



# Application-aware Energy Attack Mitigation in the Battery-less Internet of Things

Chetna Singhal

chetna.iitd@gmail.com

Indian Institute of Technology (IIT)  
Kharagpur, India

Thiemo Voigt

thiemo.voigt@angstrom.uu.se

Uppsala University and RI.SE Sweden

Luca Mottola

luca.mottola@ri.se

RI.SE Sweden, Uppsala University,  
and Politecnico di Milano, Italy

## ABSTRACT

We study how to mitigate the effects of energy attacks in the battery-less Internet of Things (IoT). Battery-less IoT devices *live and die* with ambient energy, as they use energy harvesting to power their operation. They are employed in a multitude of applications, including safety-critical ones such as biomedical implants. Due to scarce energy intakes and limited energy buffers, their executions become *intermittent*, alternating periods of active operation with periods of recharging their energy buffers. Experimental evidence exists that shows how controlling ambient energy allows an attacker to steer a device execution in unintended ways: energy provisioning effectively becomes *an attack vector*. We design, implement, and evaluate a mitigation system for energy attacks. By taking into account the specific application requirements and the output of an attack detection module, we tune task execution rates and optimize energy management. This ensures continued application execution in the event of an energy attack. When a device is under attack, our solution ensures the execution of 23.3% additional application cycles compared to the baselines we consider and increases task schedulability by at least 21%, while enabling a 34% higher peripheral availability.

## CCS CONCEPTS

• **Security and privacy** → **Embedded systems security**; *Denial-of-service attacks*; • **Computer systems organization** → **Sensor networks**.

## KEYWORDS

Intermittent computing, Energy-attack mitigation, Task scheduling, Federated energy harvesting, Battery-less IoT Application

## ACM Reference Format:

Chetna Singhal, Thiemo Voigt, and Luca Mottola. 2023. Application-aware Energy Attack Mitigation in the Battery-less Internet of Things. In *Proceedings of the Int'l ACM Symposium on Mobility Management and Wireless Access (MobiWac '23)*, October 30–November 3 2023, Montreal, QC, Canada. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3616390.3618281>


Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

*MobiWac '23*, October 30–November 3 2023, Montreal, QC, Canada

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0367-6/23/10...\$15.00

<https://doi.org/10.1145/3616390.3618281>



**Figure 1: Example intermittent execution.** Periods of active computation and periods of charging the energy buffer alternate.

## 1 INTRODUCTION

Ambient energy harvesting allows Internet of Things (IoT) devices to eliminate their dependency on traditional batteries [10]. This enables drastic reductions of maintenance costs and previously unattainable deployments, for example, due to the reduction of a device's footprint. Several battery-less IoT deployments exist, even in safety-critical settings such as biomedical implants [1, 14, 15, 19, 26, 52] and using a variety of energy sources [10].

Energy harvested from the environment, however, is highly variable in time [10], yet energy buffers, such as capacitors, need to be miniaturized as well, as they often represent a dominating factor in size. This trait clashes with the push to realize tiny devices enabling pervasive deployments. System shutdowns due to energy depletion are thus difficult to avoid and executions become *intermittent* [25]: periods of active execution and periods of energy harvesting come to be unpredictably interleaved.

**Computing intermittently.** Fig. 1 shows an example execution. The ambient charges the onboard capacitor until a voltage threshold  $V_{on}$  is reached that causes the device to power on. The device starts sensing, computing, and communicating as long as the capacitor charge remains above a threshold  $V_{off}$ .

Using existing energy harvesting techniques and resource-constrained IoT platforms, such as TI MSP430 MCUs, the net energy balance during active execution is systematically negative, so the device rapidly reaches  $V_{off}$ . When this happens, the device switches off, waiting for the capacitor to reach again the power-on threshold  $V_{on}$ . This intermittent execution may occur on tiny time scales; computing simple error correction codes on a battery-less IoT device powered through radio-frequency energy harvesting may require as many as 16 energy cycles [11].

Using resource-constrained IoT devices, applications run on bare hardware with *no operating system support* [25]. When the device powers off upon reaching  $V_{off}$ , the system state would normally be lost and the device would completely reboot the next time it reaches  $V_{on}$ . To address this issue, intermittently-computing IoT systems employ ad-hoc techniques [2, 8, 9, 11, 16, 31, 39–42, 44, 50, 51, 59, 60] to create and maintain *persistent state* on non-volatile memory (NVM). These systems operate as the device approaches

$V_{off}$ , allowing systems to retain the application state across energy failures. NVM operations, however, are extremely energy hungry and impose a significant energy overhead [43].

Experimental evidence exists that shows how intermittently-computing IoT devices are vulnerable to *energy attacks* [45], that is, energy harvesting can be used as an *attack vector*. Exerting simple control on how the ambient provisions energy allows an adversary to create situations of *livelock*, *priority inversion*, and *denial of service*, *without* requiring physical access to the devices. The simplest scenario may consist, for example, in the attacker physically blocking the radiation arriving at a solar panel that powers the IoT device, eventually impeding forward progress, yet existing literature demonstrates more subtle setups [45].

**Contribution and road-map.** We tackle the problem of *mitigating* the impact of energy attacks on system performance. One option is to design countermeasures that apply to specific attacks. For example, when an energy attack is meant to cause a livelock, one may keep track of the number of restarts from a specific point in the code and suspend the operation for a given time period if this quantity exceeds a threshold.

We take a different stand here. We develop a generic abstraction and support run-time, called Energy Attack Mitigation (EAM), that applies independent of the specific energy attack. We consider a task-based multi-capacitor architecture, further described in Sec. 2, given its use in battery-less IoT deployments [1, 17, 19]. Using EAM, as illustrated in Sec. 3, programmers specify different *application profiles* that map task execution rates to different system states. When an on-device attack detection system reports the occurrence of an energy attack [45], EAM intelligently throttles task execution rates and accordingly orchestrates the charging of different capacitors, reducing energy consumption in an attempt to retain a minimum programmer-specified. It resumes normal operation once the attack is (expected to be) over. EAM consists of the energy manager, task scheduler/peripheral control, and application manager functional blocks that are implemented in software and run on-board the IoT device MCU.

Our evaluation, reported in Sec. 4, uses a mixture of numerical simulations and emulation experiments to compare EAM against two state-of-the-art baselines, using *real-world* energy traces. We evaluate the run-time energy consumption overhead associated with the on-board execution of the EAM solution. A single execution of EAM has negligible overhead, requiring less than 2 nJ energy and taking less than 1.5  $\mu$ sec. The results indicate that EAM ensures the execution of 23.3% additional application cycles compared to the baselines we consider and increases task schedulability by at least 21%, while enabling a 34% higher component availability.

## 2 BACKGROUND AND RELATED WORK

The security literature for *battery-powered IoT* systems is extensive. Resource-constraints generally make it difficult to apply standard security mechanisms [57]; further, battery-powered IoT devices enables peculiar attacks, for example, in an attempt to drain batteries [35, 48]. To save energy, they feature low-power radios, which makes them vulnerable to denial of service attacks, for example, due to intentional jamming [32]. Specialized network stacks are also necessary for multi-hop operation, which are vulnerable to new kinds of attacks. System characteristics motivate new techniques

for attack detection and mitigation ranging from hardware-based solutions [49] to techniques employing machine learning [18, 56].

Solutions for battery-powered IoT systems, unfortunately, falls short of expectations in the case of *battery-less* IoT devices. Energy constraints are way more severe compared to their battery-powered counterpart. Limited form factors impose restrictions on the harvesting unit, limiting power supply to tens of *mW* [14, 21, 30, 52]. This creates a demand-supply gap. Systems use tiny energy buffers, such as capacitors, to tame fluctuations of energy intake and for performing operations whose power consumption exceeds the maximum harvesting capabilities. Moreover, the intermittent execution pattern adds a new dimension to the problem [25], as periods of active operation are interspersed with periods for recharging energy buffers, while the rest of the system is quiescent.


**Architectures.** Fig. 3 depicts the prevailing architecture for intermittently-computing IoT devices, as seen in both available platforms [24, 30] and concrete deployments [14, 21]. A mixed-volatile MCU [28] relies on multiple capacitors [17, 23] to tame fluctuations of energy intake and for performing operations whose power consumption exceeds the maximum harvesting capabilities. The multiple capacitors allow the system to strike a better trade-off between charging times and available energy. Smaller capacitors are the first to reach  $V_{on}$ ; as this happens, tasks that consume little energy, such as probing low-power sensors, are immediately executed. Bigger capacitors take longer to reach  $V_{on}$ ; their energy is eventually consumed by energy-hungry tasks, such as radio operations. A dedicated energy management module governs the charging of the multiple capacitors based on energy intake and task demands.

Intermittently-computing IoT devices employ techniques such as checkpointing [2, 8, 9, 11, 31, 41, 42, 50, 59] or task-based programming abstractions [16, 39, 44, 51, 54, 60] to deal with energy failures. The former consist in replicating the application state on NVM, where it is retrieved back once the system resumes with sufficient energy. The latter offer abstractions to define and manage persistent state, while taking care of data consistency in case of repeated executions of non-idempotent code [59].

We adopt a task-based programming model and a multi-capacitor architecture as their combination is demonstrated to be effective in real-world deployments [1, 17, 19]. Crucially, they also offer ample opportunities to intelligently manage energy when the system is under attack, as illustrated in Sec. 3.

**Resource management.** Orthogonal to the use of multi-capacitor architectures, several techniques exist to manage energy in battery-less IoT devices. Dynamic voltage scaling (DVS) is one such technique that maintains the required performance level by adjusting the system supply voltage to a suitable minimum [61]. Dynamic frequency scaling (DFS), differently, meets the required performance by modulating the system clock frequency [13]. Dynamic voltage and frequency scaling (DVFS) combines the two [4, 6]; custom designs of DVFS exist especially for intermittently-computing devices.

Software-based maximum power point (MPP) tracking can achieve further energy gains. It can adapt the power consumption of the system that is operating at an efficient operating voltage and maximizing forward application execution without adding any external tracking or control units [7, 20]. Custom MPP designs also exist for specific energy sources. In the case of kinetic energy [55], for



**Figure 2: Tasks in example applications: (a) HVAC system, (b) greenhouse irrigation system, and (c) controlled room ventilation.** Tasks represent individual units of processing and run with transactional semantics. The exchange data through non-volatile data queues subject to timing requirements.

example, voltage-current characteristics of the energy transducer are used to find the optimal operating point, dynamically [53].

Energy scarcity also makes the real-time task scheduling difficult [29, 33, 52]. Schedulability may be improved by dynamically scheduling the computational and energy-harvesting tasks [29]. This essentially means co-designing scheduling and energy management, which is demonstrate to be particularly effective in multi-capacitor architectures [23, 29]. This is precisely what we do in this work, yet with a different goal compared to existing literature. Instead of generally improving schedulability, we aim to reach a minimum programmer-specified task execution rate against widely varying energy intakes, possibly affected by energy attacks.

**Security.** Compared to battery-powered IoT devices, the security issues in battery-less IoT devices drastically change, largely making the area uncharted territory.

The few existing solutions focus on securing persistent state. Krishnan et al. [36] present an attack model for unsecured and cryptographically secured checkpoints. Asad et al. [5] present an experimental evaluation on the use of different encryption algorithms and ARM TrustZone. Krishnan et al. [38] build on this and propose a configurable checkpoint security setting that leverages application properties to reduce overhead. Ghodsi et al. [22] use lightweight algorithms [12] for securing checkpoints. Valea et al. [58] propose a SECure Context Saving (SECCS) hardware module inside the MCU, whereas Khrihnan et al. [37] apply Authenticated Encryption with Associated Data (AEAD) to protect checkpointing data.

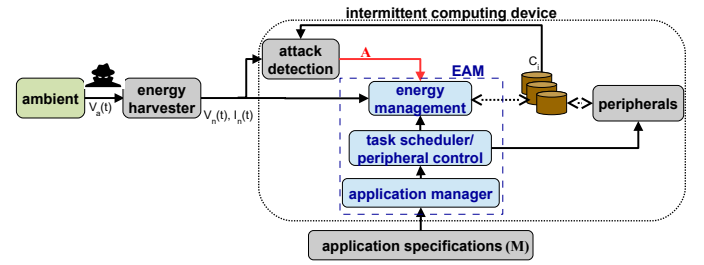
Experimental evidence of energy attacks and corresponding vulnerabilities exists [45], where machine learning is applied to develop a generic attack detection module. This work tackles the complementary problem, that is, to mitigate the effects of an energy attack on system performance once the attack is detected.

### 3 DESIGN

We use three paradigmatic applications as running examples, described in Sec. 3.1, along with the programming model we use to encode their logic. Sec. 3.2 describes the architecture of an EAM-enabled system and its run-time operation.

#### 3.1 Applications and Programming Model

We use a programming model that naturally maps to embedded sensing applications [34] and is based on well-established IoT programming concepts [46, 47]. Using this model, we consider three



**Figure 3: Hardware/software device architecture when using EAM.** EAM takes as input the output of the attack detection module, the  $V_n(t)$ ,  $I_n(t)$  signals from the energy harvester, the current capacitor charge, and the application specification  $M$ .

example applications, graphically illustrated in Fig. 2, corresponding to concrete deployments of battery-less IoT systems [1, 14, 24].

The application logic is split in individual tasks, shown as boxes in Fig. 2, which represent independent units of processing that execute asynchronously with respect to each other. Tasks have inputs and outputs, represented in Fig. 2 with links between the tasks, with attached timing requirements. For example, task *temperature sensing* in Fig. 2(a) must provide a new temperature reading to the *decision* task at least once every five minutes. Differently, task *humidity sensing* in Fig. 2(b) must provide a new humidity reading to the *decision* task at least once every 30 minutes, but no more frequently than once every five minutes [23–25].

Tasks exchange data through non-volatile data queues, which offer a way to persist state in case of an energy failure. Tasks run with transactional semantics; if they are interrupted by an energy failure before producing an output, no intermediate result is persisted and the task restarts from the beginning [2, 16, 39, 40, 60]. When a data item flows through the entire task pipeline from start to end, we say the application has *completed an execution*. The application execution rate, that is, the number of times an application completes an execution in the unit of time, is the primary performance metric we discuss in Sec. 4.

#### 3.2 Architecture and Run-time

Parameter	Value
Supported profiles	Short attack (SA), long attack (LA), normal (NML), low power (LP), critical (CTL)
Task set	{Tasks}
Task rate	Profile: {Tasks (Rate)}
Task dependency	{Finish-to-start Task Lists}

**Figure 4: Application specification M.** It includes information on the profiles supported by an application, the task set, their rate requirements, and their dependencies.

Parameter	Symbol
Attack ongoing	$a_o$
Attack ongoing accuracy	$a_{oa}$
Elapsed time	$a_{et}$
Remaining time	$a_{rt}$

**Figure 5: Energy attack parameters in information set A.** The information returned by the attack detection module as used by EAM to adapt the application behavior during an attack.

Symbol	Definition
$\rho$	Application profile
$\alpha$	Maximum $a_{rt}$ corresponding to SA
$E, E_i$	Available energy and energy stored in buffer $i$
$\Omega_0, \Omega_1$	Maximum $E$ for CTL and LP profiles, resp.
$\tau, \mathcal{T}, \tau_e$	Task, task set, and current execution task
$\mathbf{R}(\tau), \mathbf{S}(\tau)$	Task-rate and Task-state of $\tau$ , resp.
$\epsilon(\tau), \mathcal{B}(\tau)$	Energy required and energy-buffer used to execute $\tau$
$\lambda_\tau$	Energy proportion stored in $\mathcal{B}(\tau)$ for blocked, suspended $\tau$
$\Lambda_\tau, \lambda_\tau > \lambda_\tau$	Energy proportion stored in $\mathcal{B}(\tau)$ for ready, running $\tau$

**Figure 6: Notation.** The EAM parameter symbols and corresponding definitions are summarized here.

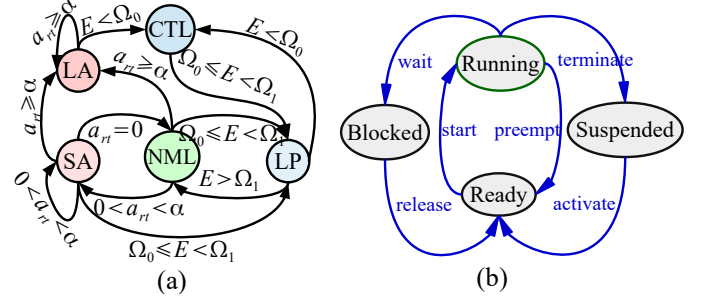
Fig. 3 illustrates the hardware/software architecture of an EAM-enabled system. EAM takes as input a specification  $\mathbf{M}$  of the application behavior, illustrated in Fig. 4. This specification includes information on what *system profile* is supported by the application among five possible options we discuss next. It also specifies the application's task set, along with their minimum and maximum required execution rates for each supported profile and their dependencies. The latter is essentially an encoding of the information graphically represented in Fig. 2.

The five available system profiles allow programmers to tune an application's behavior depending on the situation. Profile *normal* (NML) indicates regular operation. Profile *low-power* (LP) represents a general situation of energy scarcity, yet the application may still continue the execution, whereas profile *critical* (CTL) represents a situation where the system may fail to achieve progress due to extreme energy scarcity. None of these situations, however, necessarily indicates that an adversary is manipulating the energy supplies. These situations are represented by profile *short attack* (SA) and *long attack* (LA), depending on the expected attack duration. We indicate with  $\rho$  what system profile is currently active.

The adversary manipulates the energy signal  $V_a(t)$  at time  $t$  and affects the voltage and current  $V_n(t), I_n(t)$  supplied to the intermittently-computing IoT device. The attack detection module, shown in Fig. 3 may realize an attack is occurring and alert EAM by passing an information set  $\mathbf{A}$  about the attack, along with the manipulated  $V_n(t), I_n(t)$  signals. The information we expect to find in  $\mathbf{A}$  is in Fig. 5. Quantity  $a_o$  is a binary flag indicating whether an attack is ongoing, whereas quantity  $a_{oa}$  indicates how reliable is the information provided by the attack detector. Quantity  $a_{et}$  and  $a_{rt}$  respectively indicate when the attack detector thinks the attack started, and for how long it is expected to continue; these quantities are useful to determine whether profile SA or LA better represent the current situation. Existing attack detectors can provide this information accurately [45]. We consider that the harvester receives no energy during the the long and short energy attack, that is, during an ongoing energy attack ( $a_o = 1$ ) at time  $t'$ ,  $V_n(t') = 0$ .

In Fig. 3, boxes in light blue are the modules we add to a regular configuration. We describe their functioning next.

**Application manager.** The *application manager* module governs the application execution depending on the active system profile, based on the application specification  $\mathbf{M}$ . It does so by periodically probing the energy buffer to estimate the available energy  $E$  and by monitoring the information set  $\mathbf{A}$  about possible attacks.



**Figure 7: Transition diagrams for: (a) system profile and (b) task state.** System profiles model configurations that EAM enforces depending on the situation, for example, in case of an energy attack.

Fig. 7(a) shows the transition diagram for system profile. The application manager selects the system profile based on available energy  $E$  and the energy attack parameter  $a_{rt}$ . When an attack is ongoing, that is, as long as  $a_o = 1$ , if the remaining time of the attack is greater than a configuration parameter  $\alpha$ , that is,  $a_{rt} \geq \alpha$ , then the LA profile is selected; or else the SA is selected. The energy level thresholds  $\Omega_0$  and  $\Omega_1$ ,  $\Omega_1 > \Omega_0$ , are used to select between the NML, LP, and CTL profiles. When the attack is not ongoing, that is, as long as  $a_o = 0$ , if  $E > \Omega_1$  then the application manager selects the NML profile; or else if  $E < \Omega_0$  then the CTL profile is selected. If none of these applies, the LP profile is selected. The system state, including energy level and energy attack information, is periodically updated and profile selection happens based on the system current state.

**Energy manager.** We consider a multi-capacitor energy storage architecture [23], which allows the system to strike an efficient trade-off between charging times and energy efficiency. The lightweight tasks consume energy from the smaller capacitors that charge more quickly, whereas more energy-intensive tasks are executed using the energy stored in the larger capacitors [17].

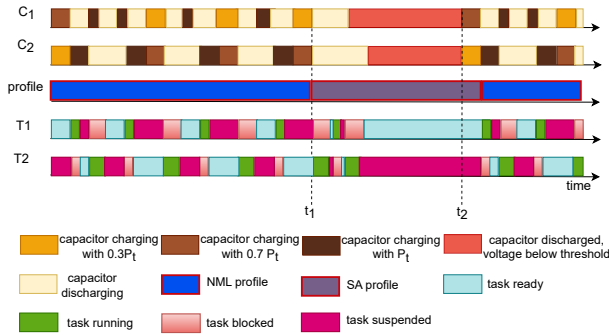
We consider a parallel resistor-capacitor circuit with the equivalent storage capacitor,  $C_i$  in parallel to the resistor equivalent to the rest of the circuit,  $R_p$ . Hence the voltage of the capacitor,  $C_i$  at time instance  $t$ , during charging, is given as:

$$V_i(t) = \sqrt{P(t) \cdot R_p - e^{-\frac{2t}{C_i R_p}} \cdot (P \cdot R_p - V_0^2)} \quad (1)$$

where  $V_0$  is the capacitor voltage at  $t = 0$  and  $P(t)$  is the power from the energy source at time  $t$ . The energy buffer,  $\mathcal{B}(\tau) \in \mathbb{N}$ ,  $\mathcal{B}(\tau) \leq m$  is the capacitor that supports execution of task  $\tau$ .

The *energy manager* module in Fig. 3 directs the incoming energy towards different capacitors depending on the task energy requirements in the different system profiles. It stores energy proportionally in the capacitors that sustain the operation of tasks that are ready to execute, that is, they have data in their input queues, and are closer to the deadline indicated in the application specification. This means that the energy manager dynamically changes the amount of energy going towards different capacitors depending on the instantaneous system needs. This is achieved by periodically checking the state of all tasks, the corresponding timing requirements given the active system profile, and the current energy content of all capacitors.





**Figure 8: Example EAM execution.** A short energy attack is detected at  $t_1$ . The system transitions to profile SA where only  $T_1$  is scheduled for execution as all incoming energy is directed to  $C_1$ .  $T_2$  is suspended in the mean time and its capacitor does not charge until the attack is over and the system transitions back to NLM.

**Task scheduler and peripheral control.** We co-design task scheduling with energy management. Tasks may be in one of the four states of Fig. 7(b). Tasks are ready when there is data in their incoming queues. One of them is selected for execution using a form of real-time periodic task scheduling (RTS) [33], subject to sufficient energy availability in the corresponding capacitor. This means that a task that is ready to execute, but has no sufficient energy to complete, is deferred until sufficient energy is available.

Tasks are set to be in a blocked state if they depend on another task that is yet to be executed, that is, they consume its output. Suspended tasks are those waiting for low-level I/O operations to complete, for example, obtaining a sensor reading, or are waiting for sufficient energy to start the execution.

**Example.** Fig. 8 shows a timing diagram for an example application that has two tasks,  $T_1$  (light weight) and  $T_2$  (energy-intensive), that are executed using the energy stored in capacitors  $C_1$  and  $C_2$ , respectively. Here,  $C_1 < C_2$ , thus  $C_1$  has shorter charging and discharging times as compared to  $C_2$ . Task  $T_2$  depends on the input from  $T_1$  and can only execute after completion of  $T_1$ . The capacitors are charged proportionally with input power based on the state of the tasks that each supports. A higher proportion of power is used to charge the capacitor that must support a ready task.

Initially, the profile is set to NML. A short attack happens from  $t_1$  to  $t_2$  and the energy available for charging the capacitors greatly reduces during this time. When the energy attack starts at  $t_1$ , the profile is set to SA. During the energy attack, the tasks are executed as many times as possible while maintaining dependency using the remaining energy in the capacitors until they both fall below the threshold. When the energy attack ends at time  $t_2$ , the capacitors start charging proportionally based on the task state. As soon as the voltage in  $C_1$  reaches  $V_{on}$ , EAM is executed and the profile is set back to NML and the periodic execution of the tasks resumes.

**Energy attack mitigation algorithm.** Algorithm 1 shows EAM’s operation through pseudocode. The inputs include the application specification  $\mathbf{M}$  and energy attack parameters  $\mathbf{A}$ . The outputs include the system profile  $\rho$ , task states  $\mathbf{S}$ , and energy stored in capacitors  $E_i, 1 \leq i \leq m$ . EAM includes three core functionality, operating as follows.

- Function AppTaskManager implements the *application manager*. It selects the system profile  $\rho$  based on the available energy and energy attack parameters  $a_o$  and  $a_{rt}$ . Thereafter, it creates the task set  $\mathbf{T}$  and assigns the task rate  $\mathbf{R}(\tau), \forall \tau \in \mathbf{T}$  using the profile-specific application specification ( $\mathbf{M}$ ).
- Function TaskScheduler implements *task scheduling and peripheral control*. It sets the task state  $\mathbf{S}(\tau)$  as ready if there is sufficient energy in the buffer  $\mathcal{B}(\tau)$  to support the execution of task  $\tau$ . In the event of an energy attack, a task state is set to ready if the remaining time of the attack is estimated to be greater than the task execution time, that is,  $a_{rt} > 1/\mathbf{R}(\tau)$  while the capacitor supporting its execution has sufficient energy. Otherwise the task state is set to blocked. A ready task is selected to be the current execution task  $\tau_e$  when its supporting capacitor has sufficient stored energy.
- Function FederatedHarvesting implements the *energy manager*. It proportionally stores the harvested energy in separate capacitors based on the state of the task it supports. If the task  $\tau$  is ready then a higher proportion of energy  $\Lambda_\tau$  is stored in the corresponding buffer, else a lower proportion of energy  $\lambda_\tau$  is stored.

## 4 EVALUATION

We illustrate the experimental setup in Sec. 4.1 and split the discussion of the results in three parts. In Sec. 4.2, we investigate the improvement enabled by EAM in the *application execution rate*, that is, the number of times an application manages to complete the data pipeline. Next, in Sec. 4.3, we study the *task scheduling* behavior to gain a deeper understanding of how EAM enables the improvements in application execution rate, and conclude the discussion with an analysis of the *run-time energy overhead* imposed by EAM.

### 4.1 Setting

We use the application in Fig. 2(a) to experimentally measure the performance of EAM and of the baselines. This is the most challenging case in Fig. 2, especially in terms of task execution rates. We describe next the metrics we use to evaluate EAM performance, along with the baselines we compare with and the tools we use.

**Metrics.** We consider five performance metrics

- The *application execution rate* is the number of times per time unit an application completes the data pipeline, while meeting all task dependencies. This is a function of the individual task execution rates, indicated in the application specification  $\mathbf{M}$  depending on the current system profile, and of the way the system manages available resources.
- The *task schedulability* measures the percentage of time that tasks in the system have sufficient resources to execute: in our specific setting, this means that the capacitor(s) in the system store sufficient energy for the task to execute and produce the corresponding output.
- The *component availability* is the percentage of time that the MCU (for computation) or the attached peripherals (for sensing, wireless transmissions, or actuation) have sufficient resources to execute. This metric essentially represents the hardware counterpart of the task schedulability figure.
- The *time and energy overhead* is the time or energy spent for running the given scheduling technique, which essentially measures how much time or energy is subtracted from application processing to manage system resources.

**Algorithm 1:** EAM: Energy-attack Mitigation

---

```

Input: A, M
AppTaskManager(A,M):
  1) Select profile  $\rho$ :
    if  $a_o = 1$  then
      if  $a_{rt} > \alpha$  then  $\rho = \text{LA}$ ;
      else  $\rho = \text{SA}$ ;
    else
      if  $\omega_o < E \leq \omega_1$  then  $\rho = \text{LP}$ ;
      else if  $E < \omega_0$  then  $\rho = \text{CTL}$ ;
      else  $\rho = \text{NML}$ ;
  2) Select task set:
    Include task  $\tau$  in task set T if  $M(\tau, \rho) > 0$ 
  3) Assign task rate:  $R(\tau) = M(\tau, \rho)$ 
  return  $\rho, T, R$ 

TaskScheduler( $\rho, T, R, A$ ):
  1) Set task state (when there is no energy attack):
  for each task  $\tau \in T$  do
    if  $(E_{\mathcal{B}(\tau)} \geq \varepsilon(\tau))$  and  $(\tau = 1 \text{ or } S(\tau - 1) = \text{'running'})$ 
      then set  $S(\tau) = \text{'ready'}$ ;
    else set  $S(\tau) = \text{'blocked'}$ ;
  2) Defer/ execute task in event of energy attack,  $a_o = 1$ :
  for each task  $\tau \in T$  do
    if  $a_{rt} > (1/R(\tau))$  and  $E_{\mathcal{B}(\tau)} > \varepsilon(\tau)$  then
      set  $S(\tau) = \text{'ready'}$ ;
      else set  $S(\tau) = \text{'blocked'}$ ;
  3) Execution task  $\tau_e$  from the task set T:
  for each task  $\tau \in T$  do
    if  $S(\tau) = \text{'ready'}$  and  $E_{\mathcal{B}(\tau)} \geq \varepsilon(\tau)$  then
      Set  $\tau_e = \tau$ 
      Set  $S(\tau) = \text{'running'}$ 
      Remove  $\tau$  from T
  return S,  $\tau_e$ 

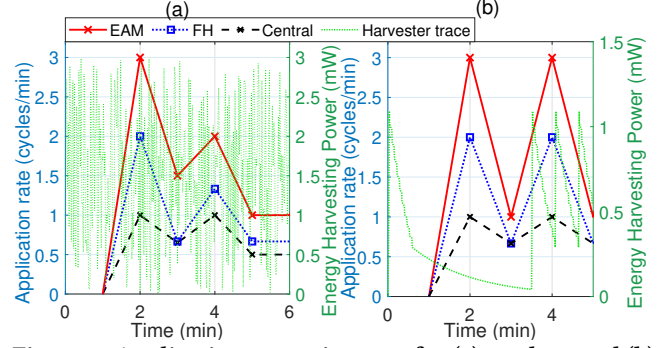
FederatedHarvesting( $\rho, T, R, S$ ):
  1) Proportion of harvested energy to buffers
  for each buffer  $i = 1$  to  $m$  do
    for each  $\tau \in T$  do
      if  $S(\tau) = \text{'ready'}$  or  $\text{'running'}$  then
        Set  $E_{\mathcal{B}(\tau)} = \Lambda_\tau \cdot E$ ;
        else Set  $E_{\mathcal{B}(\tau)} = \lambda_\tau \cdot E$ ;
  return  $E_i(n)$ 
Output:  $\rho, S, E_i, 1 \leq i \leq m$ 

```

---

The application execution rate represents the primary metric, because it measures the level of service provided to end users. Ideally, the scheduler should *at least* meet the application specification  $M$ , while any other improvement is welcome as it indicates the system operates above the minimum requirement. The other metrics, but the run-time energy overhead, are instrumental to understand the trends we observe in application execution rate. For example, a higher execution rate should be enabled by higher task schedulability and/or component availability. The run-time energy overhead is the cost we pay for managing system resources in a given way. Measuring this quantity is necessary to weigh the possible improvements in application execution rate against the needed energy cost.

**Baselines and tools.** We compare EAM with federated energy harvesting (FH) [24] and a centralized architecture (Central) [23].



**Figure 9:** Application execution rate for (a) outdoor and (b) indoor traces using EAM, FH, or Central.

Both use the same real-time energy-aware periodic task scheduling (RTS) [33, 52] as EAM. FH is a multi-capacitor architecture that statically directs incoming energy to the multiple capacitors, in a way proportional to capacitor size. This means it does not take into account possibly varying task execution rates. FH uses the same number and type of capacitors we use in EAM. Comparing with FH allows us to measure the gains enabled by making the scheduler aware of changing application requirements, based on the currently active system profile. The Central baseline, instead, allows us to study the benefits brought by combining a multi-capacitor architecture with application-aware energy management and scheduling.

To obtain quantitative results, we use a combination of numerical simulations and emulation. We use Matlab to evaluate the performance of EAM and the baselines in application execution rate, task schedulability, component availability, and component availability latency. We implement accurate models of energy harvesting and energy consumption, based on the operation of a staple MSP430-based intermittent computing platform [24]. We compute the runtime energy overhead using MSPSim, which provides time-accurate emulation of binary code for the MSP430 platform, considering a 3V capacitor supply as in existing literature [3].

We feed the numerical models with *real-world* energy traces. The traces are time series obtained from the actual output voltage across a 30k $\Omega$  resistor in a solar harvesting system using a monocrystalline high-efficiency solar panel that is partly stationary and partly moving, and operates indoor or outdoors [3]. To investigate the performance during an attack, we also emulate different types of energy attacks, either short- or long-term, by completely zeroing the input voltage trace for a variable time period [45].

## 4.2 Application Execution Rate

Fig. 9 shows an example trace across a few minutes of execution demonstrating how the application execution rate enabled by EAM is systematically higher than the baselines we consider. This observation applies to both traces. The plot also demonstrates how EAM can capture energy fluctuations more effectively than the baselines. For example, at 2 and 4 minutes in Fig. 9(a), the system ends up with an excess of energy and the application execution rate enabled by all solutions accordingly increases. However, the relative improvement is much higher for EAM than for the baselines.

Fig. 10 provides a closer look at the system behavior in case of the scenarios that correspond to the supported application profiles




Figure 10: Application execution rate and active system profile.

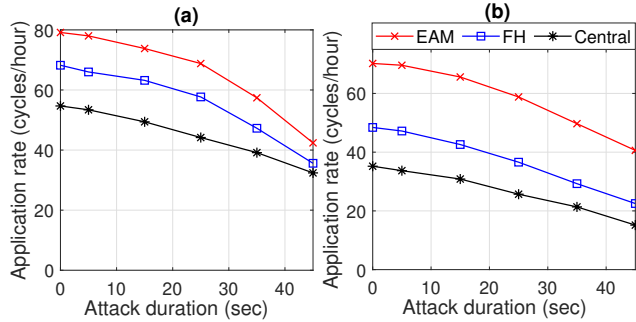


Figure 11: Application execution rate in case of an energy attack for (a) outdoor and (b) indoor traces.

including a short and a long energy attack. The application execution rate enabled by EAM is higher than the baselines in the absence of an energy attack. At time 11175 sec (186.25 min) a short attack begins and EAM maintains a higher application execution rate, resuming normal operation right after the attack is over and maintaining better performance than the baselines throughout. The same observation applies at time 11268 sec (187.81 min) when a long attack occurs.

Most importantly, Fig. 11 plots the application execution rate when we emulate an energy attack of variable duration at a random point in the energy trace, which is the same for EAM and the baselines. As expected, the performance degrades with longer energy attacks, in that the energy budget available for system operation accordingly reduces. However, EAM shows better and more robust performance: the gains over the baselines are retained regardless of the duration of the energy attack. As before, this observation applies to both traces, demonstrating the generality of the conclusion.

As Fig. 9 and Fig. 10 demonstrate that EAM performs better than the baselines already in the absence of an attack, one may conjecture that the better resilience to energy attacks shown in Fig. 11 is not




Figure 12: Application execution rate when an attack commences and all solutions have the same energy budget for (a) outdoor and (b) indoor traces.

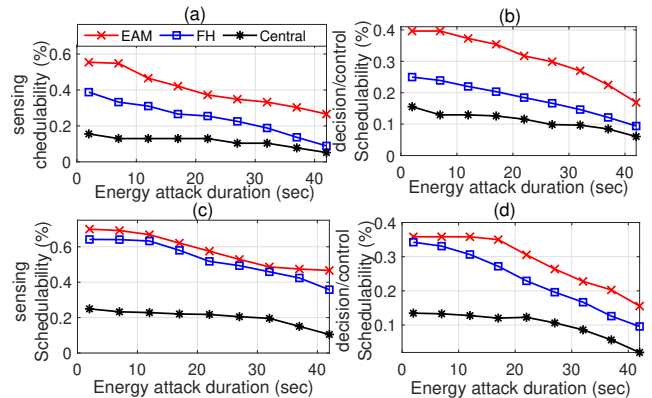


Figure 13: Task schedulability depending with variable energy attack duration for (a-b) outdoor and (c-d) indoor traces.

due to EAM's design, but comes from an energy excess gained beforehand that is eventually spent during the attack. To investigate




Figure 14: Component availability using outdoor trace.




Figure 15: Component availability using indoor trace.

this specific aspect, we run experiments where both EAM and the baselines are given the same energy budget right before the attack commences, hence nullifying any energy excess that EAM may accumulate beforehand. Fig. 12 shows the results. Even when EAM is given the same initial energy budget than the baselines and an attack commences, EAM enables up to twice the application execution rate compared to the baselines.

### 4.3 Scheduling

The better performance in application execution rate comes from better resource management, which leads to more efficient task scheduling. Fig. 13 depicts the results we obtain in task schedulability, as a function of the type of task being scheduled, the attack duration, and the energy trace. Regardless of the setting, EAM constantly ensures higher schedulability for both types of tasks and for both traces. The result is thus arguably general.

Fig. 14 and Fig. 15 report the component availability of EAM and the baselines in the first few minutes of a sample execution and

Task	Duration	Energy
Sensing	12.030 msec	19.066 $\mu$ J
Decision	10.182 msec	15.731 $\mu$ J
Control	60.150 msec	92.931 $\mu$ J

Figure 16: Energy cost of application tasks.

overall, for either energy trace. Using outdoor trace, we observe that FH yields a higher availability for the MCU than Central, yet a lower availability for the peripherals necessary for actuation. The situation is different in indoor trace, where FH improves the availability of both compared to Central. Regardless of the Trace, EAM yields higher component availability than any other baseline.

We use MSPSim to ascertain the energy consumption overhead due to EAM, compared to the energy cost of running the actual application logic. A single execution of EAM requires 1.781 nJ and takes 1.237  $\mu$ sec. The energy cost and execution time of a single application execution is reported in Fig. 16 depending on the single task [27, 28, 52]. The execution time and energy consumption of EAM are orders of magnitude smaller than any application task. The overhead is thus negligible.

## 5 CONCLUSION

We presented EAM, an application-aware energy attack mitigation system. EAM tunes the task execution rates and steers the charging of multiple capacitors to meet programmer-provided application requirements, specified as a function of system profiles. We illustrated the co-design of scheduling and energy management in EAM and quantified its performance against two baselines. We experimentally demonstrated that EAM produces higher application execution rates before, during, and after energy attacks, compared to the baselines we consider. It ensures the execution of 23.3% additional application cycles compared to the baselines we consider and increases task schedulability by at least 21%, while enabling a 34% higher component availability.

## ACKNOWLEDGEMENT

This work has been financially supported by the Swedish Foundation for Strategic Research.

## REFERENCES

- [1] M. Afanasov et al. 2020. Battery-Less Zero-Maintenance Embedded Sensing at the Mithræum of Circus Maximus. In *Proc. ACM Conf. Embedded Networked Sensor Sys. (SENSYS)*.
- [2] S. Ahmed et al. 2019. Efficient Intermittent Computing with Differential Checkpointing. In *Proc. ACM Intern. Conf. Languages, Compilers, and Tools for Embedded Sys. (LCTES)*.
- [3] S. Ahmed, A. Bakar, N. A. Bhatti, M. H. Alizai, J. H. Siddiqui, and L. Mottola. 2019. The Betrayal of Constant Power  $\times$  Time: Finding the Missing Joules of Transiently-powered Computers. In *Proc. Intern. Conf. Languages, Compilers, and Tools for Embedded Sys. (LCTES)*.
- [4] S. Ahmed, Ul A. Q., J. H. Siddiqui, L. Mottola, and M. H. Alizai. 2020. Intermittent Computing with Dynamic Voltage and Frequency Scaling. In *Proc. Intern. Conf. Embedded Wireless Sys. and Networks (EWSN)*.
- [5] H. A. Asad et al. 2020. On securing persistent state in intermittent computing. In *Intern. Workshop on Energy Harvesting and Energy-Neutral Sensing Sys.*
- [6] 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. Computer-Aided Design of Integrated Circuits and Sys.* 35, 5 (2016), 738–749.
- [7] D. Balsamo, B. J. Fletcher, A. S. Weddell, G. Karatzias, B. M. Al-Hashimi, and G. V. Merrett. 2019. Momentum: Power-Neutral Performance Scaling with Intrinsic MPPT for Energy Harvesting Computing Systems. *ACM Trans. Embed. Comput. Syst.* 17, 6 (2019).



- [8] 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. on Computer-Aided Design of Integrated Circuits and Sys.* (2016).
- [9] D. Balsamo, A. S. Weddell, G. V. Merrett, B. M. Al-Hashimi, D. Brunelli, and L. Benini. 2015. Hibernus: Sustaining Computation During Intermittent Supply for Energy-Harvesting Systems. *IEEE Embedded Sys. Letters* (2015).
- [10] N. A. Bhatti, M. H. Alizai, A. A. Syed, and L. Mottola. 2016. Energy Harvesting and Wireless Transfer in Sensor Network Applications: Concepts and Experiences. *ACM Trans. on Sensor Networks* (2016).
- [11] N. A. Bhatti and L. Mottola. 2017. HarvOS: Efficient Code Instrumentation for Transiently-powered Embedded Sensing. In *Proc. ACM/IEEE Intern. Conf. Information Processing in Sensor Networks (IPSN)*.
- [12] J. Borghoff et al. 2012. PRINCE—a low-latency block cipher for pervasive computing applications. In *Intern. Conf. the Theory and Application of Cryptology and Information Security*.
- [13] B. Brock and K. Rajamani. 2003. Dynamic power management for embedded systems [SOC design]. In *Proc. IEEE Intern. [Systems-on-Chip]*. 416–419.
- [14] Q. Chen, Y. Liu, G. Liu, Q. Yang, X. Shi, H. Gao, L. Su, and Q. Li. 2017. Harvest Energy from the Water: A Self-Sustained Wireless Water Quality Sensing System. *ACM Trans. on Embedded Computing Sys.* (2017).
- [15] H. Chiang, J. Hong, K. Kinningham, L. Riliskis, P. Levis, and M. Horowitz. 2018. Tethys: Collecting Sensor Data without Infrastructure or Trust. In *Proc. IEEE/ACM Intern. Conf. Internet-of-Things Design and Implementation (IoTDI)*.
- [16] A. Colin and B. Lucia. 2016. Chain: Tasks and Channels for Reliable Intermittent Programs. In *Proc. ACM SIGPLAN Intern. Conf. Object-Oriented Programming, Sys., Languages, and App. (OOPSLA)*.
- [17] A. Colin, E. Ruppel, and B. Lucia. 2018. A Reconfigurable Energy Storage Architecture for Energy-Harvesting Devices. In *Proc. Intern. Conf. Architectural Support for Programming Languages and Operating Sys. (ASPLOS)*.
- [18] K. AP da Costa, J. Papa, C. O Lisboa, R. Munoz, and V. H. C de Albuquerque. 2019. Internet of Things: A survey on machine learning-based intrusion detection approaches. *Computer Networks* 151 (2019).
- [19] B. Denby, K. Chintalapudi, R. Chandra, B. Lucia, and S. Noghiabi. 2023. Kodan: Addressing the Computational Bottleneck in Space. In *Proc. Intern. Conf. Architectural Support for Programming Languages and Operating Sys. (ASPLOS)*.
- [20] B. J. Fletcher, D. Balsamo, and G. V. Merrett. 2017. Power neutral performance scaling for energy harvesting MP-SoCs. In *Design, Automation and Test in Europe Conference (DATE)*. 1516–1521.
- [21] F. Fraternali, B. Balaji, Y. Agarwal, L. Benini, and R. Gupta. 2018. Pible: Battery-Free Mote for Perpetual Indoor BLE Applications. In *Proc. Conf. Sys. for Built Environments*.
- [22] Z. Ghodsi, S. Garg, and R. Karri. 2017. Optimal checkpointing for secure intermittently-powered IoT devices. In *Intern. Conf. Computer-Aided Design (ICCAD)*.
- [23] J. Hester, L. Sitanayah, and J. Sorber. 2015. Tragedy of the Coulombs: Federating Energy Storage for Tiny, Intermittently-Powered Sensors. In *Proc. ACM Conf. Embedded Network Sensor Sys.* 5–16.
- [24] J. Hester and J. Sorber. 2017. Flicker: Rapid Prototyping for the Batteryless Internet-of-Things. In *Proc. Intern. Conf. Embedded Network Sensor Sys. (SENSYS)*.
- [25] J. Hester and J. Sorber. 2017. The Future of Sensing is Batteryless, Intermittent, and Awesome. In *Proc. ACM Conf. Embedded Network Sensor Sys. (SENSYS)*.
- [26] N. Ikeda, R. Shigeta, J. Shiomi, and Y. Kawahara. 2020. Soil-Monitoring Sensor Powered by Temperature Difference Between Air and Shallow Underground Soil. *Proc. ACM on Interactive, Mobile, Wearable and Ubiquitous Tech. (IMWUT)* (2020).
- [27] Texas Instruments. 2008. MSP430F261x, MSP430F261xA Mixed Signal Microcontroller Datasheet. <https://www.ti.com/lit/ds/symlink/msp430f2618.pdf>. Revised March 2022; Accessed on February 27, 2023.
- [28] Texas Instruments. 2017 (last access: May 13th, 2020). MSP430-FR5969 datasheet. <https://www.ti.com/lit/ds/symlink/msp430fr5969.pdf>.
- [29] B. Islam and S. Nirjon. 2020. Scheduling Computational and Energy Harvesting Tasks in Deadline-Aware Intermittent Systems. In *Proc. IEEE Real-Time and Embedded Tech. and App. Symp. (RTAS)*. 95–109.
- [30] N. Jackson, J. Adkins, and P. Dutta. 2019. Capacity over Capacitance for Reliable Energy Harvesting Sensors. In *Proc. Intern. Conf. Information Processing in Sensor Networks (IPSN)*.
- [31] H. Jayakumar, A. Raha, W. S. Lee, and V. Raghunathan. 2015. QuickRecall: A HW/SW Approach for Computing Across Power Cycles in Transiently Powered Computers. *ACM Journal on Emerging Tech. in Computing Sys.* (2015).
- [32] J. Kanwar, N. Finne, N. Tsiftes, J. Eriksson, T. Voigt, Z. He, C. Åhlund, and S. Saguna. 2021. JamSense: Interference and Jamming Classification for Low-power Wireless Networks. In *IFIP Wireless and Mobile Networking Conf. (WMNC)*.
- [33] M. Karimi, H. Choi, Y. Wang, Y. Xiang, and H. Kim. 2021. Real-Time Task Scheduling on Intermittently Powered Batteryless Devices. *IEEE Internet of Things Journal* 8, 17 (Mar. 2021), 13328–13342.
- [34] M. A. Kordestani and H. Bourdoucen. 2017. A survey on embedded open source system software for the internet of things. In *Free and Open Source Software Conf.*
- [35] K. F. Krentz, C. Meinel, and H. Graupner. 2017. Countering Three Denial-of-Sleep Attacks on ContikiMAC. In *Intern. Con. on Embedded Wireless Sys. and Networks (EWSN)*.
- [36] A. Krishnan and P. Schaumont. 2018. Exploiting Security Vulnerabilities in Intermittent Computing. In *SPACE Intern. Conf.*
- [37] A. Krishnan, C. Suslowicz, D. Dinu, and P. Schaumont. 2019. Secure intermittent computing protocol: Protecting state across power loss. In *Design, Automation & Test in Europe Conf. & Exhibition (DATE)*.
- [38] A. S. Krishnan and P. Schaumont. 2022. Benchmarking And Configuring Security Levels In Intermittent Computing. *ACM Trans. Embedded Computing Sys.* (2022).
- [39] B. Lucia and B. Ransford. 2015. A Simpler, Safer Programming and Execution Model for Intermittent Systems. In *Proc. ACM SIGPLAN Conf. Programming Language Design and Implementation (PLDI)*.
- [40] K. Maeng, A. Colin, and B. Lucia. 2017. Alpaca: Intermittent Execution Without Checkpoints. *Proc. ACM Programming Languages* (2017).
- [41] K. Maeng and B. Lucia. 2018. Adaptive dynamic checkpointing for safe efficient intermittent computing. In *Proc. USENIX Symp. on Operating Sys. Design and Implementation (OSDI)*.
- [42] K. Maeng and B. Lucia. 2019. Supporting Peripherals in Intermittent Systems with Just-in-Time Checkpoints. In *Proc. ACM SIGPLAN Conf. Programming Language Design and Implementation (PLDI)*.
- [43] A. Maioli and L. Mottola. 2021. ALFRED: Virtual Memory for Intermittent Computing. In *Proc. ACM Conf. Embedded Networked Sensor Sys. (SENSYS)*.
- [44] A. Y. Majid, C. Delle Donne, K. Maeng, A. Colin, K. S. Yildirim, B. Lucia, and P. Pawelczak. 2020. Dynamic Task-Based Intermittent Execution for Energy-Harvesting Devices. *ACM Trans. on Sensor Networks* (2020).
- [45] L. Mottol, A. Hameed, and T. Voigt. 2023. Energy Attacks in the Battery-less Internet of Things. arXiv:2304.08224 [cs.CR]
- [46] L. Mottola, A. Pathak, A. Bakshi, V. K Prasanna, and G. P. Picco. 2007. Enabling scope-based interactions in sensor network macroprogramming. In *IEEE Intern. Conf. Mobile Adhoc and Sensor Sys.* IEEE.
- [47] L. Mottola, G. P. Picco, F. J. Oppermann, J. Eriksson, N. Finne, H. Fuchs, A. Gaglione, S. Karnouskos, P. Montero, N. Oertel, K. Romer, P. Spies, S. Tranquillini, and T. Voigt. 2019. makeSense: Simplifying the Integration of Wireless Sensor Networks into Business Processes. *IEEE Trans. Software Engineering* 45 (2019).
- [48] V. L. Nguyen, P. C. Lin, and R. H. Hwang. 2019. Energy depletion attacks in low power wireless networks. *IEEE Access* 7 (2019).
- [49] J. Portilla, A. Otero, E. de la Torre, T. Riesgo, O. Stecklina, S. Peter, and P. Langendorfer. 2010. Adaptable security in wireless sensor networks by using reconfigurable ECC hardware coprocessors. *Intern. J. of Distributed Sensor Networks* (2010).
- [50] B. Ransford, J. Sorber, and K. Fu. 2011. Mementos: System Support for Long-running Computation on RFID-scale Devices. *ACM SIGARCH Computer Architecture News* (2011).
- [51] E. Ruppel and B. Lucia. 2019. Transactional Concurrency Control for Intermittent, Energy-harvesting Computing Systems. In *Proc. Intern. Conf. Programming Language Design and Implementation (PLDI)*.
- [52] A. Sabovic, A. K. Sultania, C. Delgado, L. D. Roeck, and J. Famaey. 2022. An Energy-Aware Task Scheduler for Energy-Harvesting Batteryless IoT Devices. *IEEE Internet of Things Journal* 9, 22 (2022), 23097–23114.
- [53] M. M. Sandhu, K. Geissdoerfer, S. Khalifa, R. Jurdak, M. Portmann, and B. Kusy. 2020. Towards Optimal Kinetic Energy Harvesting for the Batteryless IoT. In *Proc. IEEE Intern. Conf. Pervasive Computing and Commun. Workshops*. 1–6.
- [54] C. Singhal. 2023. Sustainable Application Support in Battery-less IoT Sensing Network System. In *Proc. IEEE Intern. Conf. Commun. (ICC) SIGNIS Workshop*.
- [55] L. Spadaro, M. Magno, and L. Benini. 2016. Poster Abstract: KinetiSee - A Perpetual Wearable Camera Acquisition System with a Kinetic Harvester. In *Proc. ACM/IEEE Intern. Conf. Information Processing in Sensor Networks (IPSN)*. 1–2.
- [56] S. M. Tahsien, H. Karimipour, and P. Spachos. 2020. Machine learning based solutions for security of Internet of Things (IoT): A survey. *Journal of Network and Computer Applications* 161 (2020).
- [57] V. A. Thakor, M. A. Razaque, and M. R. A. Khandaker. 2021. Lightweight cryptography algorithms for resource-constrained IoT devices: A review, comparison and research opportunities. *IEEE Access* 9 (2021).
- [58] E. Valea, M. Da Silva, G. Di Natale, M. L. Flottes, S. Dupuis, and B. Rouzeyre. 2018. SI ECCS: SECure context saving for IoT devices. In *Intern. Conf. Design & Tech. of Integrated Sys. In Nanoscale Era (DTIS)*.
- [59] J. Van Der Woude and M. Hicks. 2016. Intermittent Computation Without Hardware Support or Programmer Intervention. In *Proc. USENIX Conf. Operating Sys. Design and Implementation (OSDI)*.
- [60] K. S. Yildirim, A. Y. Majid, D. Patoukas, K. Schaper, P. Pawelczak, and J. Hester. 2018. InK: Reactive Kernel for Tiny Batteryless Sensors. In *Proc. ACM Conf. Embedded Networked Sensor Sys. (SENSYS)*.
- [61] L. Yuan and G. Qu. 2005. Analysis of energy reduction on dynamic voltage scaling-enabled systems. *IEEE Trans. Computer-Aided Design of Integrated Circuits and Sys.* 24, 12 (2005), 1827–1837.