission and is colocated with the source of carrier signal.

1 INTRODUCTION

Over the past few years, there has been a significant interest to develop low power wireless networks that can achieve long communication ranges. Wireless standards designed for long communication ranges like SigFox, NB-IOT and LoRa are being widely deployed. A key feature of applications conceived using these standards is that they connect wireless sensors to a base station over distances of hundreds (or thousands) of meters. Long communication range is a key enabler for important applications like connected cities, precision agriculture or civil infrastructure monitoring. These applications often require large deployments of sensors. The latter are currently almost exclusively battery powered primarily due to the high energy cost of communication. For example, radios most commonly employed to enable long distance wireless links consume power in the order of mWs. While battery-free sensors operating on harvested energy are appealing, the amount of power available from ambient sources is usually orders of magnitude lower than that required to operate conventional radios.

Backscatter by reflecting ambient wireless signals enables wireless transmissions at a power consumption that is several orders of magnitude lower compared to traditional radios. Hence, backscatter can provide wireless connectivity to sensors operating on µWs of harvested power. Recent backscatter systems have demonstrated the ability to synthesise transmissions compatible with commodity wireless protocols like ZigBee [4, 6], WiFi [5] and BLE [2]. The use of commodity protocols enables novel applications such as connected implants [4]. They are, however, limited to a maximum communication range of tens of meters [5]. The short communication range is due to the high bitrates, and consequently, the wide bandwidth occupied by commodity protocols. Short communication range proves to be a hindrance for many of the application discussed earlier.

2 LONG RANGE BACKSCATTER COMMUNICATION

There have been two approaches that enable km long backscatter links at µWs of power consumption. The first uses backscatter transmissions with FSK as modulation scheme [9]. Second approach use LoRa transmissions that are based on a Chirp Spread Spectrum (CSS) modulation scheme [7]. The key advantage of these approaches is that they leverage commodity radio transceivers for receptions which *enables pervasive* deployment due to very low cost (few USD) when compared to receivers like software-defined-radios (SDRs).

LoRa-based systems achieve one of the highest sensitivity levels among different communication technologies [7] by using a CSS modulation scheme. This high sensitivity enables Talla et al. [7] to demonstrate the highest achieved range with backscatter communication of 2.8 km with tag colocated with carrier source. Further, they devise mechanisms built on CSS and the LoRa wireless standard to enable operations of multiple backscatter tags..

In comparison, FSK-based systems leverage narrow-bandwidth transmissions to increase receiver sensitivity and communication range [9]. Our system demonstrates a range of over 3.4 km by transmitting with a narrow bandwidth of 13 kHz, with carrier source located 1 m away from the tag. Even though LoRa achieves highest sensitivity levels among different radio technologies which helps achieve high communication range, LoRa does so at very low bitrates. At more practical bitrates, FSK offers similar sensitvity to LoRa[10]. Further, narrow bandwidth offers numerous advantages: *First*, it enables better usage of the spectrum, especially to support operations of multiple tags using FDMA schemes. *Second*, it results in a narrow receive bandwidth which significantly reduces out-of-band interference even at a small frequency offset. This further improves the reliability of weak backscatter transmissions and helps operate on spectrum shared with other devices. Hence, the

Table 1: Spreading codes used to improve ability to mitigate in-band interference, and also improve range.

Spreading mode	Symbol 0	Symbol 1
1X	0	1
2X	00	11
4X	1100	0011
8X	00111100	11000011

focus of my research is to use FSK backscatter to design a network with wide-area coverage.

3 RESEARCH CHALLENGES

We envision that the long range achievable with backscatter communication can enable wide-area networks of backscatter tags. Such a network would consist of a vast number of tags dispersed over a large geographical area, with multiple carrier generators providing the necessary carrier signal. The tags could communicate to each other through the carrier generators which mediate access to the medium [5, 7, 11]. Further, the tags could operate without batteries from minuscule amounts of energy harvested from the environment, e.g., using tiny solar cells. However, while the long communication range is a key enabler, there are several unsolved challenges that hinder widespread deployment. We next present our initial efforts to identify, and to solve, these challenges.

One of the key challenges to enable our vision is to make links from the tags to the backscatter receivers reliable and robust. This is difficult due to the extremely weak backscatter signals that are inherently prevalent when communicating over long distances. For example, we must be able to receive backscatter signals as weak as −124 dBm, which is 40 dB weaker than the weakest detectable signals for modern WiFi transceivers [1]. The challenge is that operations at such low signal strengths result in high bit errors.

The second key challenge to our vision is to mitigate the harmful effects of in-band-, and out-of-band interference, particularly from the carrier signal. State-of-the-art backscatter systems leverage heterodyning operations to keep carrier and backscattered signal apart in frequency to reduce interference from the carrier signal. Furthermore, the backscatter tags themselves operate in unlicensed spectrum that is shared with other devices, which also causes interference. Mitigating such interference is important to enable wide-area network.

The third challenge we identify is to support the operation of multiple uncoordinated tags at the same time. This is particularly challenging in wide-area deployments.

Final challenge we identify is to support medium access control (MAC) issues for backscatter tags. Existing receivers employed on backscatter tags are passive energy detectors that lack frequency selectivity to receive FSK transmissions [4]. Devising MAC layer under such constraints would be challenging.

4 OUR RECENT EFFORTS

We present overview of preliminary design which we present in [10] which takes a step forward to make our vision a reality.

Improving reliability and range. To improve the reliability of FSK backscatter systems, we implement widely used forward error correction (FEC) mechanisms [8]. We implement a 1/2 rate non-systematic non-recursive convolutional encoder with a constraint length of $K = 7$ and the following polynomials: ($g^{(0)} = [1, 0, 0, 1, 1, 1, 1]$, $g^{(1)} = [1, 1, 0, 1, 1, 0, 1]$). The code outputs two symbols for each data bit. We implement this particular encoder because of two reasons: *First*, convolutional codes have been widely used for almost half a century to enable reliable long distance wireless links [3]. Convolution codes can be easily generated using digital logic elements which helps power consumption at the tag to be µWs. *Second*, many radios, including the one we use, support convolution codes which helps to reduce complexity of the reception.

Bit spreading is commonly employed in protocols like WiFi and ZigBee to improve resilience to interference, and to improve the sensitivity level of the receiver due to the coding gain. To mitigate the harmful effects of in-band interference and improve range, we devise a spreading mechanism similar to DSSS. We devise a direct-sequence spread spectrum (DSSS) scheme with the spreading sequences shown in Table 1. We assign a known chip sequence to each transmitted symbol. For example, with a spreading mode of 4X, we spread the symbol 0 to the sequence 1100, and the symbol 1 to the sequence 0011. Spreading each of the symbols to a chip sequence further decreases the effective data rate. With a spreading factor of 8 and a chip rate of 3000 bit/s we can achieve data rate of 187.5 bit/s (=3000/8/2). The tag can select between different spreading modes based on the prevailing link conditions.

Supporting concurrent transmissions. FDMA schemes divide the spectrum into a number of channels, where each channel is allocated to a single link. The key challenge is to plan the frequency of the channels in such a way that they are separated enough to limit out-of-band interference from each other but close enough to use the spectrum efficiently.

We leverage the high the adjacent channel rejection ability of our transceivers to space individual channels at spacing of 50 kHz, slightly larger than the receiver filter bandwidth. At this spacing, we observe an out-of-band rejection ability of approx −40 dB. Each of these channels can be allocated to one or more devices. We devise such a scheme to support operation of multiple concurrently transmitting tags.

REFERENCES

[1] Texas Instruments. CC3200. 2017. Accessed: 01-04-2017. http://www.ti.com/lit/ds/symlink/cc3200.pdf. (2017).

[2] Joshua F Ensworth and Matthew S Reynolds. 2015. Every smart phone is a backscatter reader: Modulated backscatter compatibility with bluetooth 4.0 low energy (ble) devices. In *IEEE RFID, 2015*. IEEE, 78–85.

[3] Jerold Heller and Irwin Jacobs. 1971. Viterbi decoding for satellite and space communication. *IEEE Transactions on Communication Technology* 19, 5 (1971), 835–848.

[4] Vikram Iyer, Vamsi Talla, Bryce Kellogg, Shyamnath Gollakota, and Joshua Smith. 2016. Inter-technology backscatter: Towards internet connectivity for implanted devices. In *SIGCOMM 2016*. ACM, 356–369.

[5] Bryce Kellogg, Vamsi Talla, Shyamnath Gollakota, and Joshua R Smith. 2016. Passive wi-fi: Bringing low power to wi-fi transmissions.

[6] Carlos Pérez-Penichet, Frederik Hermans, Ambuj Varshney, and Thiemo Voigt. 2016. Augmenting IoT Networks with Backscatter-enabled Passive Sensor Tags. In *HOTWIRELESS 2016*. ACM, 23–27.

[7] Vamsi Talla, Mehrdad Hessar, Bryce Kellogg, Ali Najafi, Joshua R Smith, and Shyamnath Gollakota. 2017. LoRa Backscatter: Enabling The Vision of Ubiquitous Connectivity. *Ubicomp* (2017).

[8] David Tse et al. Fundamentals of Wireless Communication. In *CAMBRIDGE*.

[9] Ambuj Varshney, Oliver Harms, Carlos Perez Penichet, Christian Rohner, Frederik Hermans, and Thiemo Voigt. 2017. LoRea: A Backscatter architecture that achieves a long communication range. *ACM SenSys* (2017).

[10] Ambuj Varshney, Carlos Pérez-Penichet, Christian Rohner, and Thiemo Voigt. 2017. Towards wide-area backscatter networks. *ACM HOTWIRELESS* (2017).

[11] Pengyu Zhang et al. HitchHike: Practical Backscatter using Commodity WiFi. In *ACM SENSYS 2016*.

Robust and Efficient EH-WSN Simulation using Solar Radiation Data

[Extended Abstract]

James Allen
Newcastle University
Newcastle upon Tyne, UK
j.d.allen2@newcastle.ac.uk

Matthew Forshaw
Newcastle University
Newcastle upon Tyne, UK
matthew.forshaw@newcastle.ac.uk

Nigel Thomas
Newcastle University
Newcastle upon Tyne, UK
nigel.thomas@newcastle.ac.uk

ABSTRACT

Solar radiation data has been widely used to support simulation of energy-harvesting wireless sensor network systems. Simulation is an important tool in understanding the effectiveness of new energy management policies and energy-harvesting aware networking for sensor networks using solar energy as their power source. Several open-access solar radiation databases are available online which have been used as simulation input in studies. There is little consistency in the temporal granularity of the data used, and often only a handful of geographic locations are assessed. Temporal granularity of solar data may significantly affect the accuracy of reported energy performance, and using few geographic locations limits the scope of conclusions drawn to a subset of weather patterns experienced around the world. We propose research objectives and methods which will contribute to the understanding of the effect of solar data granularity on simulation accuracy, and the collation of disparate solar radiation data sets from many different sources.

CCS CONCEPTS

• **Computer systems organization → Sensor networks**; • **Hardware → Power and energy**;

KEYWORDS

Wireless Sensor Networks; Energy Harvesting; IoT

1 INTRODUCTION

A key limiting factor to the use of Wireless Sensor Networks (WSNs) is battery life, which limits their operational lifetime. Battery replacement is costly, and not always practically feasible due to their deployment environment. A great deal of attention has been directed towards the use of energy harvesting (EH) as an alternative power source for WSNs (EH-WSNs) as a solution to these problems.

EH-WSNs pose new challenges to energy management and EH-aware networking. To support the exploration of this new design space, WSN simulators have emerged which incorporate energy harvesting models, such as GreenCastalia [2], SensEH [7], HarvWSNet [10], WSNSim [15], and extensions to NS-3 [22] [6]. Among these, GreenCastalia has many commendable features which make it the most promising target for further research and development. These include its open-source license, realistic wireless channel modelling, its extensibility, and the use of the popular and mature OMNET++ event driven simulator as its foundation.

There are two broad approaches to modelling solar energy harvesting in an EH-WSN system. Firstly, calculations of the solar radiation incident on a tilted surface at a particular latitude which take into account atmospheric conditions has been well studied [8]. Combining this astronomical model with appropriate atmospheric measurements such as turbidity, horizontal visibility and cloud coverage can yield satisfactory weather system models which produce estimations of ground solar radiation [13]. These must be tuned using atmospheric data sets from the location of interest, and should be validated using historic solar radiation data sets from the location of interest. Such models are complex, but have the advantage that obstruction effects from surrounding terrain can be accurately modelled.

Secondly, given that weather models must be tuned using real atmospheric data from the location to be modelled, a simpler approach is to directly use historic measurements of solar radiation from that location. Since solar radiation data sets from a large number of locations are now available online [19] [21] [5] [20], these could be used to rapidly simulate EH-WSN performance over a wide range of locations, under the assumption of no obstructing terrain. This would provide an efficient method of testing and optimising energy management policies and EH-aware networking protocols over a wide range of weather conditions.

Solar radiation data sets have been used in various EH-WSN simulation studies, but have generally been limited to using only a small number of locations, and granularities of the data used range from 30 seconds, to 1 minute, 10 minutes, hourly and daily [1] [9] [6] [12] [17] [4]. This variability in data used is a result of the disparate nature of online solar radiation databases, their differing observation frequencies and data formats, and the time consuming nature of collating and combining multiple data sets.

Depending on system parameters such as energy storage capacity, solar panel size, sensing frequency, and the characteristics of the MAC and routing layers used, the granularity of the solar radiation data may have a significant effect on the accuracy of the energy performance reported. Simulations of locations with highly variable cloud conditions in particular may yield significant errors in reported performance when a greater granularity of data is used (daily) compared to smaller granularities (minutely). They may also report significant variations in performance compared to more stable climates, such as those with comparatively less cloud cover during a typical day. It is therefore important to be able to simulate performance over a wide range of locations, and to have understanding of the impact of the temporal granularity of the data used.

2 RESEARCH OBJECTIVES

Our contributions to EH-WSN simulation research will be as follows.

Effect of granularity of solar radiation data on simulation results: We will perform a number of simulations using different granularities of solar radiation data (minutely, 30 minute, hourly and daily) across a number of scenarios to determine the effect of granularity on reported energy performance. This will provide researchers with a reference guide by which appropriate solar radiation data sets can be chosen for a given scenario. Our conclusions will be further supported by comparison to real EH-WSN deployments at Newcastle University.

Effect of location on simulation results: Extensions to GreenCastalia will be developed which allow user access to a wide range of open-access solar radiation data sources. Meteonorm provide a consolidated database of a wide range of sources [16], but is not open-access, and thus less suitable for open publication of repeatable research results. By developing integration modules to a wide range of open data sources, we aim to publish EH-WSN simulation results over a wider range of locations than has been published before. This should provide researchers with a more rigorous assessment of the variation in performance of EH-WSN designs deployed in different locations across the world. By incorporating methods from a related study which synthesises lower granularity data (minutely) from higher granularity data (hourly) [3] in a way which captures the fluctuations and variability observed in minutely data, we can further extend the range of data sets available which can be run at minutely resolution.

Scalability of large scale simulations: It is anticipated that simulation time will increase significantly when using a large number of solar radiation data sources across a number of scenarios. We will investigate the performance of GreenCastalia during our simulations over a wide range of locations and scenarios to identify any opportunities to improve efficiency using parallel execution of tasks. This may result in the development of parallel infrastructure support for GreenCastalia which does not currently exist.

3 METHODS

Modules for GreenCastalia have been developed which simulate events and corresponding energy consumption traces of TelosB sensor motes using a BoX-MAC-2 [18] MAC layer and a Collection Tree Protocol (CTP) [11] routing layer. These will be used to run simulations with varying parameters, such as solar panel size, energy storage capacity and sensor sampling rate, and with varying network topologies. Existing minutely radiation data sets will be downsampled to provide comparable data sets at different granularities.

TelosB motes will be programmed with BoX-MAC-2 and CTP stacks and deployed at Newcastle University to provide a sensor network test bed against which to compare simulation results.

Following a survey of available online solar radiation databases, common data source features will be identified, and a GreenCastalia interface defined against which to develop an extensible collection of sub-modules which automate downloading data sets from a variety of locations and provide any necessary data transformations. Performance will be analysed using CPU, memory and time elapsed metrics. Integration with Akaroa [14] may developed to provide parallel execution support to GreenCastalia as necessary.

4 ACKNOWLEDGMENTS

This work was supported by the Engineering and Physical Sciences Research Council, Centre for Doctoral Training in Cloud Computing for Big Data [grant number EP/L015358/1].

REFERENCES

[1] Giuseppe Ateniese, Giuseppe Bianchi, Angelo T Capossele, Chiara Petrioli, and Dora Spenza. 2016. Low-cost Standard Signatures for Energy-Harvesting Wireless Sensor Networks. *ACM Trans. Embed. Comput. Syst.* 16, January (2016).

[2] David Benedetti, Chiara Petrioli, and Dora Spenza. 2013. GreenCastalia: An Energy-Harvesting-Enabled Framework for the Castalia Simulator. *Proceedings of the 1st International Workshop on Energy Neutral Sensing Systems* (2013).

[3] J M Bright, C J Smith, P G Taylor, and R Crook. 2015. Stochastic generation of synthetic minutely irradiance time series derived from mean hourly weather observation data. *Solar Energy* 115 (2015), 229–242.

[4] Bernhard Buchli, Felix Sutton, Jan Beutel, and Lothar Thiele. 2014. Towards enabling uninterrupted long-term operation of solar energy harvesting embedded systems. In *Lecture Notes in Computer Science*. Vol. 8354. 66–83.

[5] City of Edmonton. 2017. City of Edmonton Open Data Portal. (2017). https://data.edmonton.ca

[6] S Climent, A Sanchez, S Blanc, J V Capella, and R Ors. 2016. Wireless sensor network with energy harvesting: modeling and simulation based on a practical architecture using real radiation levels. *Concurrency Computat.: Pract. Exper* 28 (2016), 1812–1830.

[7] Riccardo Dall'Ora, Usman Raza, Davide Brunelli, and Gian Pietro Picco. 2014. SensEH: From simulation to deployment of energy harvesting wireless sensor networks. *Proceedings of the 39th Annual IEEE Conference on Local Computer Networks Workshops* (2014), 566–573.

[8] J V Dave, P Halpern, and H J Myers. 1975. Computation of Incident Solar Energy. *IBM Journal of Research and Development* 19, 6 (1975), 539–549.

[9] Moumita Deb and Roy Sarbani. 2016. Link Quality Modeling for Energy Harvesting Wireless Sensor Networks. *Proceedings of the First International Conference on Intelligent Computing and Communication* (2016), 669–678.

[10] Amine Didioui, Carolynn Bernier, Dominique Morche, and Olivier Sentieys. 2013. HarvWSNet: A co-simulation framework for energy harvesting wireless sensor networks. *2013 International Conference on Computing, Networking and Communications (ICNC)* (2013), 808–812.

[11] Omprakash Gnawali, Rodrigo Fonseca, Kyle Jamieson, Maria Kazandjieva, David Moss, and Philip Levis. 2013. CTP: An Efficient, Robust, and Reliable Collection Tree Protocol for Wireless Sensor Networks. *ACM Transactions on Sensor Networks* 10, 1 (2013), 16.

[12] L Hanssen and J Gakkestad. 2010. Solar Cell Size Requirement for Powering of Wireless Sensor Network Used in Northern Europe. *Proceedings of the International Workshops on PowerMEMS* (2010), 17–20.

[13] Jaein Jeong and David Culler. 2012. Predicting the Long-Term Behavior of a Micro-Solar Power System. *ACM Trans. Embed. Comput. Syst.* 11, 2 (2012), 35:1–35:38.

[14] Don Mcnickle, Krzysztof Pawlikowski, and Greg Ewing. 2010. Akaroa2: A Controller of Discrete-Event Simulation which exploits the Distributed Computing Resources of Networks. *Proc. European Conf. on Modelling and Simulation (ECMS 2010)* (2010).

[15] Geoff V. Merrett, Neil M. White, Nick R. Harris, and Bashir M. Al-Hashimi. 2009. Energy-aware simulation for wireless sensor networks. *2009 6th Annual IEEE Communications Society Conference on Sensor, Mesh and Ad Hoc Communications and Networks, SECON 2009* (2009).

[16] Meteonorm. 2017. Meteonorm. (2017). www.meteonorm.com

[17] C Moser, L Thiele, D Brunelli, and L Benini. 2007. Adaptive Power Management in Energy Harvesting Systems. *Design, Automation Test in Europe Conference Exhibition* (2007), 773–778.

[18] David Moss and Philip Levis. 2008. *BoX-MACs : Exploiting Physical and Link Layer Boundaries in Low-Power Networking*. Technical Report. Stanford Information Networks Group, Stanford.

[19] NREL. 2017. National Renewable Energy Laboratory. (2017). http://midcdmz.nrel.gov

[20] PVGIS. 2017. Photovoltaic Geographical Information System. (2017). http://re.jrc.ec.europa.eu/pvgis

[21] Newcastle University. 2017. Newcastle Urban Observatory. (2017). http://uoweb1.ncl.ac.uk

[22] He Wu, Sidharth Nabar, and Radha Poovendran. 2011. An Energy Framework for the Network Simulator 3 (ns-3). *Proceedings of the 4th International ICST Conference on Simulation Tools and Techniques* 3 (2011), 222–230.

Exploring Energy Efficient State Retention in Transiently-Powered Computing Systems

Extended Abstract

Theodoros D. Verykios, Domenico Balsamo, Geoff V. Merrett
School of Electronics and Computer Science
University of Southampton
Southampton, UK
{T.Verykios,D.Balsamo,gvm}@soton.ac.uk

ABSTRACT

Batteries have traditionally been used to power embedded electronic devices. However, requirements such as a long lifetime, low cost, and weight, pose significant challenges to battery-powered systems. Energy harvesting offers the potential for embedded systems to operate without batteries. Nonetheless, harvesting has been traditionally coupled with large energy buffers such as supercapacitors to tackle the instability of the source. Transiently-powered computing systems enable computation to be sustained despite the source variability, without the need for additional energy storage. To make this feasible, the system state (e.g. registers and RAM) needs to be saved to Non-Volatile Memory (NVM) before a power outage, and restored once power is available again. Existing transient systems save the entire state of the system upon power failure and do not consider the properties of different NVM technologies, leading into a sub-optimal state retention process. As a consequence, the time and energy spent towards useful computation are decreased significantly, affecting the forward progress that the system can achieve. The aim of this research is to introduce novel methods to reduce the time and energy overhead of the state retention process, exploring solutions both in the software and hardware domain.

CCS CONCEPTS

• **Computer systems organization** → *Embedded software*; Embedded hardware; • **Hardware** → Memory and dense storage;

KEYWORDS

Transient Computing; State Retention; Embedded Systems; Energy Harvesting

1 INTRODUCTION

The number of battery-powered electronic systems has increased dramatically with the vast deployment of mobile, autonomous and wearable devices. Even though batteries act as a virtually unlimited power source for a certain period of time before becoming exhausted, they come at a dimensional, financial, and environmental cost which cannot be neglected. The nature of some applications such as implantable bio-sensors [7] and underground WSNs [5] implies limited access and, consequently, maintenance becomes a serious challenge. In addition, even though CMOS devices have been downscaled significantly over the last years, the energy density of batteries has not been increased accordingly. The realisation of some devices such as implantable bio-sensors (which need to be miniaturised) was challenging due to the size of batteries. Therefore, the need for systems that can operate without batteries has emerged [1].

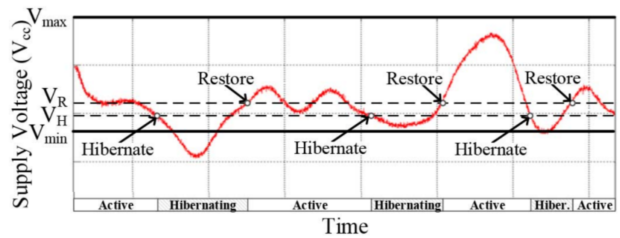

Figure 1: Operation of a typical transient system [2]

2 TRANSIENT SYSTEMS

Energy harvesting (EH) systems scavenge energy from environmental sources such as light, vibration, motion or temperature to power themselves, instead of relying on batteries [3]. Nevertheless, factors such as the weather condition, availability of light, or the intensity of vibration can significantly affect the energy availability. Relying solely on these sources can therefore result in the system being unable to sustain computation.

Transiently-powered computing systems are storage-less systems that enable computation to be sustained, despite the variable and unstable energy harvested from the environment [6]. Transient systems retain their state in Non-Volatile Memory (NVM) upon a power failure to cope with frequent power interruptions. This implies that the main memory and registers are saved before a power outage, and restored when the power is available again, as shown in Figure 1 which shows the operation of a typical transient system. Several software-based approaches such as Hibernus [2] have been proposed for transient computing; however, they all save the entire volatile state without considering what is actually required. Furthermore, using a universal policy, without regard for the NVM technology, results in considerable time and energy overhead for the state retention process.

Recently, Bhatti et al. [4] proposed a selective policy for efficient state retention which dynamically identifies the unallocated space and only saves to Flash memory the parts of the main memory being used by the main application. We implemented this "Allocated State" policy on two platforms with different NVM technologies (FRAM and Flash). Figure 3 shows the Allocated State policy applied to different applications, where the time and energy required to save

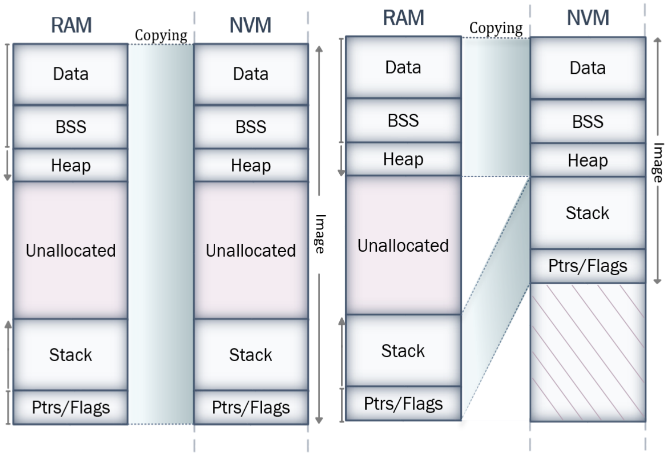

(a) Complete State [2] **(b) Allocated State [4]**

Figure 2: Existing state retention approaches illustrating a) Complete State and b) Allocated State

the system state are experimentally measured. This policy allows for substantial energy time and energy savings when used with NVM technologies that do not require erasing, such as FRAM. Figures 2a and 2b demonstrate that the cost for saving is proportionally reduced with the size of allocated memory, when compared to saving the entire memory (up to 85.1% reduction when the memory usage is 18%). However, when this policy is implemented using Flash memory, it is far less effective as shown in Figures 2c and 2d. This is because the overhead due to the erasing process (typical for Flash memory) that is needed before saving the system state, is not taken into account.

3 SELECTIVE POLICIES

Our current research focuses on devising novel selective policies for efficient state retention to improve the energy efficiency of transiently-powered computing systems. These policies exploit the characteristics of different NVM technologies, to ensure that the state retention is an energy and time efficient operation. In addition, the memory usage by the main application is taken into account when designing these policies. Therefore, the policies are tailored to the specific characteristics of each NVM type and the memory usage and consequently, the energy savings can be maximised.

A fundamental challenge is the integration of these policies in state-of-the-art transient systems in order to show the benefits against a universal policy. Using these policies, the overhead for saving/restoring the system state can be reduced which can lead into more energy efficient transient systems.

One of the most important parameters of transient systems is the threshold voltages at which the system starts its state retention process or restores the system state in order to continue its operation. Our future work will be concerned with optimising these thresholds by integrating these policies, in order to dynamically maximise the useful computation time and therefore, ensure that transient systems become more energy efficient.

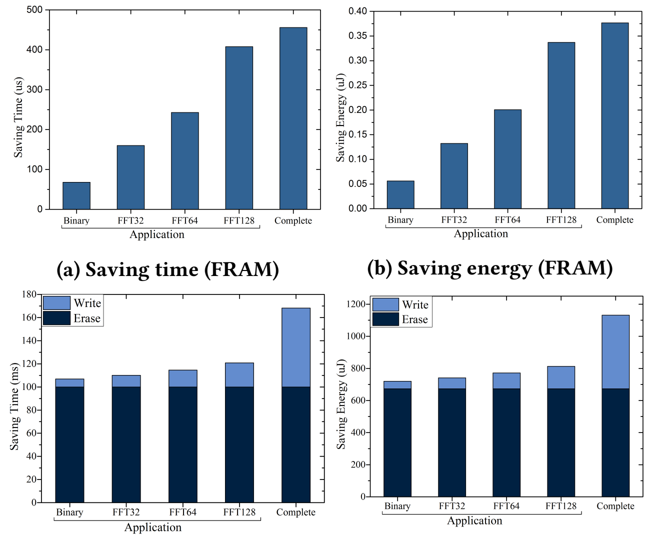

(a) Saving time (FRAM) **(b) Saving energy (FRAM)**

(c) Saving time (Flash) **(d) Saving energy (Flash)**

Figure 3: Time and energy overhead of Allocated State with FRAM and Flash memories

4 CONCLUSIONS

In this extended abstract, the need for energy efficient state retention policies was highlighted. Software based policies are investigated which are targeted for different NVM technologies, exploiting their individual properties and the way the memory is used by the main application. Future work will focus on predicting the energy needs before saving the system state in order to maximise the time and energy spent on useful computation.

ACKNOWLEDGMENTS

This work was supported in part by UK Engineering and Physical Sciences Research Council (EPSRC) Platform Grant EP/P010164/1. Experimental data accompanying this work can be found at DOI:10.5258/SOTON/D0276 (http://doi.org/10.5258/SOTON/D0276).

REFERENCES

[1] D. Balsamo, A. Das, A. S. Weddell, D. Brunelli, B. M. Al-Hashimi, G. V. Merrett, and L. Benini. 2016. Graceful performance modulation for power-neutral transient computing systems. *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.* 35, 5 (2016), 738–749.

[2] D. Balsamo, A. S. Weddell, A. Das, A. R. Arreola, D. Brunelli, B. M. Al-Hashimi, G. V. Merrett, and L. Benini. 2016. Hibernus++: a self-calibrating and adaptive system for transiently-powered embedded devices. *IEEE Trans. Comp.-Aided Des. Integr. Circuits Syst.* 35, 12 (2016), 1968–1980.

[3] S. Beeby and N. White. 2010. *Energy harvesting for autonomous systems.* Artech House.

[4] N. A. Bhatti and L. Mottola. 2016. Efficient State Retention for Transiently-powered Embedded Sensing.. In *EWSN 2016.* Graz, Austria, 137–148.

[5] S. Kahrobaee and M. C. Vuran. 2013. Vibration energy harvesting for wireless underground sensor networks. In *IEEE Int'l Conf. on Communications (ICC) 2013.* IEEE, 1543–1548.

[6] K. Ma, Y. Zheng, S. Li, K. Swaminathan, X. Li, Y. Liu, J. Sampson, Y. Xie, and V. Narayanan. 2015. Architecture exploration for ambient energy harvesting nonvolatile processors. In *IEEE 21st HPCA 2015.* IEEE, San Francisco, CA, USA, 526–537.

[7] J. Olivo, S. Carrara, and G. De Micheli. 2011. Energy harvesting and remote powering for implantable biosensors. *IEEE Sensors Journal* 11, 7 (2011), 1573–1586.

Hardware and Software System Support for Intermittent Energy-harvesting Devices

A case-study on a solar-powered nano-satellite

Alexei Colin
Carnegie Mellon University, Pittsburgh, PA, USA
acolin@andrew.cmu.edu

ABSTRACT

Energy-harvesting devices are embedded computers that operate without batteries by gathering energy from their environment. Without wires and batteries, sensing devices can be deployed in locations otherwise inaccessible due to extreme cold, such as in space. However, a microcontroller and sensors cannot be powered directly from a weak and unstable energy harvester, requiring custom power system design Furthermore, applications developed for traditional battery-powered embedded devices may not execute correctly on energy-harvesting platforms due to power failures that may leave program state inconsistent in memory. We propose to improve accessibility of energy-harvesting platforms to embedded developers through system support at hardware and software levels. In this work we apply a state-of-the-art intermittent computing system and develop a custom hardware power system to build a solar-powered board-scale nano-satellite.

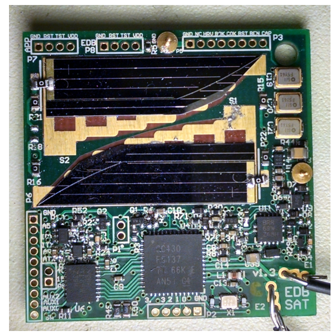

Figure 1: EDBsat: a solar-powered nano-satellite (front).

1 INTRODUCTION

Sensing devices for monitoring human health, structural integrity, or exploring the planet from orbit are constrained by their energy source. A battery or a wired connection to the grid limits the lifespan and versatility of a device. For example, in a satellite, wires are not an option and batteries undergo irreversible damage at low temperatures without heaters. Emerging *energy-harvesting circuits* offer an alternative power supply that does not suffer from the same limitations. An energy-harvester can power a computing device with mechanical, solar, radio, or thermal energy available in its environment. Unfortunately, these novel power sources cannot seamlessly replace traditional ones, because their unpredictable power output jeopardizes correct operation of the software.

Energy availability alone is not sufficient for successful execution of the program, because a power loss may interrupt the execution unexpectedly, corrupting the program state in non-volatile memory. As a consequence, *embedded software designed for traditional platforms is not portable to intermittently-powered platforms.* To develop reliable and efficient software for intermittently-powered platforms programmers need a hardware platform with a system software stack that is aware of intermittent execution. In this work we give a brief overview of prior work in language and a profiling tool, and describe their application to EDBsat, a board-scale solar-powered satellite, with a focus on the satellite's hardware power system.

2 BACKGROUND

Energy-harvesting devices differ from traditional embedded devices in the need for hardware circuits for harvesting and storing energy and in the intermittent execution of code on the processor.

2.1 Energy-harvesting Power Systems

Before being able to consume the incoming energy that arrives unpredictably and at a low voltage, a device must accumulate it into a capacitor and possibly boost or buck the voltage. The power system circuit on the device is responsible for gathering the ambient energy, charging the capacitor, and connecting the microcontroller and sensors to the capacitor once it is charged. The challenge in power system design is to use the least amount of volume to supply a voltage to the microcontroller, sensors, and radio that is sufficiently high for the duration of each operation that must take place atomically. For example, the radio on the EDBsat requires 2V for the duration of a packet transmission (250 ms).

2.2 Intermittent Execution

To run on energy-harvesting devices, software must be robust with respect to an *intermittent execution model.* In this model, code executes in short bursts interrupted by power loss, in contrast to the familiar continuous execution model where computation proceeds uninterrupted until the program finishes. State stored in volatile structures (i.e., registers, SRAM) is lost with each power failure. Unless the system implements a mechanism to resume after a power loss, it cannot run *reliably* and *efficiently* on an energy-harvesting platform. Several such mechanisms were proposed based on saving checkpoints of volatile state into non-volatile memory and restoring them on every reboot after power loss [3]. A recent alternative to checkpointing, Chain [2] is based on idempotent tasks that statically multi-version data to avoid overwriting their inputs. In this work, we develop the satellite payload sensing application in Chain and deploy it onto the application processor.

3 A SOLAR-POWERED NANO-SATELLITE

We demonstrate a co-designed energy-harvesting hardware platform with a system software stack, by building *EDBsat*, a solar-powered nano-satellite, shown in Figure 1. The satellite is constrained by the KickSat form factor to be one circuit board of size 35x35x4 mm. The primary objective is to collect samples from a temperature sensor, magnetometer, and accelerometer, average them over several time scales, and transmit the time series over radio. A secondary objective is to record a profile of the energy harvested by solar panels as a histogram of the capacitor voltage at different points in the program. We provision the satellite with two MCUs – one for the application, and one for the radio stack and the energy profiler – and a custom power system that distributes the energy between them.

4 SYSTEM SOFTWARE STACK

To create systems support for intermittently-powered platforms we must recognize their fundamental differences from traditional platforms. From the programmer's perspective, moving to an intermittently-powered platform requires rethinking allocation of program state in volatile and non-volatile memory, and ensuring progress of the computation. Chain [2], provides strong guarantees about memory consistency and forward progress during intermittent execution. The satellite application was written using the Chain primitives from the ground up to inherit the consistency guarantee.

Consistency of program state in memory is a necessary but not a sufficient condition for successful execution on an energy-harvesting device. The program may fail to terminate if one of its tasks requires more energy than the capacitor stores – a situation that occurred during the development of the satellite. While this failure mode is observable as the lack of radio packets, the non-terminating task, and the specific path through that task is not observable without an instrumentation tool. To debug non-termination as well as to profile energy available in the environment, we built into the satellite a custom implementation of the Energy-Interference Debugger (EDB) [1]. In contrast to its original design as stand-alone debugging tool powered from the USB port of a computer, the EDB module on the satellite is powered from an energy harvester and delivers recorded information via the radio channel. *In situ* tracing and profiling in an energy-harvesting context helped identify an insufficiently sized capacitor during development.

5 ENERGY-HARVESTING POWER SYSTEM

Intermittent power failures affect not only the processor but also the peripherals, e.g. sensors and radios. For example, a radio transmission of a packet does not reach the destination if it is interrupted by a power failure. To satisfy this atomicity constraint, we developed a power system that can store enough energy for a packet transmission in a small footprint. Figure 2 shows the design of the power system. The system distributes power from one harvester (an array of four solar panels) to two isolated load domains: the application MCU and sensors and the radio MCU which also runs the energy profiler. The diodes prevent a capacitor from one domain from discharging into the other domain, and the switch prevents the EDB domain from loading the harvester — thereby starving the application from energy — during profiling. To fit within the size constraints we use low-profile supercapacitors and compensate for their limited current capacity by boosting the discharge voltage and maximize stored energy by boosting the charge voltage to the maximum rating of the capacitor.

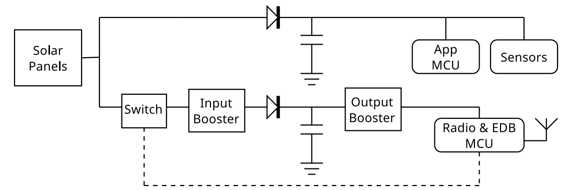

Figure 2: Power system hardware design

Time	Temp	Mag	Accel
Short	30	(7, 0, -5)	(1, -1, -3)
Long	30	(7, 0, -6)	(1, -1, -3)

Task	#HighE	#LowE
Boot	1	0
Sample	1	29
Avg-1	1	29
Avg-2	0	31

Table 1: Radio packets with sensor data and energy profile histograms from the satellite.

6 PRELIMINARY RESULTS

We have designed and manufactured the EDBsat circuit board and deployed onto it the application software written in Chain and a customized version of the EDB firmware [1]. To verify the sensing and energy profiling functionality we received and decoded the radio packets from the device over one hour while it was powered by light from an incandescent bulb. Table 1 shows a received packet with averages of sensor values down-sampled to 4 bits each over two time scales, and the energy profile for 4 instrumented points in the program. The profile shows that the program booted once and had sufficient incoming energy to execute all its tasks without rebooting, but starting each task with a partially depleted capacitor most of the time.

7 CONCLUSION

Energy-harvesting devices have the potential for sensing in adverse environments where batteries do not function well, such as in space. However, the weak and unstable power supply from an energy harvester requires a custom hardware power system and a software stack to tolerate intermittent execution. We have developed a versatile power system for a solar-powered satellite and applied prior work on intermittent systems to deploy a sensing application. In the future we will investigate mechanisms for programmatic control of the amount of energy to accumulate in the capacitor.

REFERENCES

[1] Alexei Colin, Graham Harvey, Brandon Lucia, and Alanson P Sample. 2016. An energy-interference-free hardware-software debugger for intermittent energy-harvesting systems. *ACM SIGPLAN Notices* 51, 4 (2016), 577–589.

[2] Alexei Colin and Brandon Lucia. 2016. Chain: tasks and channels for reliable intermittent programs. In *OOPSLA*. ACM, 514–530.

[3] Brandon Lucia and Benjamin Ransford. 2015. A simpler, safer programming and execution model for intermittent systems. In *ACM SIGPLAN Notices*, Vol. 50. ACM, 575–585.

Configurable Buck-Boost Converter for RF Energy Harvesting and Transfer

G. C. Martins and W. A. Serdijn
Section Bioelectronics, Delf University of Technology
Delft, The Netherlands

Biomedical devices are often powered from batteries, which have a limited life time and/or must be charged from time to time. However, with the decreasing power consumption of the biomedical devices, energy harvesting techniques or wireless power transfer can be used as an alternative to batteries [3].

There are several modalities of energy harvesting, which include thermoelectric, solar and radio-frequency energy harvesting [6]. The harvester element captures energy from a source and generates an output voltage that depends on the available power. This voltage must be boosted and regulated in order to provide a stable supply for the electronics. The input resistance of this voltage booster should be set to optimize the harvester efficiency [5]. This optimum value is tracked using Maximum Power Point Tracking (MPPT) techniques. The voltage boosting converter must be designed to allow for changes in the input resistance while keeping a good conversion efficiency.

In this work, we specifically tackle the problem of wireless energy harvesting and transfer [4]. In Fig. 1, we present the block diagram of an energy converter that adapts itself to its varying input power. Two DC-DC converter blocks make the interface between the rectifier and the load. The first DC-DC converter extracts energy from the rectifier output and charges a storage capacitor. When enough energy is stored, it can be processed by the second DC-DC converter and relayed onto the load, which represents the target circuitry to be powered. If the available power is sufficiently high to continuously power the load, i.e., if no duty cycling is necessary, the DC-DC converter that charges the capacitor can be bypassed and the power is directly transferred from the rectifier to the load through the second DC-DC converter. Prior to the beginning of the harvesting operation, a start-up circuit starts the DC-DC converters and MPPT block.

In order to boost the voltage, a switched inductor DC-DC converter or a switched capacitor DC-DC converter can be used. While charging a storage capacitor, the output voltage will be increasing from the ground voltage up to the maximum possible. In the case of a switched capacitor converter, the maximum efficiency is obtained for specific values of the output voltage only, which depends on the selected converter configuration. In contrast, the switched inductor topologies can fundamentally operate with high efficiency for any output voltage.

In order to keep the input resistance of the DC-DC converter fixed while charging a capacitor, the input current must be independent of the output voltage. This can be performed using a buck-boost converter operating in Discontinuous Conduction Mode (DCM) and open loop, such as the non-inverting buck-boost converter in Fig. 2(a). It has an input current I_{in} (see Fig. 2(b)) that is independent of the output voltage V_{out}. This behaviour cannot be achieved with the straight-forward boost converter. The average input resistance of the buck-boost converter in DCM, neglecting

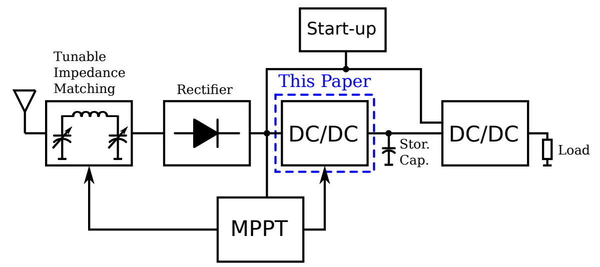

Figure 1: Adaptive power conversion chain simplified block diagram.

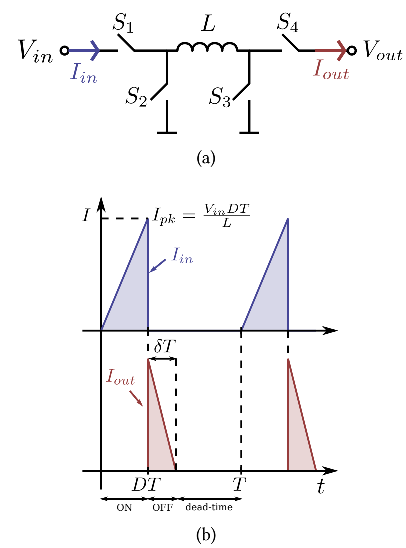

Figure 2: (a) Buck-boost converter circuit principle; (b) Input and output current in DCM.

the input voltage ripple, is given by:

$$R_{in,avg} = \frac{V_{in}}{I_{in,avg}} = \frac{2L}{D^2T}, \tag{1}$$

in which D is the duty cycle of the switching control signal (DT is the ON-time of the converter), L is the inductor value and T is the switching period.

Even though the same converter can be used in other energy harvesting systems, we optimized its design for a wireless energy harvesting system that employs a single-stage differential-drive CMOS rectifier [2]. The optimum output voltage of the employed rectifier varies from 0.38 V at 1- μW output power to 1.3 V at 1- mW

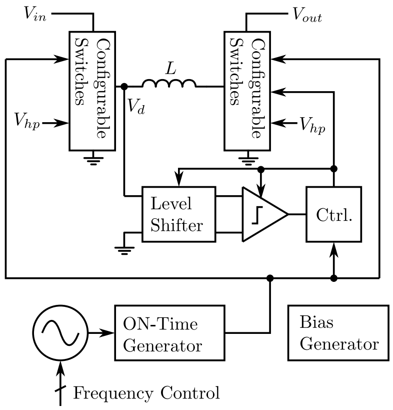

Figure 3: Block diagram of the proposed converter.

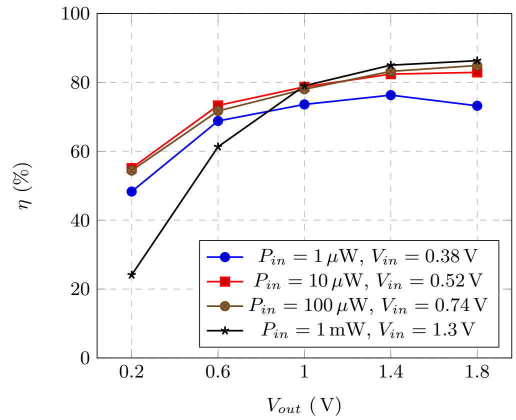

Figure 4: Efficiency results of the buck-boost converter for various output voltages.

output power. The technology used in this design is a standard 0.18-μm technology, with a maximum V_{DD} of 1.8 V. Those are the basic specifications in which the buck-boost converter must operate. The proposed converter's block diagram is presented in Fig. 3.

The power switches S_1 to S_4 (in Fig. 2(a)) are implemented by switches connected in parallel (not shown here). Wider switches are used for operation when $P_{in} > 500\,\mu$W and smaller switches otherwise. When the input power is low, only the smaller switches are used in order to reduce the power needed to drive them. Alternatively, we could reduce the switching losses exclusively by decreasing the switching frequency. This, however, would increase the peak value of I_L, which, in turn, would lead to increased conduction losses. Unless a larger inductor is used, which is undesired.

An adaptively biased amplifier is presented in [1], in which the bias current increases with the differential input V_{in} in order to increase slew rate. We use a similar approach, but apply it to a comparator to increase the bias current when V_{in} approaches zero, in order to decrease the time the comparator takes to bring the output from ground to V_{dd}.

Since the main application of this buck-boost converter is to charge a storage capacitor in an energy harvesting system while maintaining good efficiency for different available power levels, we present the efficiency results for a varying output voltage from 0.2 to 1.8 V in Fig. 4. The peak efficiency of the buck-boost converter is 76.3% at an input power of 1 μW and 86.3% at 1 mW. While charging a capacitor from 0.2 to 1.8 V, the total efficiency is 73.8% at $P_{in} = 1\,\mu$W. The lower efficiency at lower output voltages is observed because the OFF-time is larger (voltage drop across the inductor is lower) and the zero current detection circuitry is turned on for a longer time. Additionally, the output power switches are not optimized for low output voltages. This extra conduction loss is more prominent at higher P_{in}, because I_L is larger.

REFERENCES

[1] M. G. Degrauwe, E. A. Vittoz, and H. J. de Man. Adaptive biasing CMOS amplifiers. *IEEE J. Solid-State Circuits*, 17(3):522–528, 1982.
[2] K. Kotani, A. Sasaki, and T. Ito. High-efficiency differential-drive cmos rectifier. *IEEE J. Solid-State Circuits*, 44(11):3011–3018, 2009.
[3] Y. Manoli. Energy harvesting - from devices to systems. In *Proc. IEEE Eur. Solid State Circuits Conf. (ESSCIRC)*, pages 27–36, Sept. 2010.
[4] G. C. Martins and W. A. Serdijn. Adaptive buck-boost converter for RF energy harvesting and transfer in biomedical applications. In *2016 Biomedical Circuits and Systems Conference (BioCAS)*. IEEE, Oct. 2016.
[5] T. Paing, J. Shin, R. Zane, and Z. Popovic. Resistor emulation approach to low-power RF energy harvesting. *IEEE Trans. Power Electronics*, 23(3):1494–1501, 2008.
[6] R. Vullers, R. Schaijk, H. Visser, J. Penders, and C. van Hoof. Energy harvesting for autonomous wireless sensor networks. *IEEE Solid-State Circuits Magazine*, 2(2):29–38, 2010.

Think BIG be small: A Vision of Space IoT

Sujay Narayana
Delft University of Technology, The Netherlands
sujay.narayana@tudelft.nl

ABSTRACT

Progress in low power miniaturized electronics and wireless technologies have enabled many innovative applications. Of particular interest is the Internet of Things (IoT) that has dominated the world of ICT in the last half a decade enabling many smart systems and applications. At the same time, we also see much enthusiasm with respect to space missions and applications including terrestrial applications, apart from exploring the universe. We believe that innovations in the domain of IoT will significantly influence the space related activities - both research and development. In this article, we chart out the innovations in space and the vision with respect to embedded and wireless systems for space applications. In particular, we bring in the notion of Sensor Wireless Actuator Networks in Space (SWANS) and Space Pixels to explain IoT in space. We explain with examples what we envision for the next decade and also the challenges therein.

CCS CONCEPTS

• **Hardware** → Sensors and actuators; Wireless integrated network sensors; Communication hardware, interfaces and storage;

KEYWORDS

Space IoT, Internet of Things, Sensor Wireless Actuator Network in Space, swarms, nanosatellite, Space Pixels

1 INTRODUCTION

With the advent of low-power miniaturized electronics supporting high-end computations and advances in wireless technologies, it is certain that Internet of Things (IoT) has become one of the key technological enablers for smart-* systems in this decade. As an enabler for the Cyber Physical Systems, IoT is making a great impact on our lifestyle and the way we interact with others, environment, and even machines. Current IoT comprises of embedded devices, and a network of such devices having sensors and actuators, leveraging the existing wired or wireless infrastructures for communication and control of physical environment and electronic systems. The next step, which is just around the cornerstone, is to involve space in IoT applications. This step will help to solve few challenges in terrestrial IoT deployments. A single satellite in space can communicate with many sensor nodes and gateways over a vast area on Earth simultaneously than a single gateway on the ground. With a network of such satellites orbiting around Earth, it is possible to get a global coverage for IoT devices even in areas such as Arctic and Antarctic regions, mountains, oceans and remote places that have little or no infrastructure. Companies such as Magnitude Space [2] are in the process of employing such methods not only to cover remote areas but also provide global coverage for IoT devices even in densely populated cities. Magnitude Space, with their proposed system "Low-power Global Area Network", envisions to use many satellites speak directly to their proprietary sensor nodes connected to IoT devices on ground, and upload data to the cloud [5].

A satellite is already a collection of different sensor and actuator systems such as gyroscopes, magnetometers, solar panels, antennas, star and sun sensors, reaction wheels, magneto-torquers, and propulsion systems. They are connected by wires (also called *harness*) within a satellite, which adds to the weight and size (note that the mass and size of a satellite are important parameters that deeply impacts its cost including its launch) of the satellite. Replacing the wires with wireless without sacrificing its performance will make a satellite a sensor wireless actuator network in space (SWANS). Together with its external communication system, we extend SWANS to include a group of satellites in order to enable several applications including *swarms of satellites*. We envision that, in the near future, SWANS will impact the space industry significantly including the traditional space applications. The main driver for SWANS is the miniaturized, inexpensive, low-power network of satellites. We provide some examples here: (a) Swarm of autonomous small satellites (or tiny wireless smart sensor networks) deployed around moon or planets can collect more data with respect to deep space exploration than any observatory station on Earth. (b) Precise measurements for smoke and aerosol concentration in the atmosphere can be done easily and effectively using a swarm of tiny satellites orbiting in very Low Earth Orbits (LEO) than having a bulky equipment on Earth. (c) With swarm of robots, space robotics can be improved significantly to control robots from Earth with much lower latencies for targeted autonomous applications than the state of the art. These are a few major visions of the space industry for the coming years [3].

Currently, most of the existing satellites in space are bulky, expensive but ultra-reliable and designed for a particular mission. Instead, a swarm of networked, small, low-cost satellites can form a wireless sensor actuator networks in a distributed way to achieve the same desired features [4]. These swarm of small satellites can communicate with each other and also with ground stations in order to perform better than a big satellite. Such swarms in space can provide advantages such as redundancy, fault tolerance, and low-cost production and deployment, autonomous, incremental deployment and massively distributed. Hence, small satellite technology can be a key element to the future vision of Space IoT.

Moving one step further, we have the Starshot program by Stephen Hawking[1], a future project where a satellite in the form of a wafer-sized chip is shot to distant starts at ultra-high-speeds away from the Earth, using highly concentrated laser beams fired from the Earth. This shows that there is a need for much advancement in miniaturization and sensor technology to make the mission successful. For missions such as Starshot program, there has to be a huge advancement in the existing satellite technology to produce a tiny sensor and actuator chips that are ultra-low-power which can communicate with one another and with ground stations. Thinking big, to be small, the complete satellite should be embedded in a single pixel shaped miniaturized chip which we call "*Space Pixels*". With a size in the range of 1 mm x 1 mm and weigh less than 1 gm, Space pixels can be analogous to Smart Dust[6] containing low-power, processing, sensor, actuator, and communication units onboard. Their weight and size make it easy to release thousands of them in different orbits by a high altitude balloon from ISS or a small satellite where experiments such as RADAR reflectivity, atmospheric coloring and gas concentration, and meteor impacts can be done almost instantly. Not only in Earth orbits, a swarm of Space Pixels can also be used in deep space explorations. They can be used for near-field atmospheric measurements in other planets and space objects such as moons, asteroids, and space stations by throwing them into the orbit from landers or rover robots. In the case of remote sensing, multiple Space Pixels with micro cameras can get a bigger image in interplanetary explorations. The source of energy for these pixels can be solar or RF based. In the case of solar, there can be a tiny solar panel and in the case of RF, a large satellite or lander/orbiter on other planets can power these with continuous RF bursts. Space pixels can indeed be an important vision for IoT. Space pixels are always constrained by weight, size, power and communication capabilities and much exploration is yet to be done to meet these requirements. Thus space related scientific endeavors have umpteen possibilities and challenges. Future SWANS needs cutting-edge technology in the field of miniaturization, energy harvesting, and distributed sensor networks.

2 CHALLENGES

We list the important challenges to be faced during the development of SWANS/Space Pixels.

Energy Harvesting: As the system gets miniaturized, the amount of energy harvested also reduces while the power consumption may not reduce proportionally. Considering SWANS, the sources of energy include stars, electromagnetic waves and atmospheric gases (if available). However, the harvested energy from the sources varies drastically over location and time, for instance, Space Pixels may not have actuators to stabilize its rotation in the orbit, thus harvesting is intermittent.

Energy storage: Devices such as Space Pixels must tolerate the extreme temperature in space. Hence, usage of batteries may not be a solution because of its temperature and mass requirements. Thus, intermittent energy harvesting and battery less operation makes the design more challenging.

Distributed systems: Future SWANS should have a tight coupling between distributed sensors and actuator systems, and distributed computing and communication. The idea is to build and assemble independent modules that work towards a common goal. Within a space object, such as Space Pixels, or amongst multiple such objects, many modules should be able to accomplish a huge mission working in tandem independently but synchronously. Modules working well individually do not guarantee that the overall system is reliable, robust and adaptive to the harsh environments in space.

Miniaturization: Since access to space is expensive and requires enormous amount of resources (fuel, infrastructure, etc.), making space objects as small as possible is a must. Thus miniaturization of every subsystem is being targeted by academia and industry. However, miniaturization also poses challenges - constraints on harvested energy through small solar panels, power generation, control, and regulation. Though a small amount of energy is sufficient to power up systems such as Space Pixels, we do not have any miniaturized energy harvesters till date to fit entire power system in a tiny chip. With miniaturization, the hardware limitations such as radiation mitigation and thermal control, etc., becomes highly challenging.

Handling a large amount of data: Any space mission these days requires high resolution data (for instance, data from remote sensing cameras). The data handling in large space systems involves maintenance and powering of the memory modules however the tough problem is to communicate the data back to Earth. In miniaturized space systems, it is also tedious to store data for longer duration because of less space, lower power, and less protection from solar flares, etc. Until now, not much work has been done on exploring this challenge; with distributed systems and requirement of high resolution data, data management and communication become an important bottleneck.

Localization: Localization is a huge challenge in space, not only because of the vast volume to cover but also due to the fact that many space objects are moving at enormous speeds. Further, on other planets, mapping with miniaturized robots/systems is difficult as they do not contain any position sensors like GPS. Stitching the data from each of them to get a broader picture is highly challenging. Further, such localization algorithms require high amount of computations, which is scarce in miniaturized space objects.

REFERENCES

[1] 2017. Inside the Breakthrough Starshot Mission to Alpha Centauri. (2017). https://www.scientificamerican.com/article/inside-the-breakthrough-starshot-mission-to-alpha-centauri/ [Online;].

[2] 2017. MagnitudeSpace - Insights of Things. (2017). https://magnitudespace.com/ [Online;].

[3] 2017. STMD: Game Changing Development. (2017). https://www.nasa.gov/directorates/spacetech/game_changing_development/index.html [Online;].

[4] Nils Olsen, Claudia Stolle, Rune Floberghagen, Gauthier Hulot, and Alexey Kuvshinov. 2016. Special issue "Swarm science results after 2 years in space". *Earth, Planets and Space* 68, 1 (2016), 172. DOI: https://doi.org/10.1186/s40623-016-0546-6

[5] U. Raza, P. Kulkarni, and M. Sooriyabandara. 2017. Low Power Wide Area Networks: An Overview. *IEEE Communications Surveys Tutorials* PP, 99 (2017).

[6] B. Warneke, M. Last, B. Liebowitz, and K. S. J. Pister. 2001. Smart Dust: communicating with a cubic-millimeter computer. *Computer* 34 (2001).

Wireless Energy and Data Transfer in IoT Sensor Nodes

A. Urso and W. A. Serdijn
Section Bioelectronics, Delf University of Technology
Delft, The Netherlands

ABSTRACT

A big step towards cost-reduction of the electronics involved in a Wireless Sensor Network (WSN) is avoiding the use of batteries. In fact, for those applications where human access is very limited (e.g. implantable biomedical microsystems, electronic devices placed in remote area), the cost of the batteries, and their replacement, might exceed the cost of the electronics involved. Therefore there is the need for remotely-powered battery-less WSNs. In Figure 1, a general architecture of a typical autonomous WSN is depicted. When powered, the microcontroller reads and processes data from sensors, which can subsequently be sent to a transmitting antenna. Data can also be received and further processed by the microcontroller unit. In order to wirelessly transmit and receive data and receive power, an antenna or multiple antennas are needed. Since there are three RF input/output ports, the WSN can have up to three antennas. As depicted in Fig. 1(b), based on how the available antennas are connected, 5 different scenarios can be considered:

a) A WSN having one antenna. The antenna is shared between the transmitter, the receiver and the Energy Harvester (EH). In this scenario, these three blocks all have to work at the same frequency.
b) A WSN having two antennas. One antenna is used by the EH, while the other is shared between the transmitter and receiver.
c) A WSN having two antennas. One antenna is used by the transmitter, while the other is shared between the receiver and the EH.
d) A WSN having two antennas. One antenna is used by the receiver, while the other is shared between the transmitter and the EH.
e) A WSN having three antennas. One antenna for the receiver, one for the transmitter and one for the EH.

Scenario e) has the most degrees of freedom. This means that the most suitable frequency for each block can be used. For instance, the energy can be received at lower frequencies to achieve higher efficiency, while the data can be transferred at higher appropriate frequencies to achieve higher data rates.

Figure 2 shows the typical power profile of the output energy of an EH and the power consumption of the most power-hungry blocks of a WSN, the transmitter and the receiver. One can immediately notice that the power consumption required during the transmitting and receiving operations is significantly higher than the energy harvester output. Fortunately the power-hungry blocks operate only for a very short time. This leads to a scenario in which there is the need to continuously store the energy and power up the electronics in a duty cycled fashion.

With technology improvement, the total power consumption of WSNs tends to reduce whereas the power conversion efficiency of energy harvesters tends to increase, as depicted in Figure 3. It can be noticed once more that, the power consumption of the electronics goes well above what can be provided by the energy harvester. The reduction of the power consumption of the electronics (e.g. reduction of their operating voltages) is usually obtained at the expenses of immunity to circuit noise which is an important specification of all the electronics circuit. This makes continuous operation very difficult to achieve. By reducing the power consumption of the electronics the upper curve of Figure 3 can be pulled down. On the other hand, by increasing the power conversion efficiency of the energy harvester, i.e. pull up the curve representing the energy harvester's output. If one day, the two curves will meet each others,

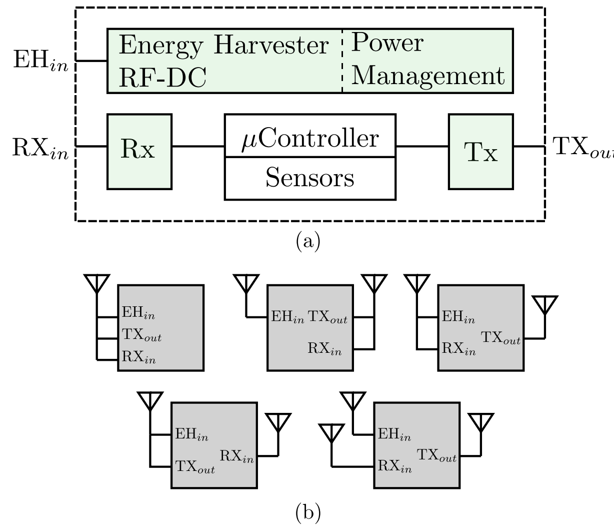

Figure 1: (a) Detailed block diagram of an autonomous wireless sensor node, (b) 5 possible configurations based on the number of available antennas [1].

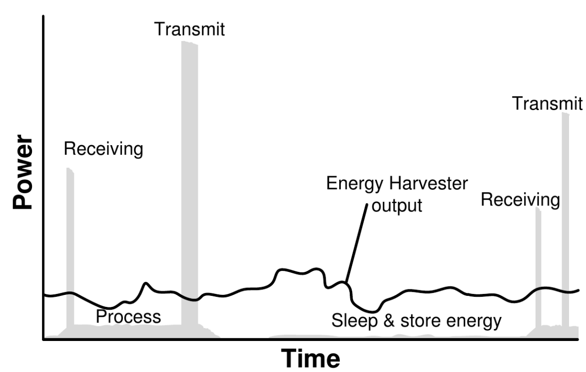

Figure 2: Typical power profile of a remotely-powered WSN [2].

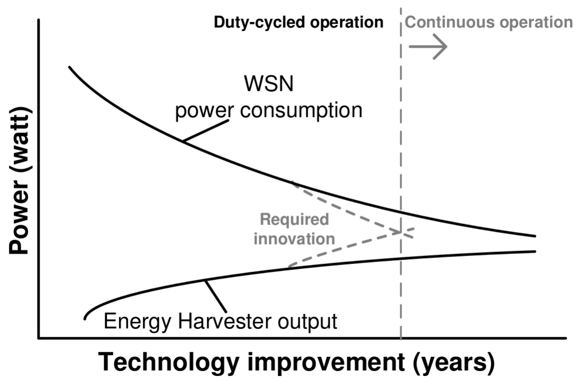

Figure 3: Typical power consumption scenario of a WSN [2].

the load can be operated continuously. However, due to the fact that the energy comes from an unreliable source, there will still be the need to store the energy.

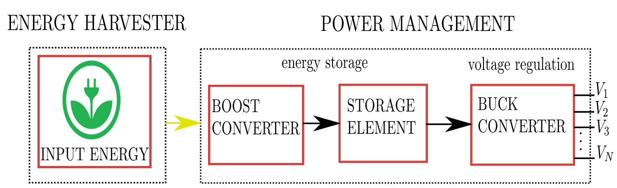

Figure 4: Detailed block diagram of the EH and power management block.

In Figure 4, the EH and the power management are illustrated in details. It is assumed that the EH performs the RF-DC conversion. However, its output voltage depends on the type of EH used and is usually in the range of few mV to few hundreds of mV. A boost DC-DC converter is used to bring the voltage up so that enough energy can be stored in the storage element. Nowadays supercapacitors are preferred over batteries thanks to their environmental-friendly propriety. Then a buck converter is used to power up the N loads, each of them with different voltages and power requirements.

REFERENCES

[1] G. C. Martins, A. Urso, A. Mansano, Y. Liu, and W. A. Serdijn, "Energy-Efficient Low-Power Circuits for Wireless Energy and Data Transfer in IoT Sensor Nodes," 2017. [Online]. Available: http://arxiv.org/abs/1704.08910
[2] M. Stoopman, "Circuit design for highly sensitive RF-powered wireless sensor nodes," Ph.D. dissertation, Delft University of Technology, Sep. 2016.

Aachener Informatik-Berichte

This list contains all technical reports published during the past three years. A complete list of reports dating back to 1987 is available from:

http://aib.informatik.rwth-aachen.de/

To obtain copies please consult the above URL or send your request to:

Informatik-Bibliothek, RWTH Aachen, Ahornstr. 55, 52056 Aachen,
Email: biblio@informatik.rwth-aachen.de

2014-01 * Fachgruppe Informatik: Annual Report 2014
2014-02 Daniel Merschen: Integration und Analyse von Artefakten in der modellbasierten Entwicklung eingebetteter Software
2014-03 Uwe Naumann, Klaus Leppkes, and Johannes Lotz: dco/c++ User Guide
2014-04 Namit Chaturvedi: Languages of Infinite Traces and Deterministic Asynchronous Automata
2014-05 Thomas Ströder, Jürgen Giesl, Marc Brockschmidt, Florian Frohn, Carsten Fuhs, Jera Hensel, and Peter Schneider-Kamp: Automated Termination Analysis for Programs with Pointer Arithmetic
2014-06 Esther Horbert, Germán Martín García, Simone Frintrop, and Bastian Leibe: Sequence Level Salient Object Proposals for Generic Object Detection in Video
2014-07 Niloofar Safiran, Johannes Lotz, and Uwe Naumann: Algorithmic Differentiation of Numerical Methods: Second-Order Tangent and Adjoint Solvers for Systems of Parametrized Nonlinear Equations
2014-08 Christina Jansen, Florian Göbe, and Thomas Noll: Generating Inductive Predicates for Symbolic Execution of Pointer-Manipulating Programs
2014-09 Thomas Ströder and Terrance Swift (Editors): Proceedings of the International Joint Workshop on Implementation of Constraint and Logic Programming Systems and Logic-based Methods in Programming Environments 2014
2014-14 Florian Schmidt, Matteo Ceriotti, Niklas Hauser, and Klaus Wehrle: HotBox: Testing Temperature Effects in Sensor Networks
2014-15 Dominique Gückel: Synthesis of State Space Generators for Model Checking Microcontroller Code
2014-16 Hongfei Fu: Verifying Probabilistic Systems: New Algorithms and Complexity Results
2015-01 * Fachgruppe Informatik: Annual Report 2015

2015-02 Dominik Franke: Testing Life Cycle-related Properties of Mobile Applications

2015-05 Florian Frohn, Jürgen Giesl, Jera Hensel, Cornelius Aschermann, and Thomas Ströder: Inferring Lower Bounds for Runtime Complexity

2015-06 Thomas Ströder and Wolfgang Thomas (Editors): Proceedings of the Young Researchers' Conference "Frontiers of Formal Methods"

2015-07 Hilal Diab: Experimental Validation and Mathematical Analysis of Cooperative Vehicles in a Platoon

2015-08 Mathias Pelka, Jó Agila Bitsch, Horst Hellbrück, and Klaus Wehrle (Editors): Proceedings of the 1st KuVS Expert Talk on Localization

2015-09 Xin Chen: Reachability Analysis of Non-Linear Hybrid Systems Using Taylor Models

2015-11 Stefan Wüller, Marián Kühnel, and Ulrike Meyer: Information Hiding in the Public RSA Modulus

2015-12 Christoph Matheja, Christina Jansen, and Thomas Noll: Tree-like Grammars and Separation Logic

2015-13 Andreas Polzer: Ansatz zur variantenreichen und modellbasierten Entwicklung von eingebetteten Systemen unter Berücksichtigung regelungs- und softwaretechnischer Anforderungen

2015-14 Niloofar Safiran and Uwe Naumann: Symbolic vs. Algorithmic Differentiation of GSL Integration Routines

2016-01 * Fachgruppe Informatik: Annual Report 2016

2016-02 Ibtissem Ben Makhlouf: Comparative Evaluation and Improvement of Computational Approaches to Reachability Analysis of Linear Hybrid Systems

2016-03 Florian Frohn, Matthias Naaf, Jera Hensel, Marc Brockschmidt, and Jrgen Giesl: Lower Runtime Bounds for Integer Programs

2016-04 Jera Hensel, Jürgen Giesl, Florian Frohn, and Thomas Ströder: Proving Termination of Programs with Bitvector Arithmetic by Symbolic Execution

2016-05 Mathias Pelka, Grigori Goronzy, Jó Agila Bitsch, Horst Hellbrück, and Klaus Wehrle (Editors): Proceedings of the 2nd KuVS Expert Talk on Localization

2016-06 Martin Henze, René Hummen, Roman Matzutt, Klaus Wehrle: The SensorCloud Protocol: Securely Outsourcing Sensor Data to the Cloud

2016-07 Sebastian Biallas : Verification of Programmable Logic Controller Code using Model Checking and Static Analysis

2016-08 Klaus Leppkes, Johannes Lotz, and Uwe Naumann: Derivative Code by Overloading in C++ (dco/c++): Introduction and Summary of Features

2016-09 Thomas Ströder, Jürgen Giesl, Marc Brockschmidt, Florian Frohn, Carsten Fuhs, Jera Hensel, Peter Schneider-Kamp, and Cornelius Aschermann: Automatically Proving Termination and Memory Safety for Programs with Pointer Arithmetic

2016-10 Stefan Wüller, Ulrike Meyer, and Susanne Wetzel: Towards Privacy-Preserving Multi-Party Bartering

2017-01 * Fachgruppe Informatik: Annual Report 2017

2017-02 Florian Frohn and Jürgen Giesl: Analyzing Runtime Complexity via Innermost Runtime Complexity

2017-04 Florian Frohn and Jürgen Giesl: Complexity Analysis for Java with AProVE

2017-05 Matthias Naaf, Florian Frohn, Marc Brockschmidt, Carsten Fuhs, and Jürgen Giesl: Complexity Analysis for Term Rewriting by Integer Transition Systems

2017-06 Oliver Kautz, Shahar Maoz, Jan Oliver Ringert, and Bernhard Rumpe: CD2Alloy: A Translation of Class Diagrams to Alloy

2017-07 Klaus Leppkes, Johannes Lotz, Uwe Naumann, and Jacques du Toit: Meta Adjoint Programming in C++

* These reports are only available as a printed version.

Please contact `biblio@informatik.rwth-aachen.de` to obtain copies.