Energy-Driven Computing:

Rethinking the Design of Energy Harvesting Systems

Geoff V Merrett

20 April 2022

University of Southampton

UNIVERSITY OF SOUTHAMPTON

University of Southampton

- ~25,000 students
- Top 100 universities worldwide (QS‘22)
- Founding member of UK’s Russell Group

School of Electronics and Computer Science

- ~2,000 students
- ~250 PhD research students
- ~100 academics/faculty
- Top 3 in UK for EEE
- 14 research groups/centres

LOCALLY-POWERED SYSTEMS

- A trillion connected devices by 2035?
- How do we power them?
  - Mains power?
  - Batteries?
    - Sustainability
    - Maintenance
    - Lifetime vs volume

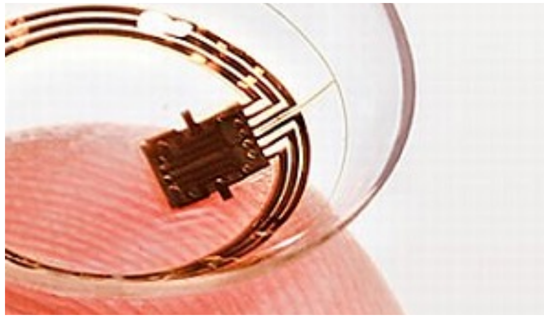

image: RS fabrications, http://www.rsfabrications.com/wp-content/uploads/2009/09/cimg2821.jpg
image: http://commons.wikimedia.org/wiki/File:Canary_Wharf_at_night,_from_Shadwell_cropped.jpg
image: Australian Transport Safety Bureau, http://www.atsb.gov.au/media/2489702/ro2010001.jpg

ENERGY HARVESTING

varies temporally

varies spatially

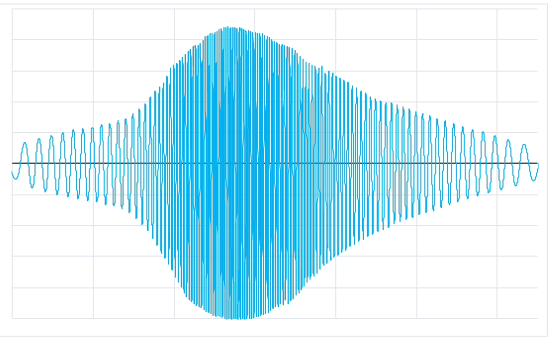

D. Balsamo et al. Hibernus++: a self-calibrating and adaptive system for transiently-powered embedded devices. IEEE TCAD, 1-13.

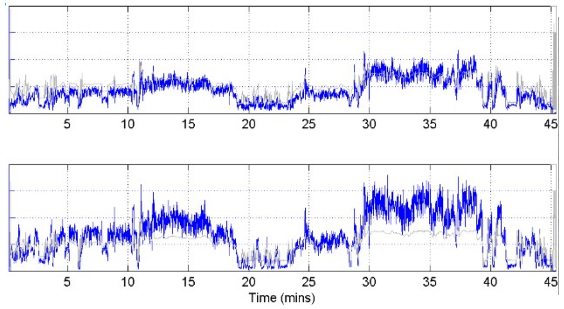

Beeby, S.P. et al. A comparison of power output from linear and non-linear kinetic energy harvesters using real vibration data. Smart Materials and Structures, 22, (7), 075022.

Power / Energy

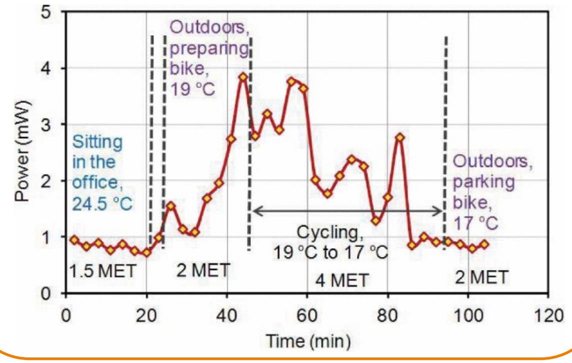

V. Leonov, "Thermoelectric Energy Harvesting of Human Body Heat for Wearable Sensors," IEEE Sensors Journal, vol.13, no.6, pp.2284-91, June13

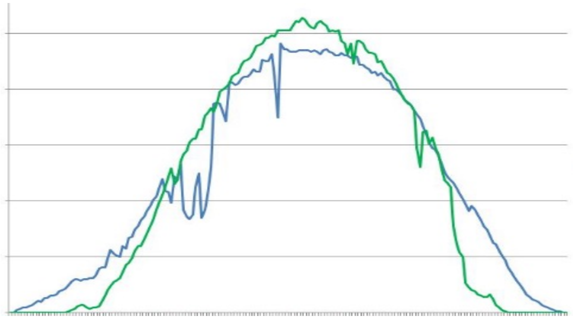

http://solar.rainham-kent.co.uk

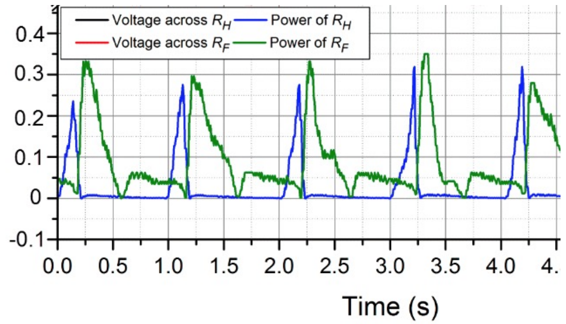

Zhao, J. et al. "A Shoe-Embedded Piezoelectric Energy Harvester for Wearable Sensors," Sensors 2014, 14, 12497-12510.

Highly variable supply + variable consumption!

networking/comms

harvester

power conversion

sensors

user

memory

An *energy harvester* is one part of a *system*

compute

application

user interface(s)

actuators

peripherals

energy storage

design tools

University of Southampton

ENERGY-NEUTRAL COMPUTING

Schematic

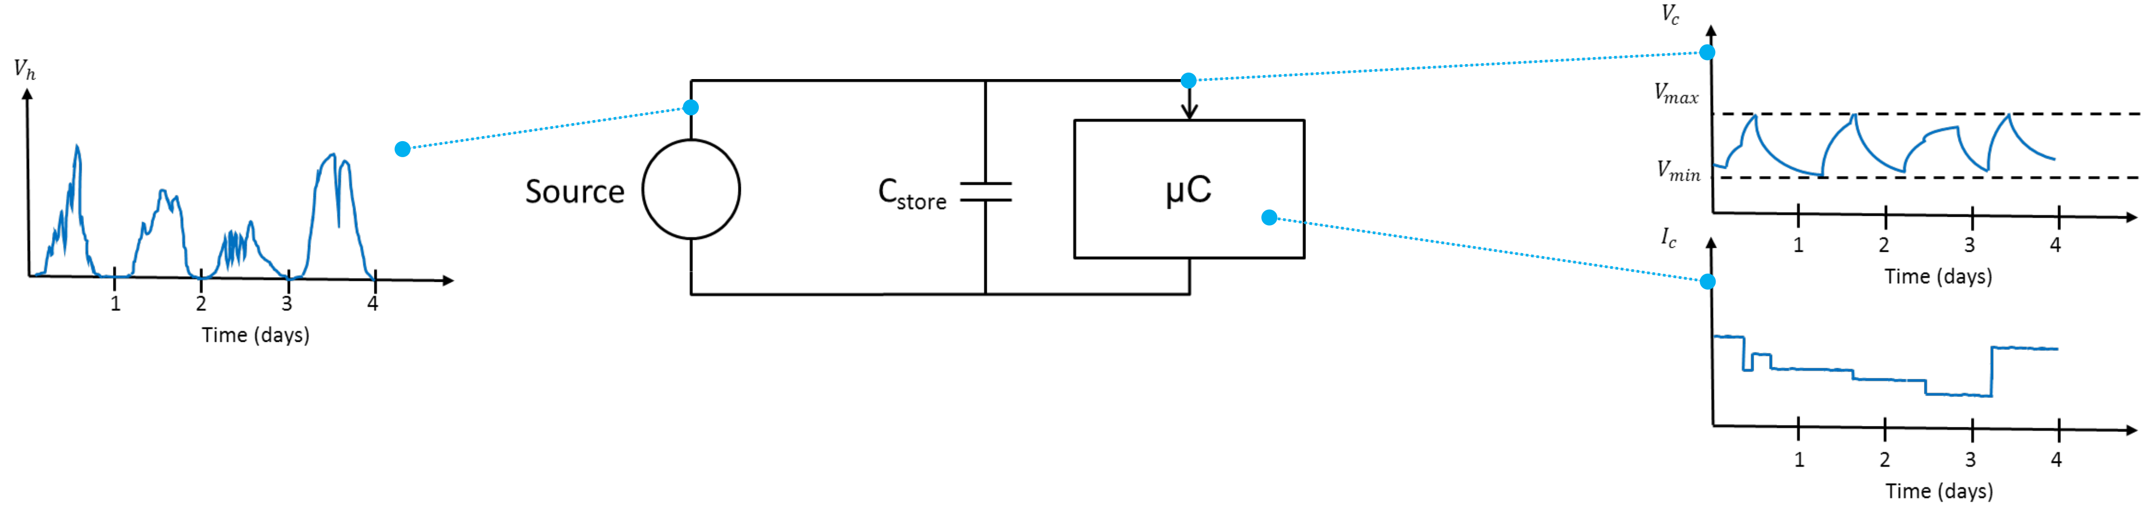

Block Diagram

→ Power Conversion → Computation (Load)

$$\int_{(n-1)\cdot T}^{n\cdot T} P_h(t)dt = \int_{(n-1)\cdot T}^{n\cdot T} P_c(t)dt$$

ENERGY-DRIVEN COMPUTING

- What’s wrong with energy storage and complexity?

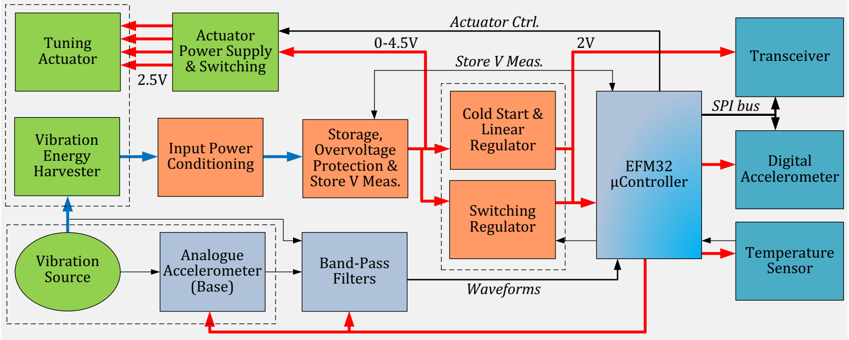

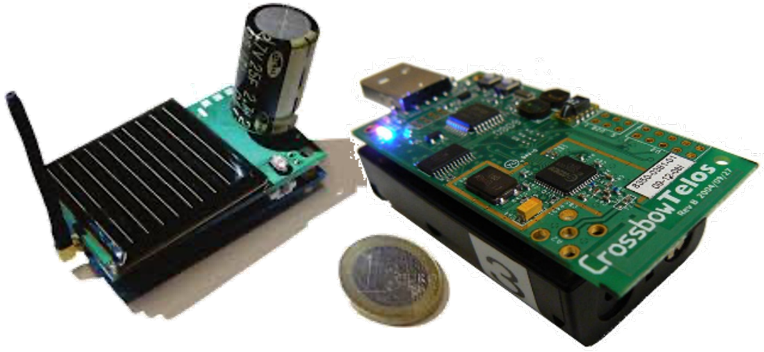

- Emerging applications demanding small dimensions, volumes, weight, cost, etc

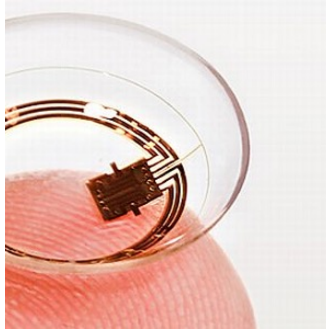

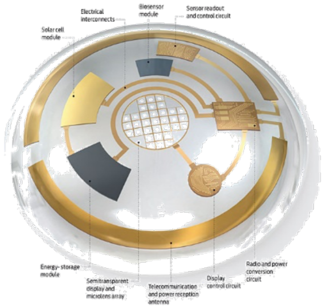

- Properties of energy storage devices (sustainability, maintenance, etc)

ENERGY-DRIVEN COMPUTING

- Rethinking the design of EH systems

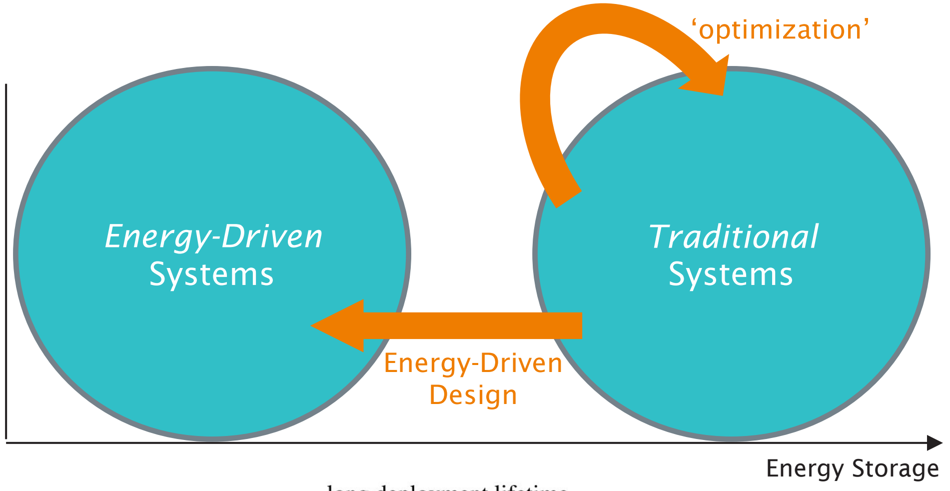

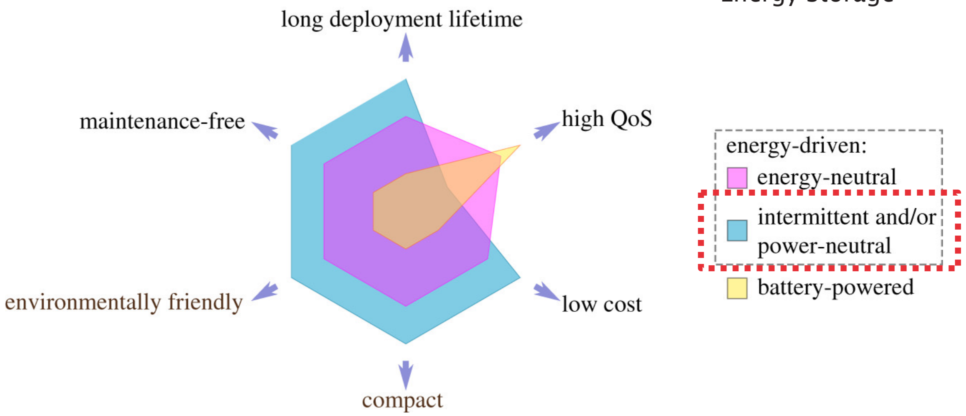

THE ROYAL SOCIETY PUBLISHING

All Journals

Brought to you by UNIVERSITY OF SOUTHAMPTON

Sign in | Search

Home | Content | Information for | About us | Sign up | Propose an issue

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A

MATHEMATICAL, PHYSICAL AND ENGINEERING SCIENCES

You have access

Check for updates

View PDF

Tools | Share

Cite this article

Section

Abstract

1. Introduction

2. Energy- and power-neutral computing

3. Intermittent computing systems

4. Unsolved challenges in energy-driven computing

5. Discussion

Data accessibility

Competing interests

Funding

Footnotes

Discussion

Energy-driven computing

Sivert T. Sliper, Oktay Cetinkaya, Alex S. Weddell, Bashir Al-Hashimi and Geoff V. Merrett

Published: 23 December 2019

https://doi.org/10.1098/rsta.2019.0158

Abstract

For decades, the design of untethered devices has been focused on delivering a fixed quality of service with minimum power consumption, to enable battery-powered devices with reasonably long deployment lifetime. However, to realize the promised tens of billions of connected devices in the Internet of Things, computers must operate autonomously and harvest ambient energy to avoid the cost and maintenance requirements imposed by mains- or battery-powered operation. But harvested power typically fluctuates, often unpredictably, and with large temporal and spatial variability. Energy-driven computers are designed to treat energy-availability as a first-class citizen, in order to gracefully adapt to the dynamics of energy harvesting. They may sleep through periods of no energy, endure periods of scarce energy, and capitalize on periods of ample energy. In this paper, we describe the promise and limitations of energy-driven computing, with an

Details | References | Related | Figures

This Issue

07 February 2020

Volume 378, Issue 2164

Theme issue 'Harmonizing energy-autonomous computing and intelligence'

organised and compiled and edited by Rishad Shafik and Alex Yakovlev

Article Information

DOI: https://doi.org/10.1098/rsta.2019.0

PubMed: 31865884

Published by: Royal Society

Print ISSN: 1364-503X

Online ISSN: 1471-2962

History:

Manuscript accepted 30/09/2019

PDF

Help

University of Southampton

Intermittent Computing

“ *This vision of computing moves in the direction of biology, where some organisms can go dormant in the right conditions. Viruses routinely do this, inserting their DNA into a host’s genome until restored to activity later when their replication is more viable. Similarly, seeds or spores can remain dormant in the environment for years, protected and using no energy, before being reactivated.*

Buchanan, M. Machines learn from biology.
Nature Physics 16, 238 (2020).

INTERMITTENT COMPUTING

INTERMITTENT COMPUTING

www.transient.ecs.soton.ac.uk

Redesigning self-powered embedded systems around energy

- Application compatibility
  - Some applications are naturally compatible
  - Some applications need to be reconsidered from an energy-driven perspective
  - Some application will not be compatible
- Typically IETF Class 1 or 2 constrained networked devices

Monjolo
DeBruin *et al.*, SenSys '13

Energy-driven computing
Sliper *et al.*, Philosophical Transactions of the Royal Society A, 378 (2164), 1-4

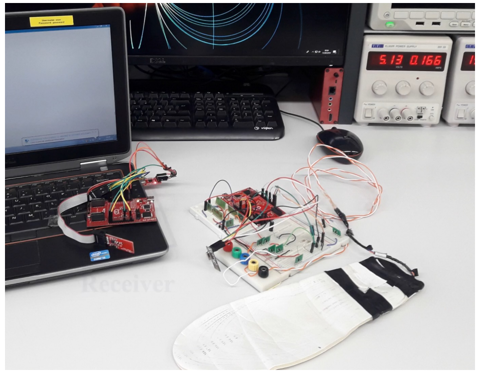

Pedometer
Arreola *et al.*, IEEE CAS '17

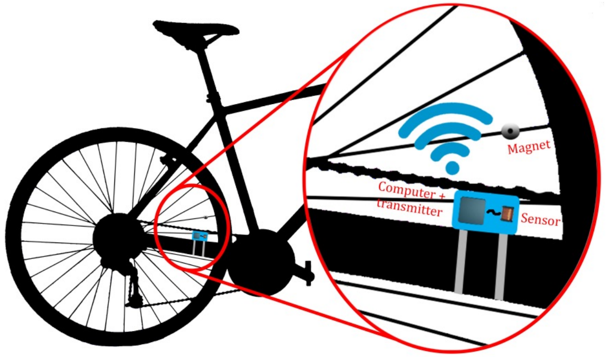

Cycle Computer
Senkans *et al.*, ENSsys '17

Self-Powered Game Boy
De Winkel *et al.*, UbiComp '20

INTERMITTENT COMPUTING

www.transient.ecs.soton.ac.uk

Fail-safe Backup-restore

- Provides a fail-safe if a normally 'always-on' system is interrupted
- Overheads during normal operation need to be minimised

Intermittency expected regularly

- Power interruptions are an expected part of operation
- Overheads during both normal operation and surrounding power interruptions need to be minimised

Energy-driven computing
Sliper *et al.*, Philosophical Transactions of the Royal Society A, 378 (2164), 1-4

Approaches to save volatile state

Static Approaches

- Checkpoints added to application

Task-based Approaches

- Application divided into small tasks
- State saving during transitions

Reactive Approaches

- Save state on power failure

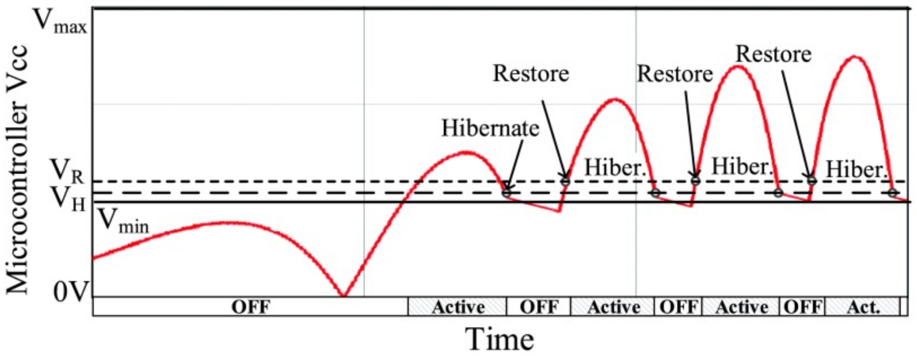

Balsamo *et al.*
IEEE Embedded Systems Letters, 2014

INTERMITTENT COMPUTING

Simplify the system; design for intermittency

Power Supply → Power Conversion → Energy Storage → Power Conversion → Computation (Load)

Compute across power outages

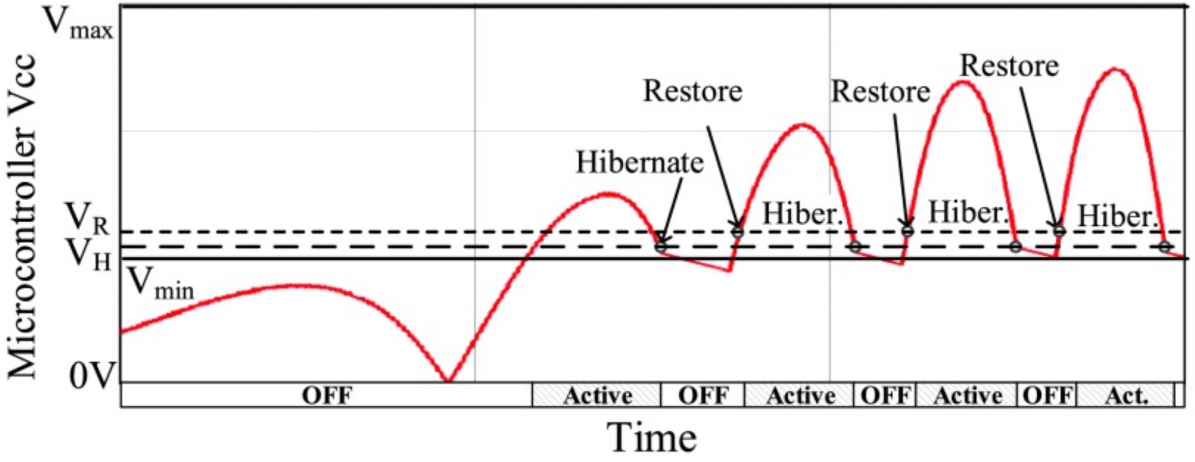

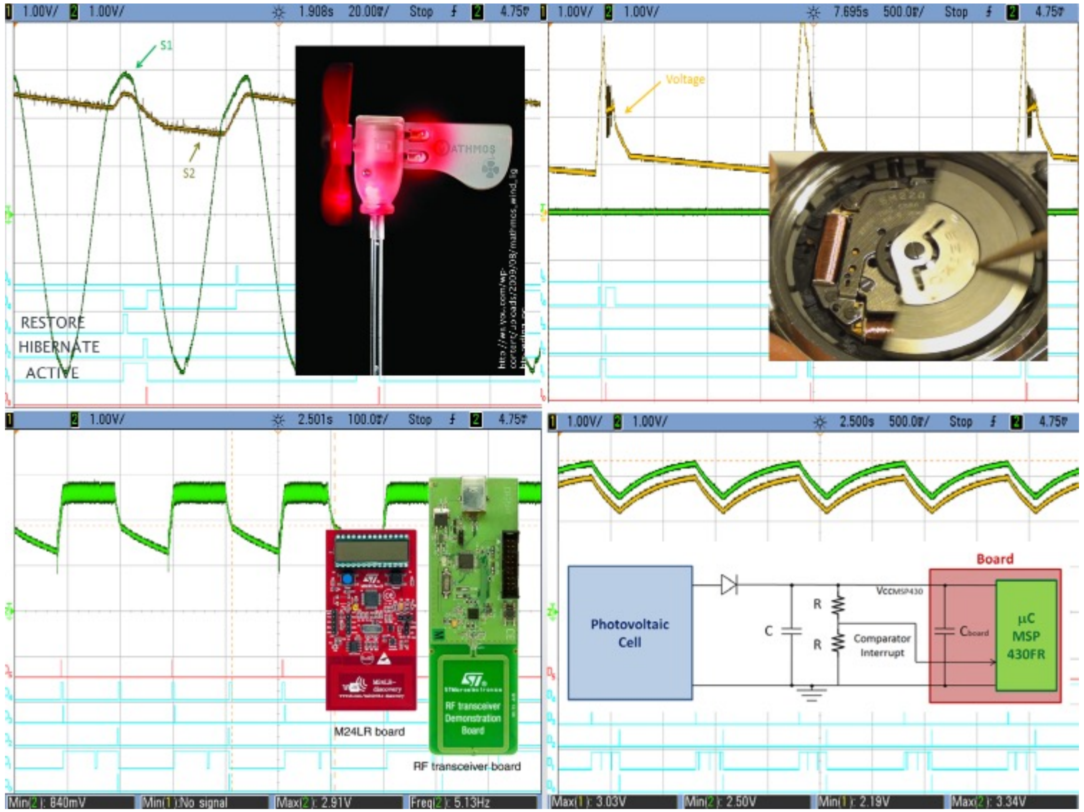

D. Balsamo, A.S. Weddell, A. Das, A. Rodriguez Arreola, D. Brunelli, B.M. Al-Hashimi, G.V. Merrett, L. Benini, (2016) Hibernus++: a self-calibrating and adaptive system for transiently-powered embedded devices. *IEEE TCAD*, 1-13.

CHALLENGES IN INTERMITTENT COMPUTING

www.transient.ecs.soton.ac.uk

Maximising performance

- Maximising forward progress
- Reducing checkpointing overheads

Ensuring state correctness

- Mitigating idempotency violations
- Avoiding live-locks

Design tools

- Masking complexity from programmer
- OS support
- Modelling and simulation

External peripherals

- Consistency

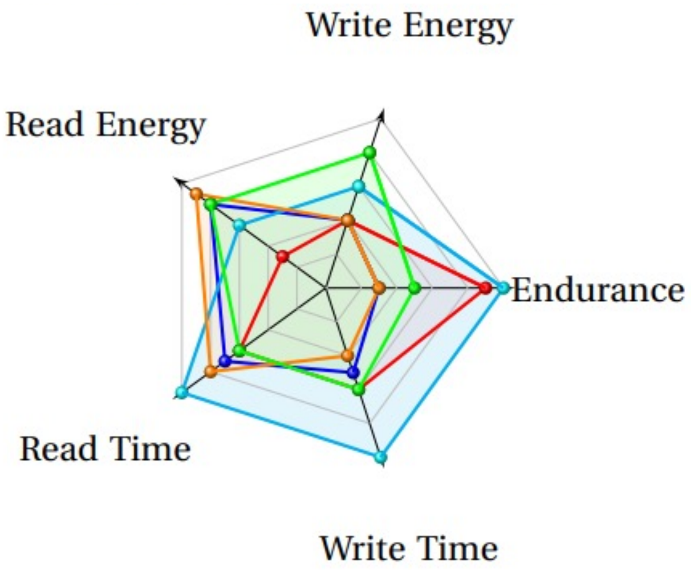

Technology	Write/Read Time (per bit)	Clock cycles (Write)	Clock cycles (Read)
NAND Flash	200μs/25.2μs	4800	605
SRAM	2.2ns/2.1ns	<<1	<<1
FRAM	120ns/120ns	2.9	2.9
STT-MRAM	250ns/10ns	6	0.24
SOT-MRAM	1.4ns/1.1ns	0.03	0.03
ReRAM	10μs/5ns	240	0.12
PCM	150ns/48ns	3.6	1.15

Comparing NVM technologies through the lens of Intermittent Computation, Daulby *et al.*, ENSsys 2020

INTERMITTENT COMPUTING

www.transient.ecs.soton.ac.uk

Compute/Memory

- Self calibration for runtime threshold optimisation (*hibernus++*)
- Adaptive restore based on EH properties (*hibernus++*)
- Efficient state retention (Selective Policies, *ManagedState*)
- Fine-grained power adaption (*PowerNeutrality*)

Peripherals/Sensors/Communication

- Hibernation and restore of external peripheral state (*RESTOP*)
- Support for communication and mesh networking

Applications/Users/Interfaces/Design Tools

- Application case studies, e.g. cycle computer, fitness monitor, wall clock, etc
- Design tools, e.g. *ENSPECT*, *FUSED*, Device Sizing, Support for *Arm Mbed*

Reducing Volatile State

STATE RETENTION POLICIES

Allocated State (symmetric/asymmetric NVMs)

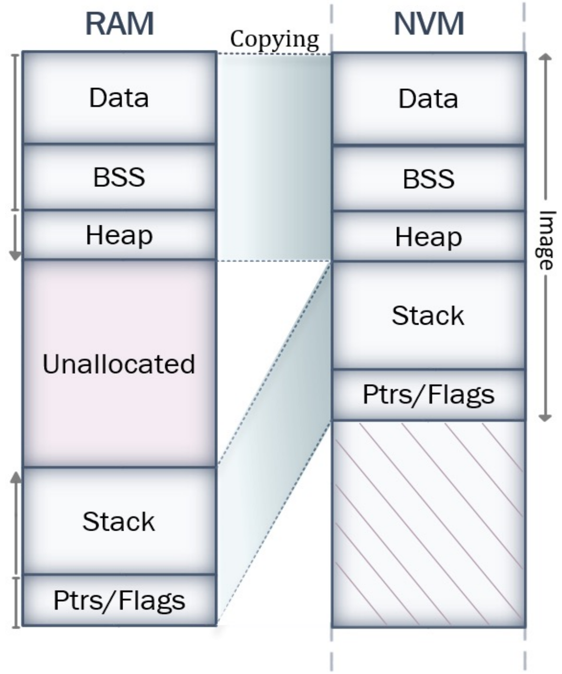

NVMs with no erase cost

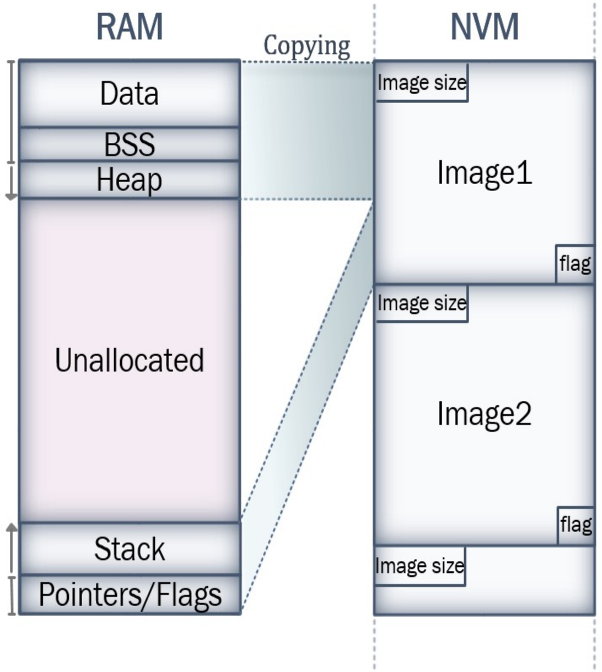

NVMs with erase cost

Multiple Blocks (asymmetric NVMs)

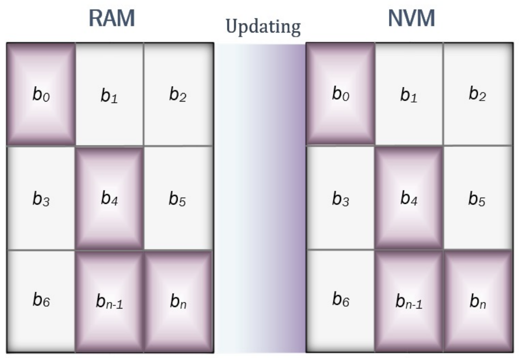

NVMs with no erase cost

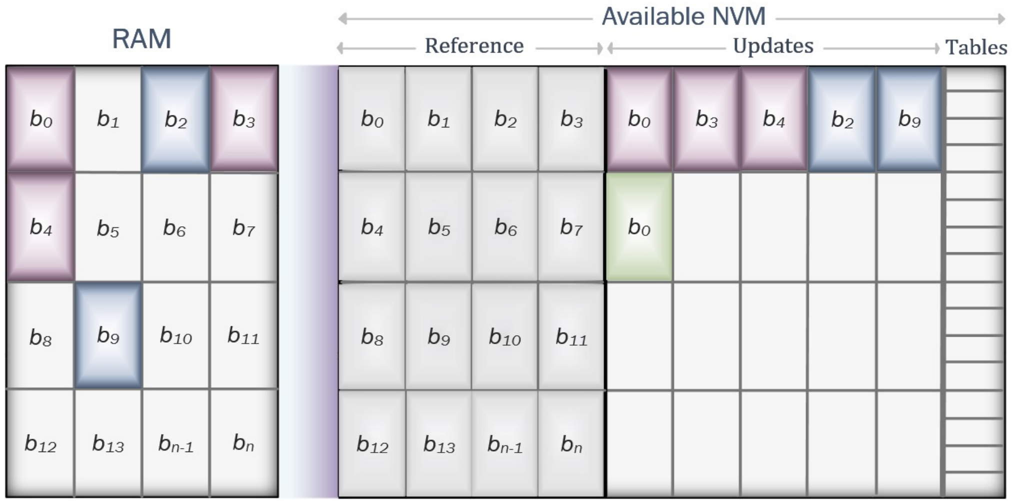

NVMs with no erase cost

Verykios, Theodoros D., Balsamo, Domenico and Merrett, Geoff V. (2018) **Selective policies for efficient state retention in transiently-powered embedded systems: exploiting properties of NVM technologies.** *Sustainable Computing: Informatics and Systems.*

“MANAGED STATE”

A lightweight memory management software-layer that loads pages of memory on demand, and saves only modified pages.

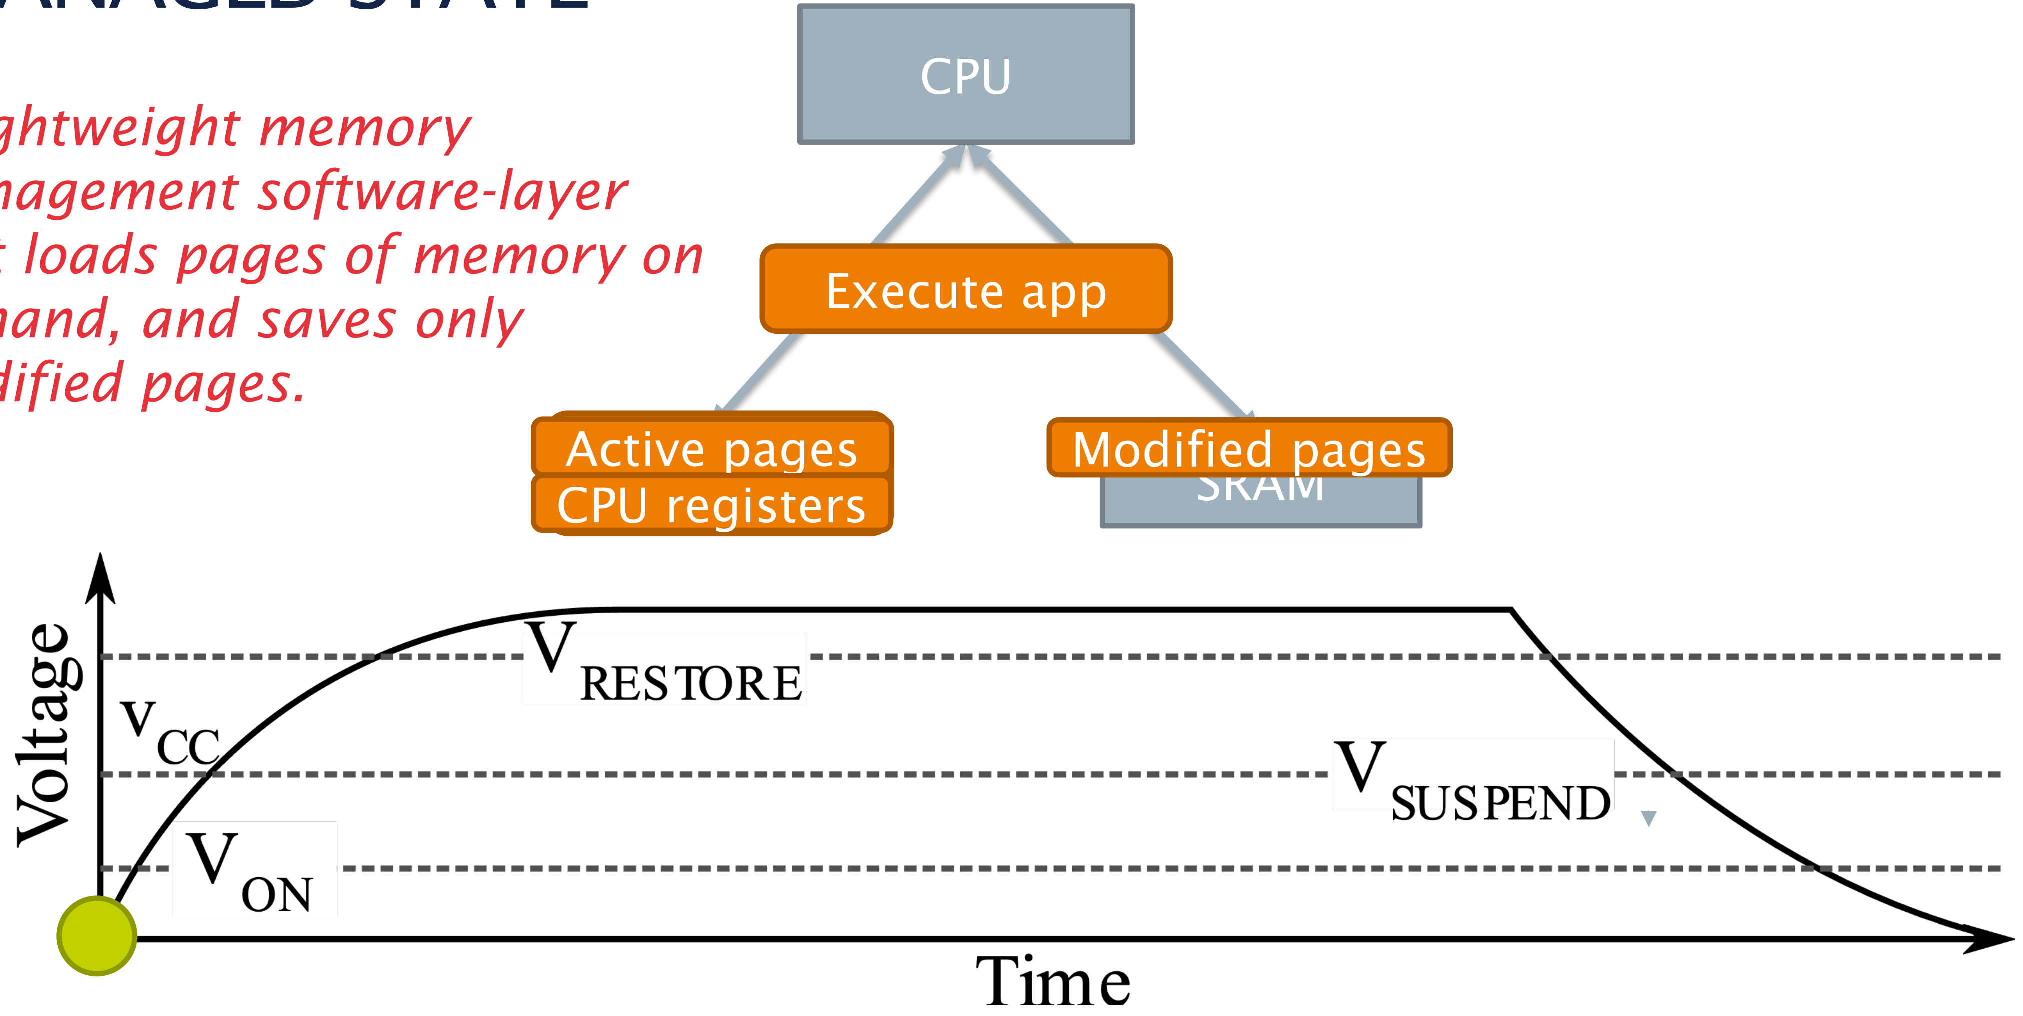

Sliper, Sivert T., Balsamo, Domenico, Nikoleris, Nikos, Wang, William, Weddell, Alexander and Merrett, Geoff (2019) **Efficient state retention through paged memory management for reactive transient computing**. *Design Automation Conference*, Las Vegas, United States. 02 - 06 Jun 2019. 6 pp .

MANAGED STATE: RESULTS

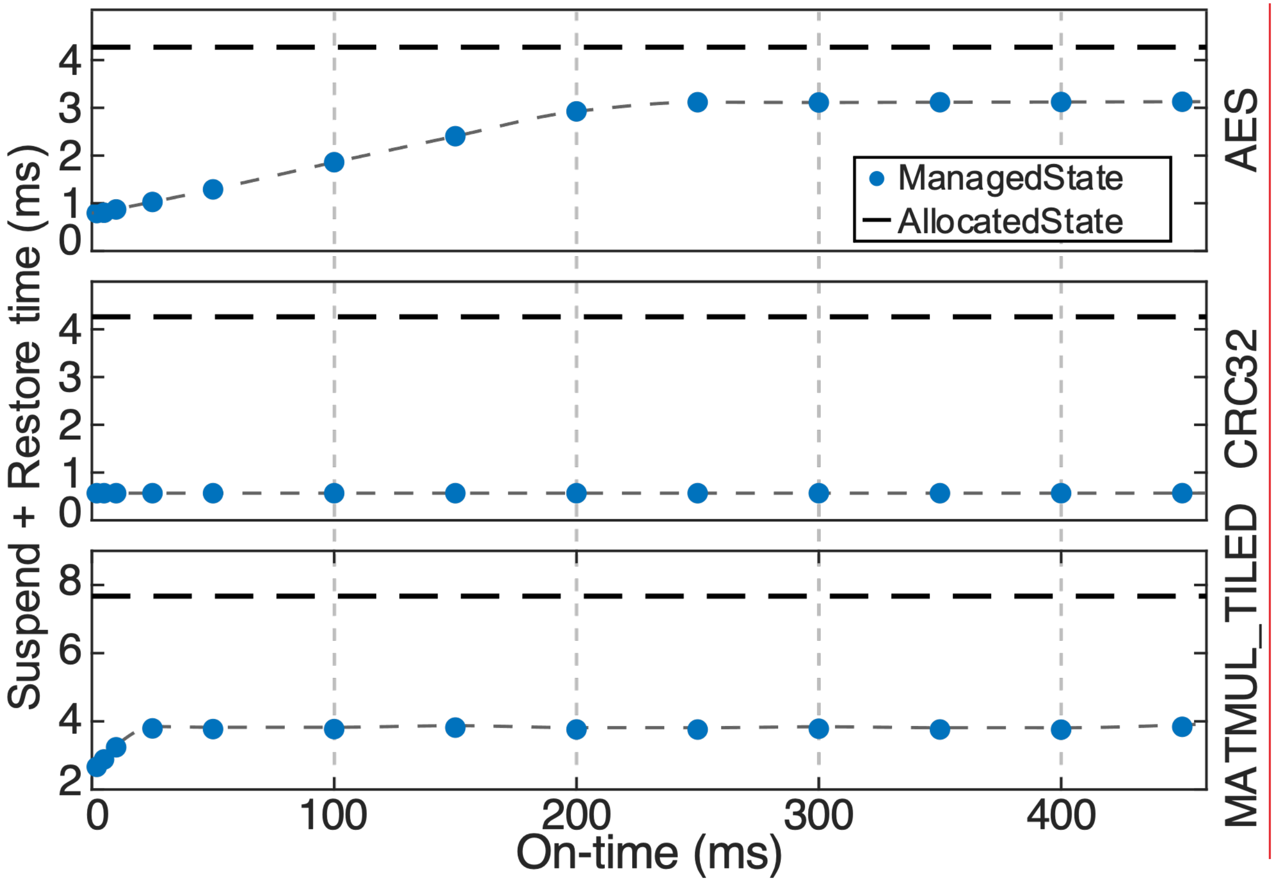

Suspend + Restore time in relation to on-time

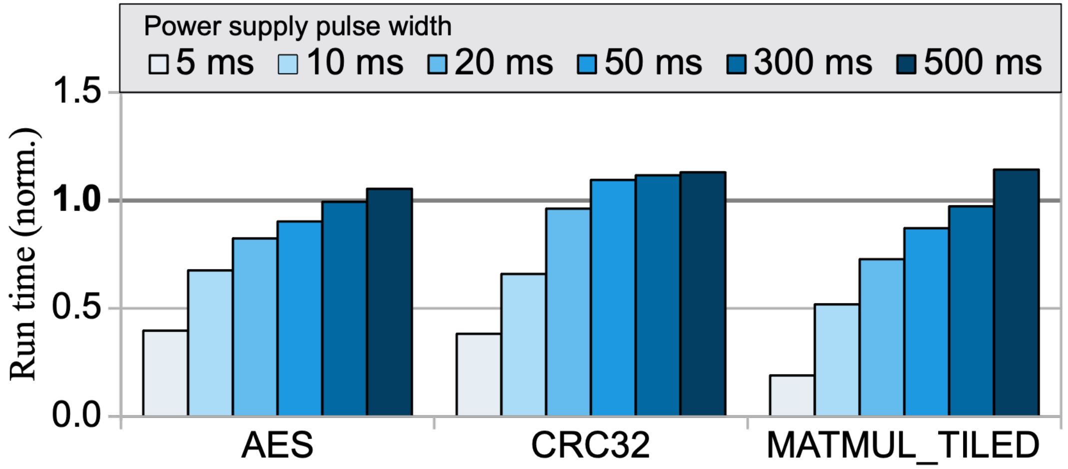

Completion time relative to AllocatedState. Powered by voltage pulses of variable width

Sliper, Sivert T., Balsamo, Domenico, Nikoleris, Nikos, Wang, William, Weddell, Alexander and Merrett, Geoff (2019) **Efficient state retention through paged memory management for reactive transient computing.** *Design Automation Conference*, Las Vegas, United States. 02 - 06 Jun 2019. 6 pp .

HARDWARE SUPPORT - MEMIC

- A memory system targeting modern MRAM-enabled low-power microcontrollers.

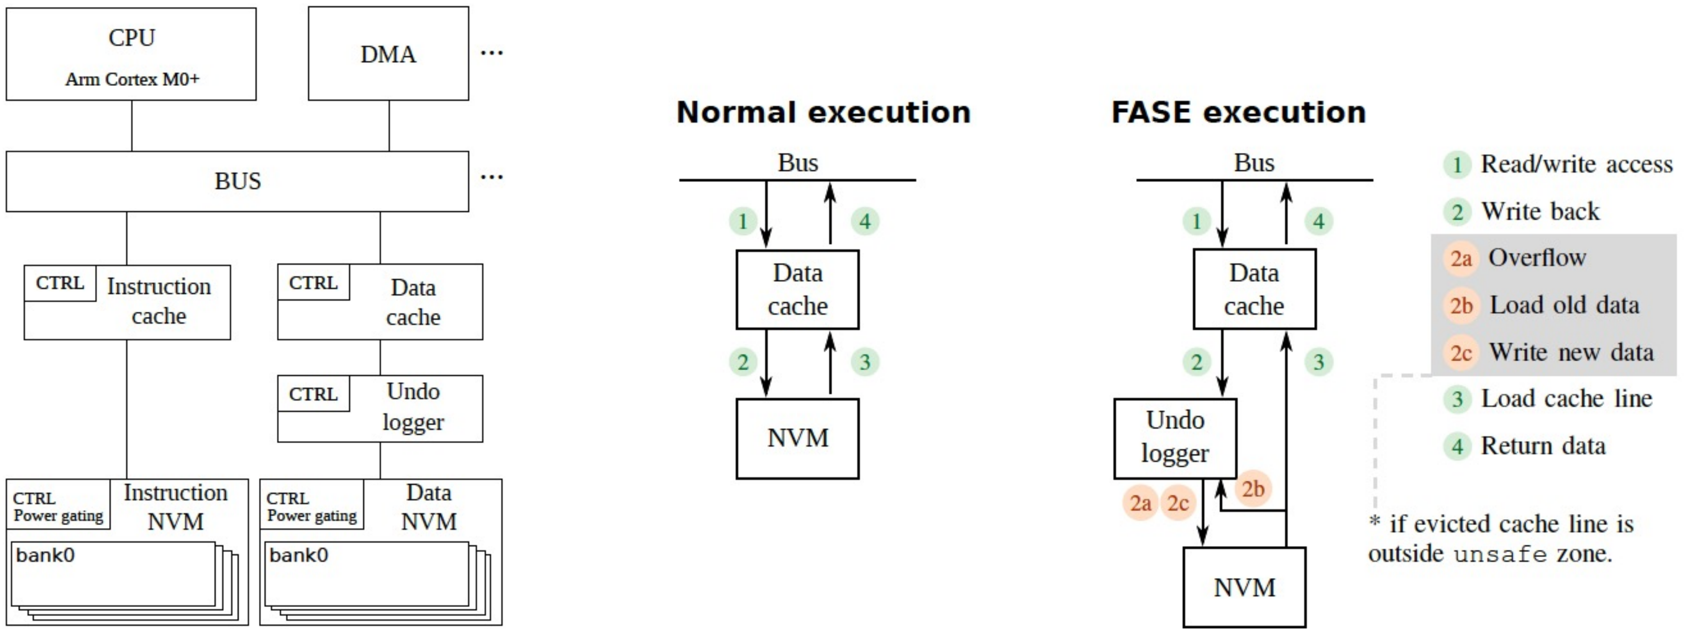

(a) Top-level architecture of *MEMIC*. (b) Data access during normal execution. (c) Data access during FASE execution.

- Improves workload completion time by 13–39%, using 13–39% less energy, and operates under condition where state-of-the-art systems fail

S. Sliper et al., Pragmatic Memory-System Support for Intermittent Computing using Emerging Non-Volatile Memory, IEEE TCAD (in press)

University of Southampton

Intermittent Networking

Talking to an empty room

S.W.C. Bing *et al.* Demo Abstract: An Energy-driven Wireless Bicycle Trip Counter with Zero Energy Storage. ACM SenSys '18.

EXISTING NETWORKING APPROACHES

- High power **coordinator nodes/hub**
  - Insufficient range
  - Inflexible and single point of failure

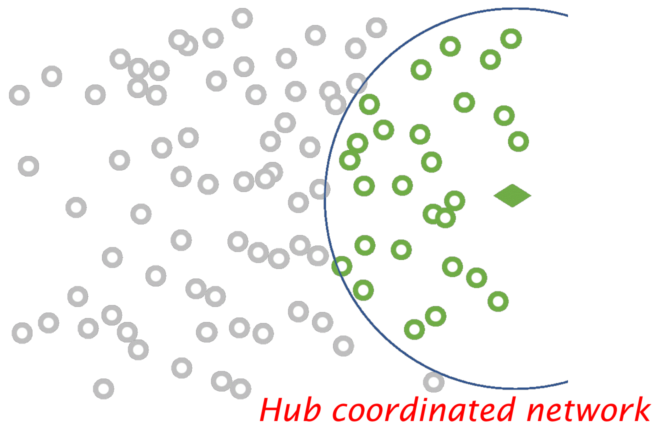

- **Scheduled communication** slots across the network
  - Requires accurate timers

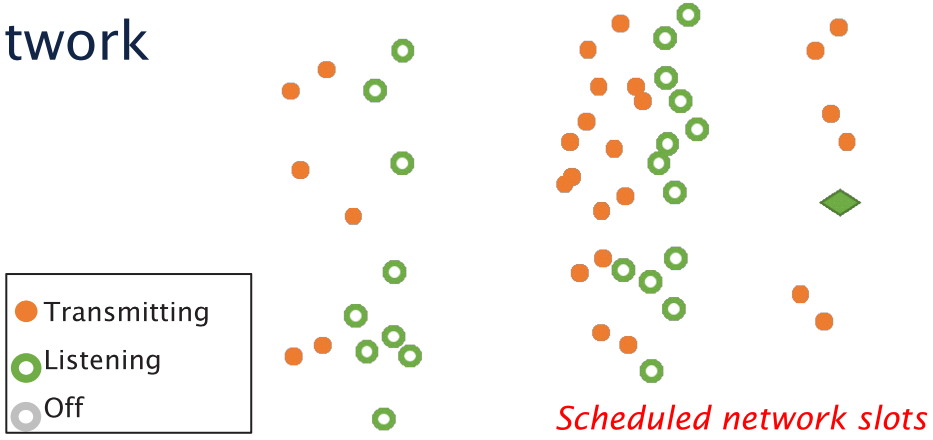

E. Longman, M. El-Hajjar, O. Cetinkaya, G.V. Merrett, (2022) **Mesh networking for intermittently-powered devices: Architecture and challenges.** *IEEE Network.*

WAKE-UP RADIOS

We model Wake-up Radio enabled nodes

- Wake-up radios allow **µW listening** for messages
  - Reduced sensitivity so higher transmit consumption
  - No synchronization required
- Transmitting node sends **high power wake-up**, then data

Node Topology

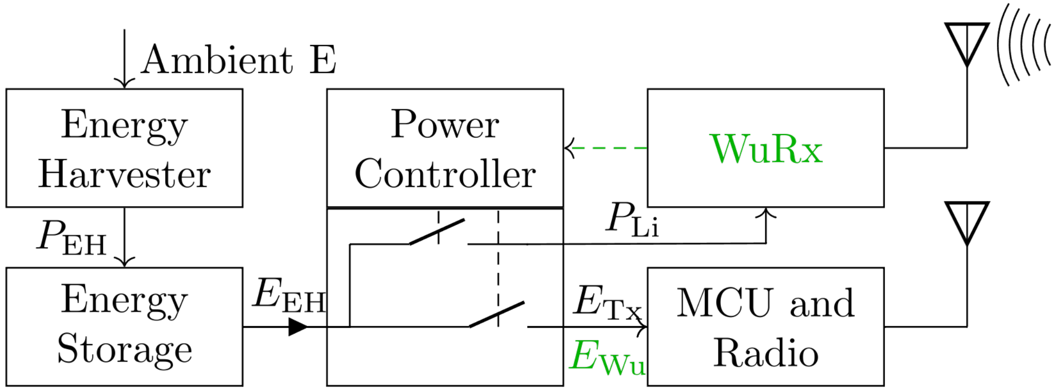

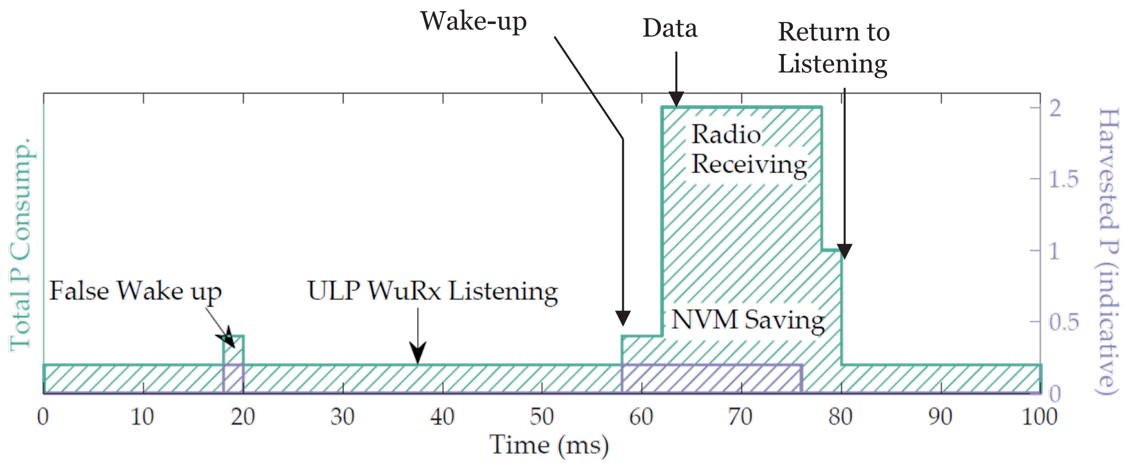

Receiver power consumption

E. Longman, O. Cetinkaya, M. El-Hajjar, G.V. Merrett, (2021) **Wake-up radio-enabled intermittently-powered devices for mesh networking: A power analysis**. *IEEE Consumer Communications and Networking Conference*, 09 - 12 Jan 2021.

HOMOGENEOUS INTERMITTENT MESH NETWORKING

- Sensor node powered from **Small EH source** of P_{EH}
- **Minimal** energy **storage**
- **Homogenous nodes** operate **intermittently** as EH allows
  - Represented by energy delivered in "bursts" E_{EH}
  - Tx & Rx have constant energy
  - WuRx has constant listening power, P_{Li}, less than 100"µW"
- Aim to **maximize Goodput, G**

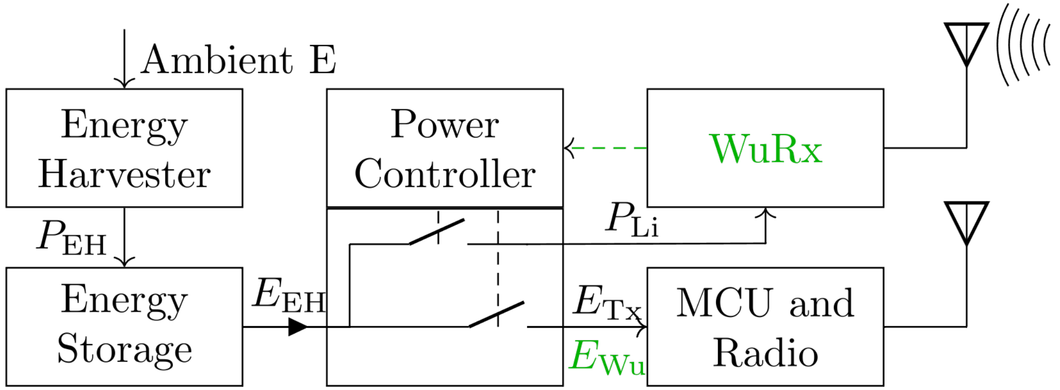

Node Topology

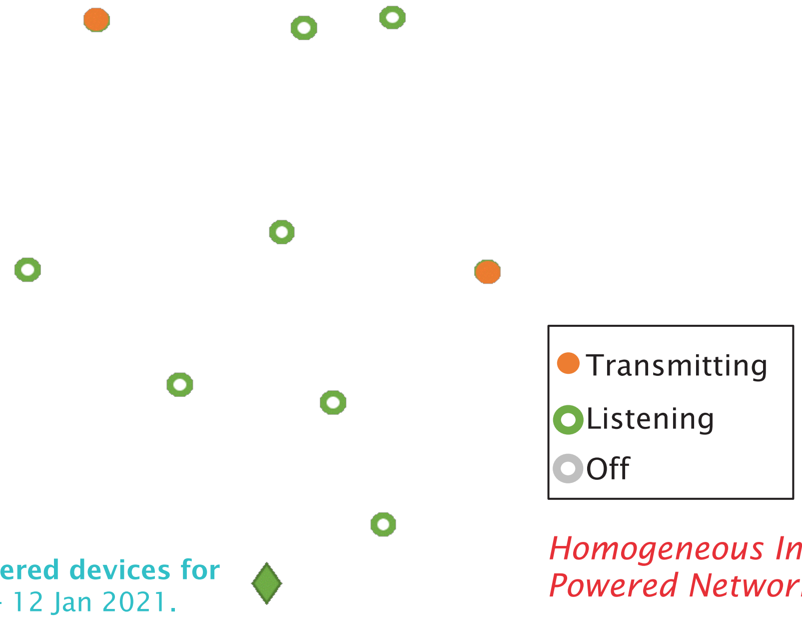

Homogeneous Intermittently-Powered Network

E. Longman, O. Cetinkaya, M. El-Hajjar, G.V. Merrett, (2021) **Wake-up radio-enabled intermittently-powered devices for mesh networking: A power analysis**. *IEEE Consumer Communications and Networking Conference*, 09 - 12 Jan 2021.

DON'T OVER TRANSMIT, TRANSMIT OPTIMALLY

- Higher energy transmission → Lower power wake-up radio required
- When $E_{Wu} = 0$, **equally splitting E_{EH} maximizes** the goodput
- Considering **wake-ups reduces** the **optimum** transmission load

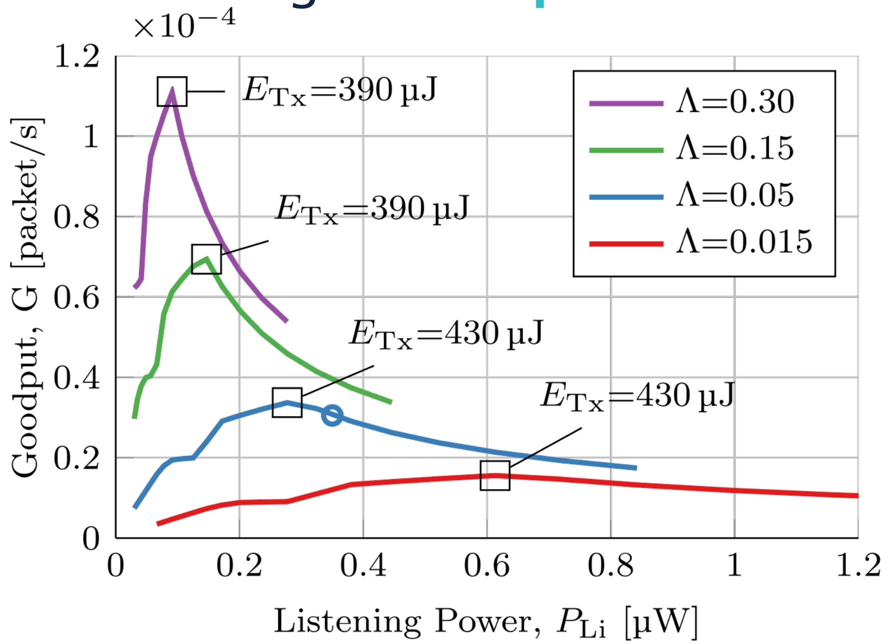

Maximum goodput for Tx and Rx Configurations

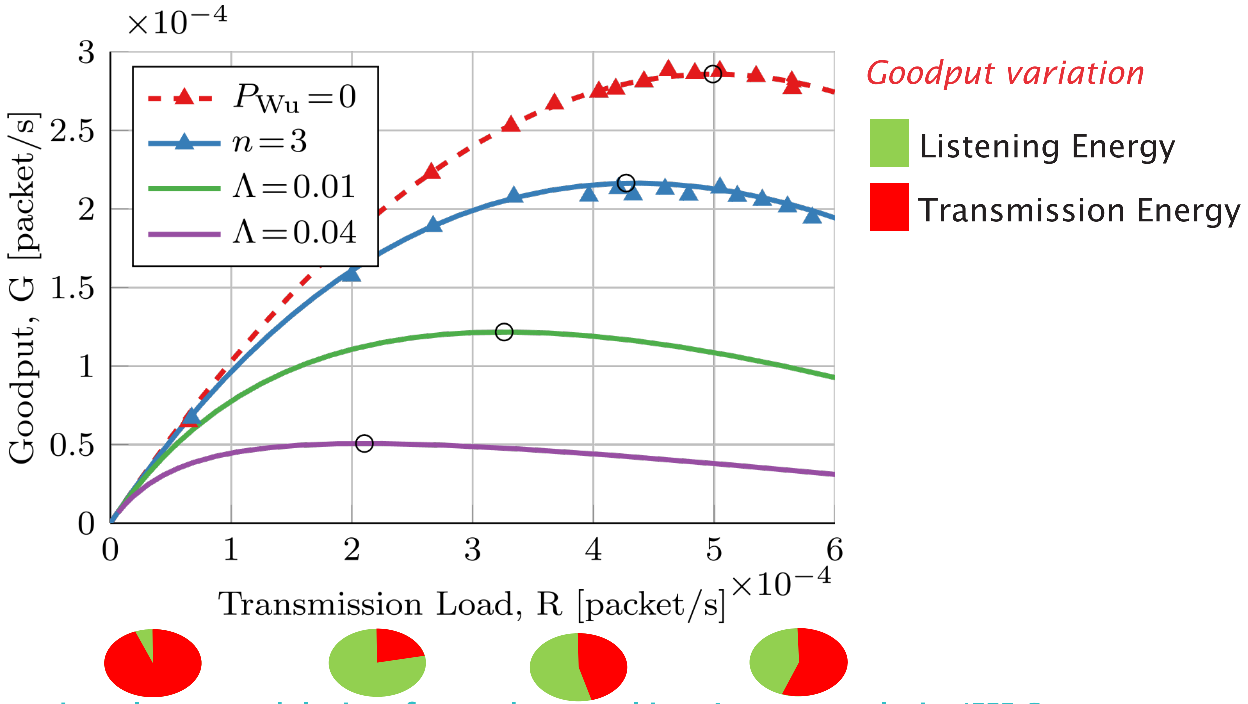

E. Longman, O. Cetinkaya, M. El-Hajjar, G.V. Merrett, (2021) **Wake-up radio-enabled intermittently-powered devices for mesh networking: A power analysis.** *IEEE Consumer Communications and Networking Conference*, 09 - 12 Jan 2021.

MULTIHOP ROUTING IN INTERMITTENT NETWORKS

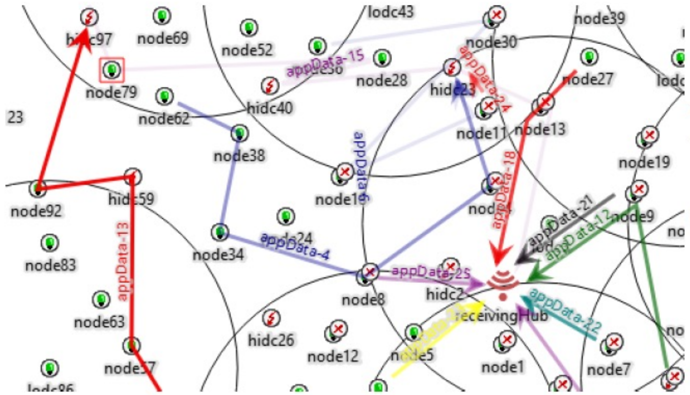

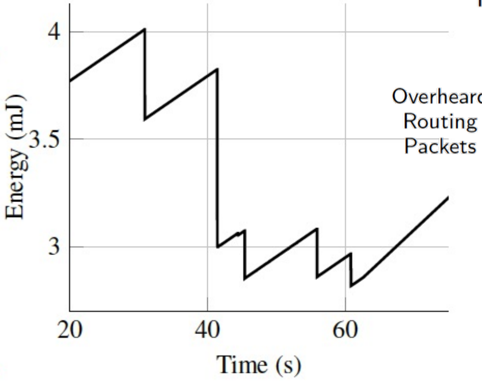

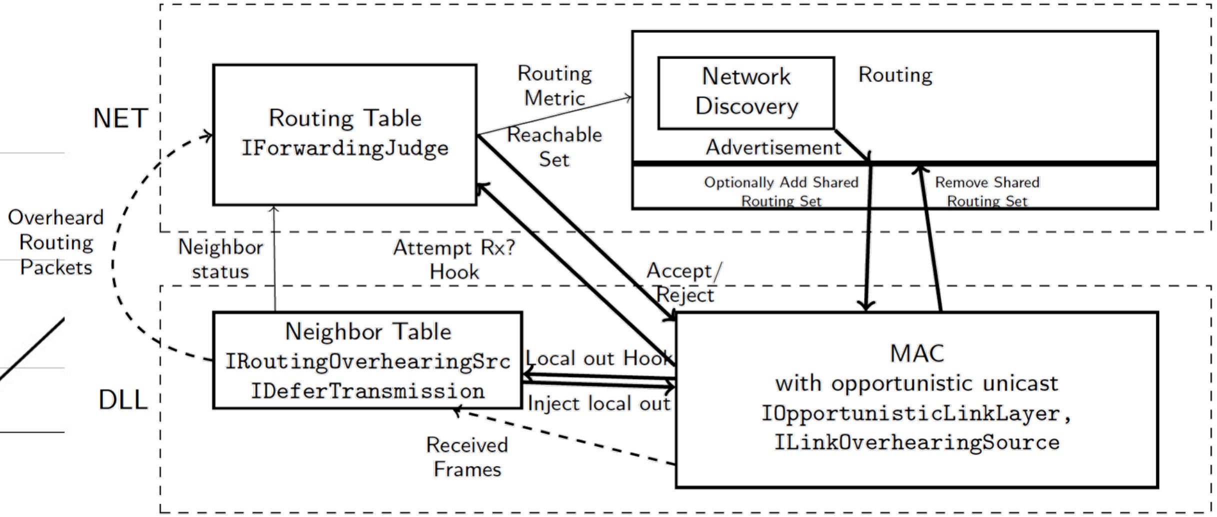

E. Longman, M. El-Hajjar, G.V. Merrett, (2021) **Intermittent opportunistic routing components for the INET framework**. In *8th OMNeT++ Community Summit 2021*. 5 pp .

Power-Neutral Computing

Energy neutrality under a microscope!

POWER NEUTRAL COMPUTING

- In Power-Neutral computing, $P_c(t) \cong P_h(t)$
- We can approximate power-neutral behaviour if $V_C(t) \approx k \;, \;\forall t$

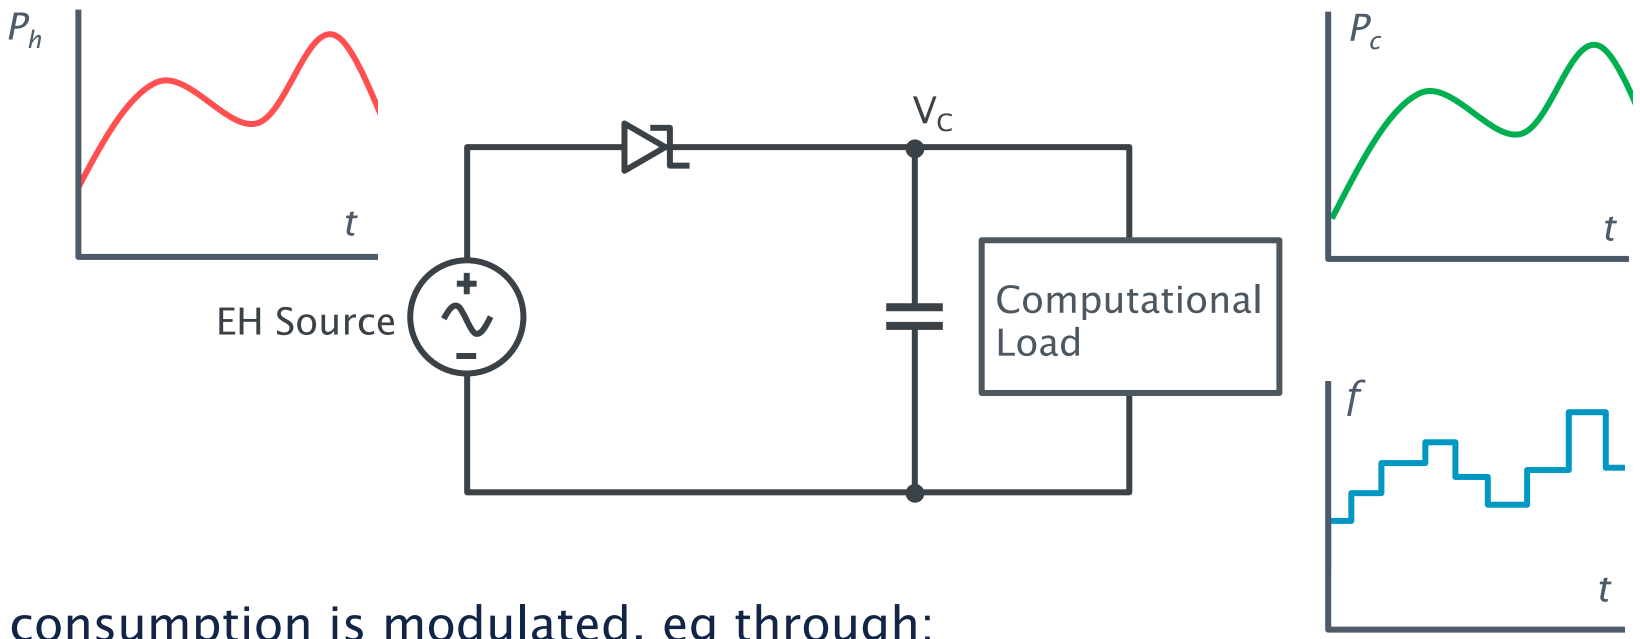

- Power consumption is modulated, eg through:
  - Core frequency and/or voltage
  - Power gating processor elements

B. Fletcher, D. Balsamo, G.V. Merrett, (2017) **Power neutral performance scaling for energy harvesting MP-SoCs**. In *2017 Design, Automation & Test in Europe Conference & Exhibition (DATE).*

POWER NEUTRAL COMPUTING

- What happens if V_C remains constant ($V_C(t) \approx k \ , \ \forall t$)?
- MPPT approaches are needed as $V_C(t) \neq V_{H\_MPP}(t) \ , \ \forall t$
- 'Software-only' MPPT modulates k: $V_C(t) = V_{H\_MPP}\ (t)$

Estimating V_{cc} and P_c, we can identify MPP

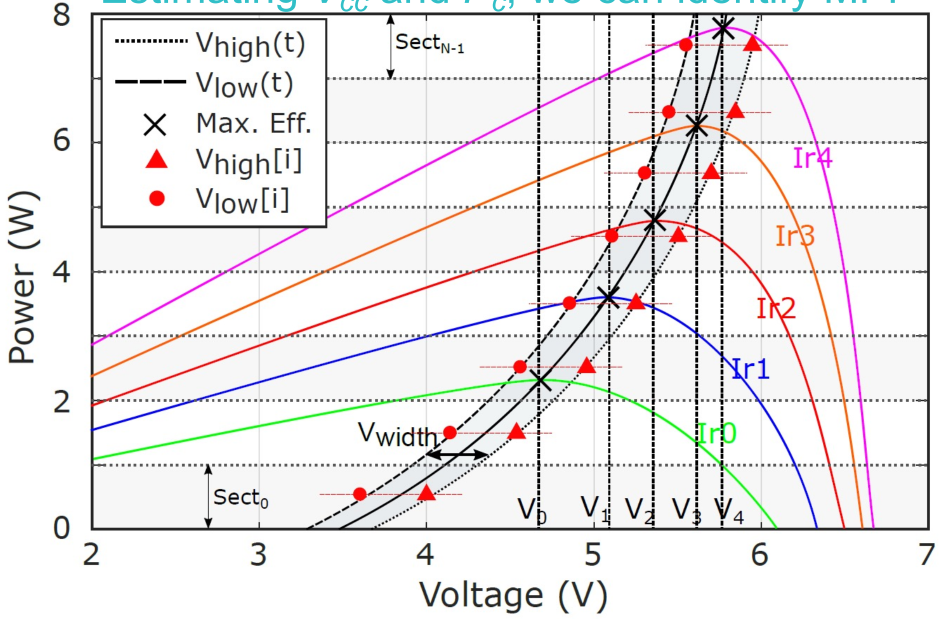

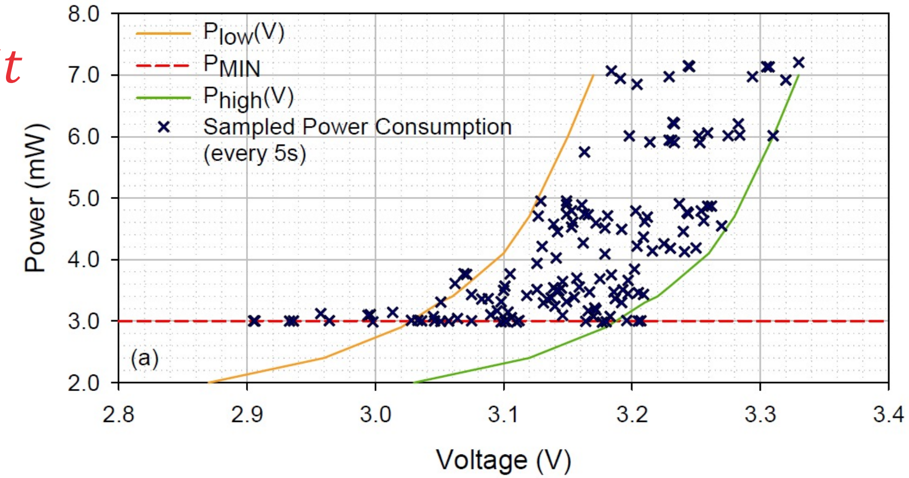

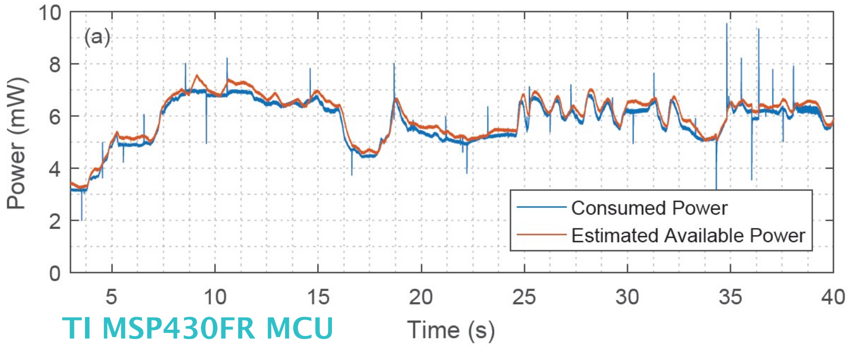

TI MSP430FR MCU

D. Balsamo, B. Fletcher, A.S. Weddell, G. Karatziolas, B.M. Al-Hashimi and G.V. Merrett, (2019) **Power neutral performance scaling with intrinsic MPPT for energy harvesting computing systems.** *ACM Transactions on Embedded Computing Systems*, 17 (6), 93:1-93:25

EXPERIMENTAL RESULTS

Operating from a PV Energy Harvester

- Power-Neutral Behaviour

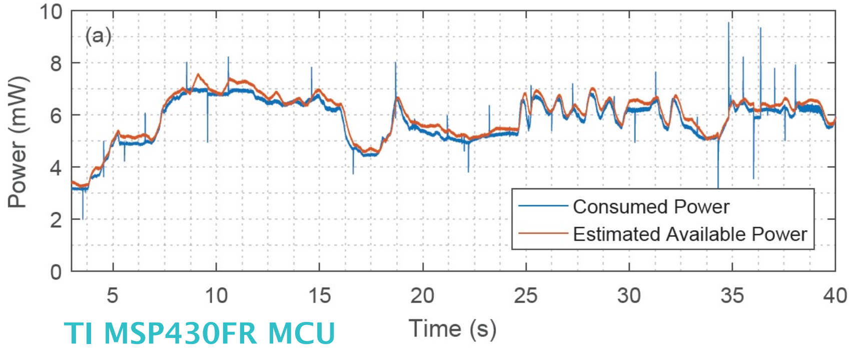

TI MSP430FR MCU

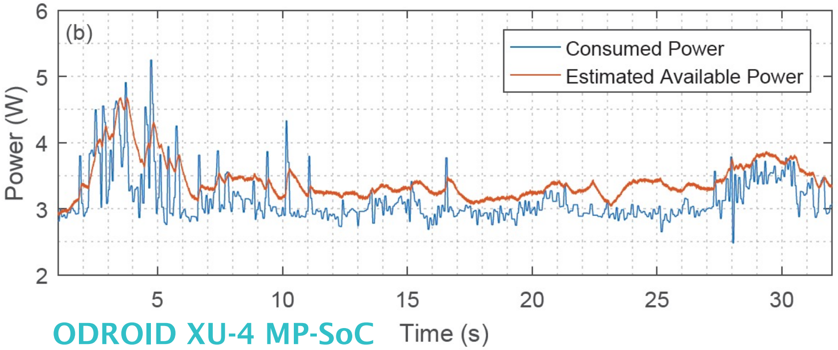

ODROID XU-4 MP-SoC

- Application Forward Progress

Performance Scaling Technique	Number of FFTs per second
Static Approach [16]	1.07
MCU Power-Neutral Approach [17]	1.65
Momentum	1.83

TI MSP430FR MCU

Performance Scaling Technique	Billions of Instructions Per Second
Linux Powersave	0.69
MP-SoC Power-neutral Approach [2]	1.17
Momentum	1.27

ODROID XU-4 MP-SoC

D. Balsamo, B. Fletcher, A.S. Weddell, G. Karatziolas, B.M. Al-Hashimi and G.V. Merrett, (2019) **Power neutral performance scaling with intrinsic MPPT for energy harvesting computing systems.** *ACM Transactions on Embedded Computing Systems*, 17 (6), 93:1-93:25

Design Tools for Energy-Driven Systems

CLOSED-LOOP PERFORMANCE & ENERGY SIMULATION

Download: www.arm.ecs.soton.ac.uk/technologies/fused

Developing energy-driven computing systems is difficult

- Operation driven directly by availability of energy
- Introspection/debug inevitably affects stored/harvested energy (and therefore operation)
- Typical embedded systems development can’t cope with a DUT that frequently powers off
- Repeatability of EH is problematic

Fused: Closed-loop Performance and Energy Simulation of Embedded Systems
Sliper *et al.*, ISPASS 2020

FUSED (Full-system Simulation of Energy-Driven Computers)

- Open source full-system simulator for energy-driven computers
- SystemC for digital and mixed-signal simulation, modelling microcontroller and mixed-signal circuitry
- Models power supply/consumption and execution in a closed loop, modelling the interaction between.
- Accurate power model obtained correlates microarchitectural events with real power measurements
- Enables hardware-software codesign and design space exploration.

FUSED

Download: www.arm.ecs.soton.ac.uk/technologies/fused

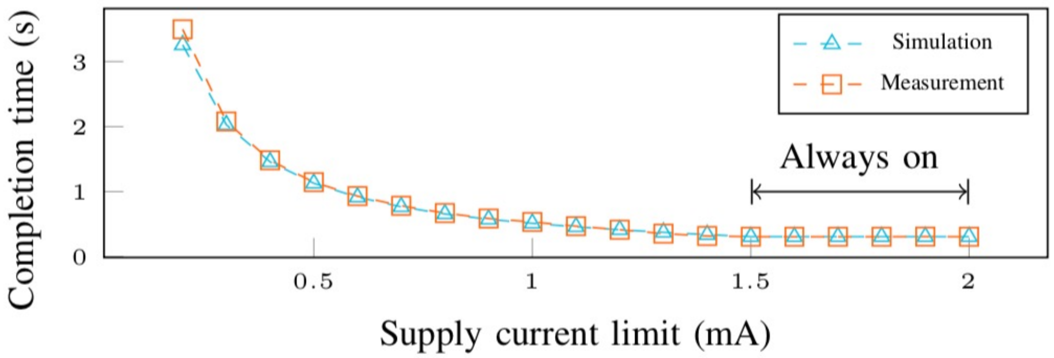

Fig. 10. Completion time of AES encryption when running intermittently, powered by a current-limited power source.

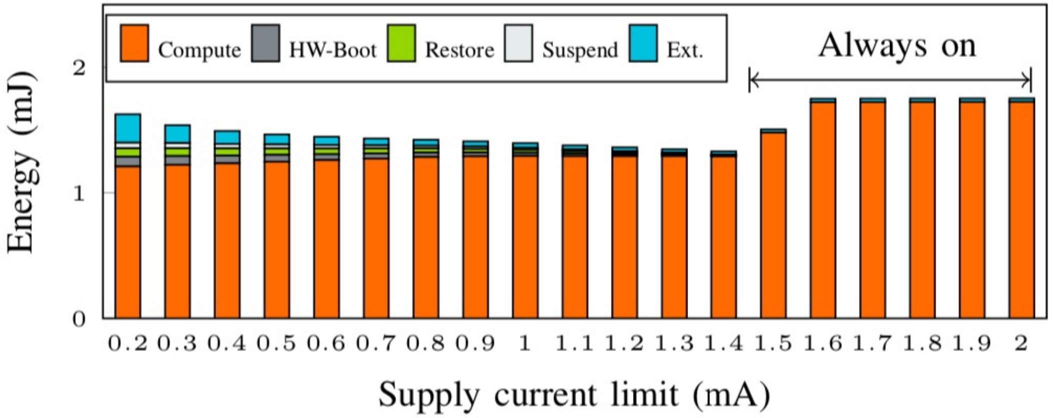

Fig. 11. Full-system energy consumption when running AES encryption intermittently, powered by current-limited power source. The energy consumption is divided into stacked bars for the external circuitry (*ext.*), hardware boot (*HW-Boot*), and the operational phases *restore*, *compute* and *suspend*.

Sivert Sliper, William Wang, Nikos Nikoleris, Alex S. Weddell, Geoff V. Merrett, "**Fused: Closed-loop Performance and Energy Simulation of Embedded Systems**," *ISPASS 2020*

FUSED

Download: www.arm.ecs.soton.ac.uk/technologies/fused

Virtual prototyping of complete EH systems

- Ambient energy environment
- Energy harvesters
- Power management circuitry
- On/off-chip peripherals

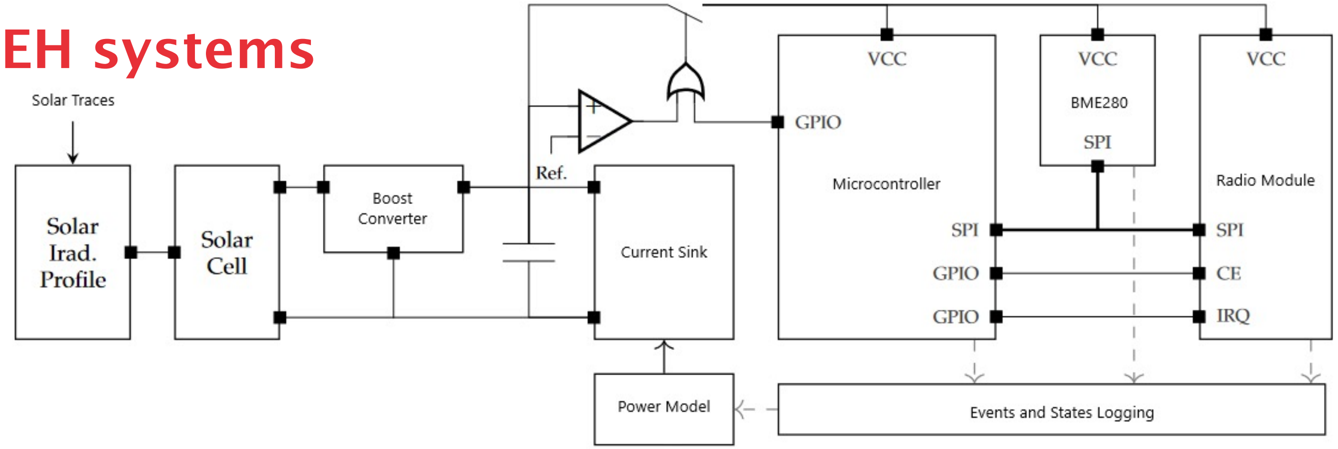

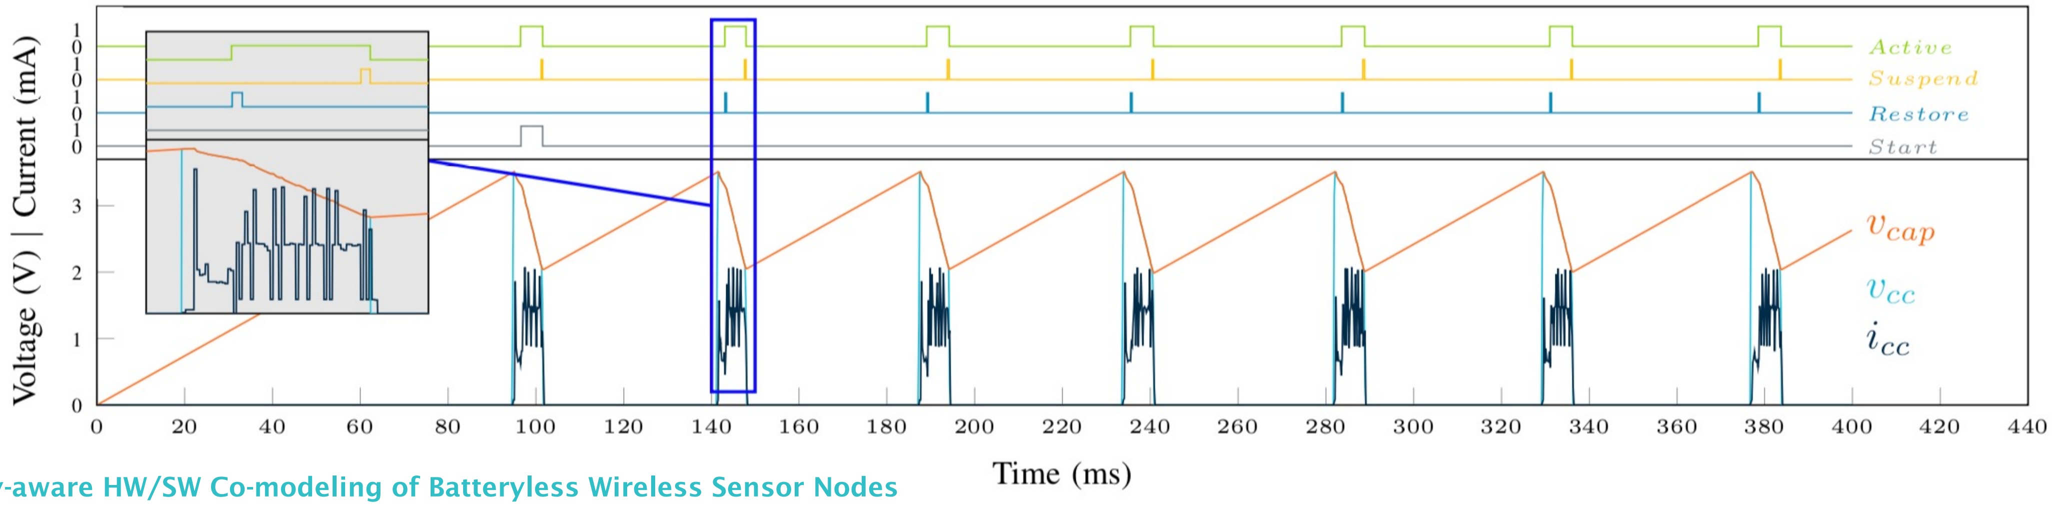

Energy-aware HW/SW Co-modeling of Batteryless Wireless Sensor Nodes
Wong *et al.*, ENSsys 2020

DISCUSSION

- We need to rethink the way that we design self-powered *systems*
- *Most* work has focused on compute.
- IC approaches need to adapt to the range of emerging NVM technologies
- Advances still needed in design tools.
- Networking and communication pose significant challenge to ED systems
  - Links are inherently (unpredictably) intermittent, and typically unsynchronized with each other
  - Is IP-networking in an ED system a realistic target?

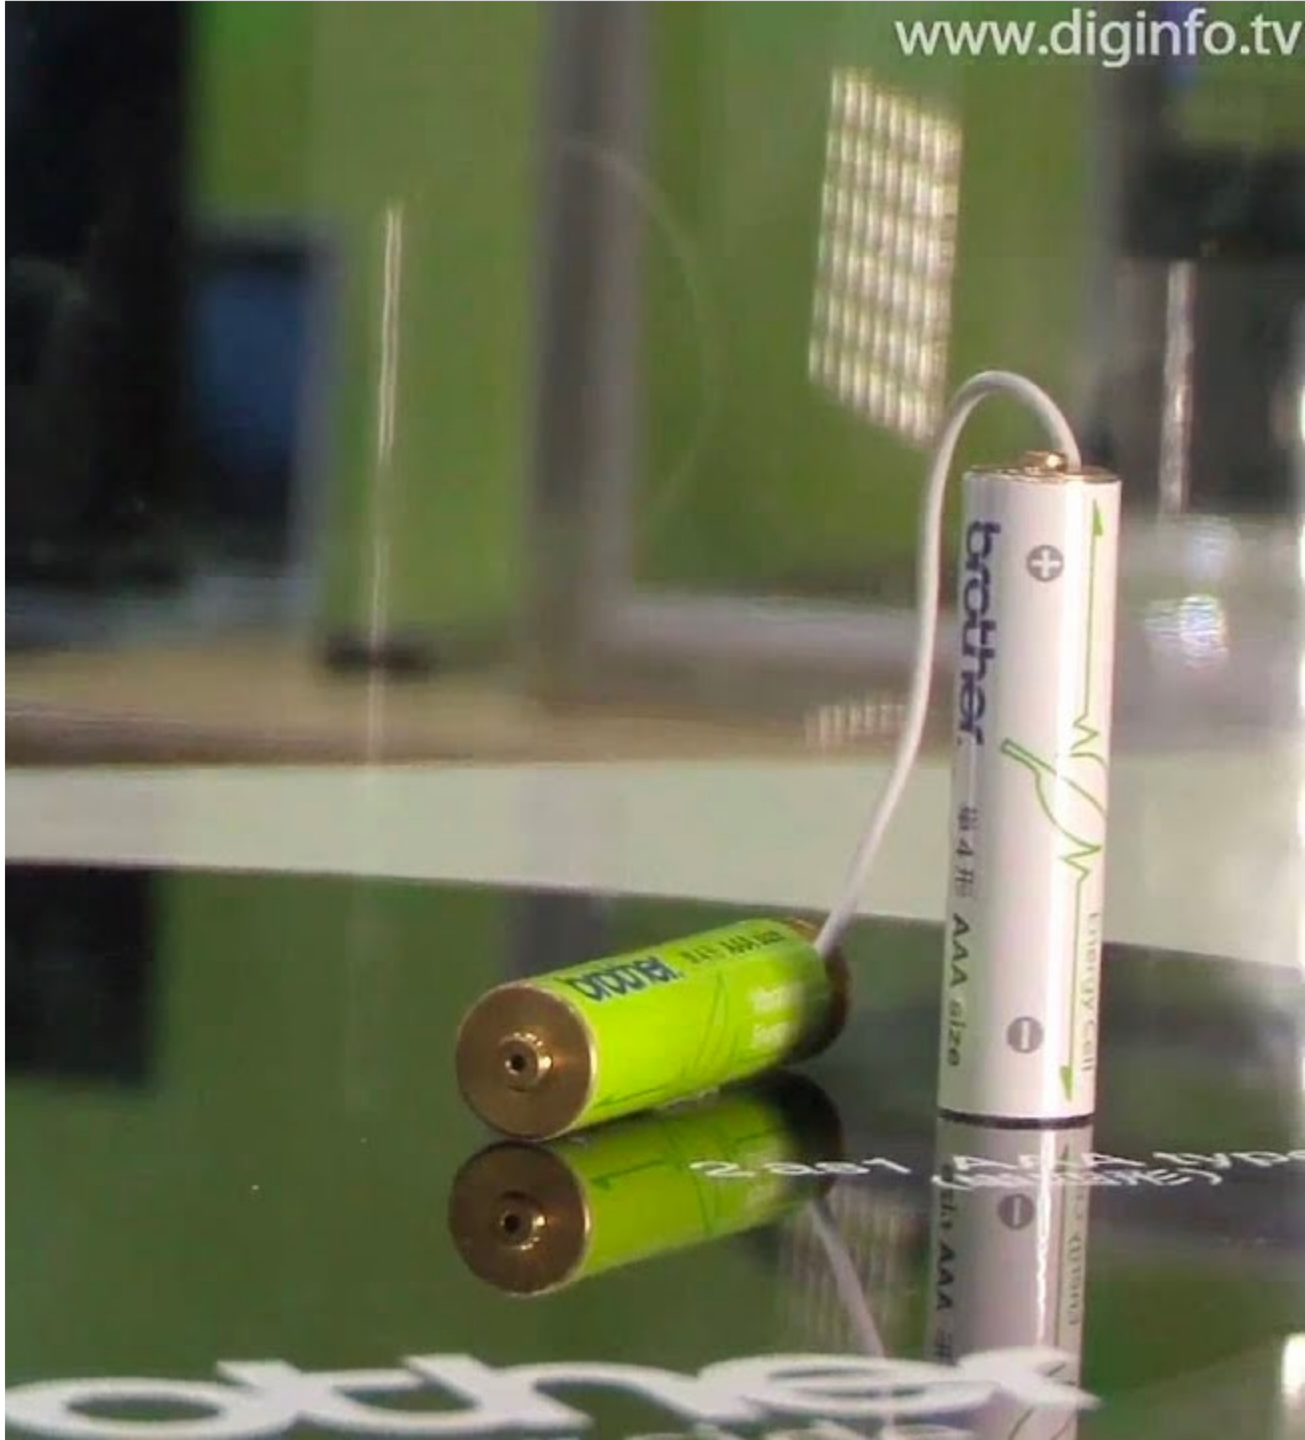

YOUR QUESTIONS

Professor Geoff Merrett | gvm@ecs.soton.ac.uk
Head of Centre for IoT and Pervasive Systems