

THE UNIVERSITY OF CHICAGO

SEASONS: SIGNAL AND ENERGY AWARE SENSING ON INTERMITTENT
SYSTEMS

A DISSERTATION SUBMITTED TO
THE FACULTY OF THE DIVISION OF THE PHYSICAL SCIENCE DIVISIONS
IN CANDIDACY FOR THE DEGREE OF
MASTERS IN COMPUTER SCIENCE

DEPARTMENT OF COMPUTER SCIENCE

BY

POUYA MAHDI GHOLAMI

CHICAGO, ILLINOIS

DEC. 14TH, 2022

TABLE OF CONTENTS

ABSTRACT	iii
1 INTRODUCTION	1
2 BACKGROUND & MOTIVATION	4
3 <i>SEASONS</i>	7
3.1 Adaptive Sampling via Time Buffering	7
3.2 Expiration Mitigation via Sample Queue Dynamics	8
4 <i>SEASONS</i> FRAMEWORK	10
4.1 ASA and the Wrapper	10
4.2 Scheduler	12
4.3 Controller	12
5 EXPERIMENTAL RESULTS	14
5.1 Experimental Setup	14
5.1.1 Accuracy Improvements w/ <i>SEASONS</i>	16
5.1.2 Latency Sensitivity of <i>SEASONS</i>	17
5.1.3 Overhead of <i>SEASONS</i>	18
6 RELATED WORK	19
7 CONCLUSION	20
8 FUTURE WORK	21
8.1 Generalized Queue Dynamics for Granular, Intermittent Application Degradation	21
8.2 Granular Inference on the Edge	21
REFERENCES	23

ABSTRACT

Both energy-aware, batteryless intermittent systems and signal-aware adaptive sampling algorithms (ASA) aim to minimize sensor data accuracy under energy constraints in edge devices. Intuitively, combining both into a signal- & energy-aware solution would yield even better accuracy. Unfortunately, ASAs and intermittent systems rely on conflicting energy availability assumptions. So, a straightforward combination cannot achieve their combined benefits.

Therefore, we propose *SEASONS*, the first framework for signal- and energy-aware intermittent systems. *SEASONS* buffers signal data in time, monitoring queue dynamics to ensure the data is reported within a user-specified latency constraint. *SEASONS* improves sensor data accuracy by 31% without increasing energy.

CHAPTER 1

INTRODUCTION

Batteryless intermittent sensing systems are a recent advancement in embedded edge sensing Maeng and Lucia [2020], Lucia and Ransford [2015], Desai et al. [2022], Colin and Lucia [2016], Gobieski et al. [2019b], Hester et al. [2017], Yildırım et al. [2018], Denby and Lucia [2020], Afanasov et al. [2020], Bakar et al. [2021], Balsamo et al. [2015]. Intermittent systems sample environmental signals to provide accurate sensor data to a central server while using only harvested ambient energy stored in small capacitors.

Batteryless operation removes the many disadvantages of batteries (*e.g.* battery replacement, system size, and environmental concerns), making intermittent systems suitable for many domains from small satellites Denby and Lucia [2020] to infrastructure monitoring Afanasov et al. [2020]. However, accurate sensing is challenging given the unreliability of ambient energy sources such as solar. Thus, *energy-awareness*—*i.e.* adapting to scarce and variable energy—is an essential for intermittent runtime systems.

Signal-awareness, achieved via a family of Adaptive Sampling Algorithms (ASA), is another advancement in the field of sensing systems Silva et al. [2017], Chatterjea and Havinga [2008], Alippi et al. [2007], Murad et al. [2020], Campos et al. [2017]. ASAs balance the trade-off between energy and accuracy in an optimal manner by analyzing the sensed environmental signals. They ensure that energy is used to collect *significant samples*—*i.e.*, those that have the most influence on accuracy—and they avoid spending energy on less significant samples. Compared to non-adaptive techniques, ASAs improve the accuracy of sensed data and reduce the system energy requirements. Overall, both intermittent systems and ASAs aim to provide accurate sensor data despite energy limitations. Intuitively, these approaches should be combined into a signal- and energy-aware intermittent system to achieve the advantages of both.

Unfortunately, despite their shared goals, these two approaches rely on incompatible as-

sumptions. ASAs achieve their goal through sampling rate variation: significant sensor data are sampled at higher frequency, whereas less significant ones are sampled less. Since the sampling rate directly affects energy consumption, varying sampling rates produces fluctuating energy consumption. Windows of significant samples consume more energy and the less significant ones do the opposite. ASAs assume that these periods of high and low energy draw will balance over a battery’s lifetime (lasting potentially hours or days). However, intermittent sensors use capacitors that only store energy for milliseconds of operation Lucia and Ransford [2015]. Thus, they cannot sustain long periods of high energy consumption nor can they store the extra energy from low energy consumption windows. In short, existing intermittent systems cannot support the variable energy consumption required by ASAs.

To address this problem, we introduce *SEASONS*, a framework that enables signal-aware ASA operation on energy-aware intermittent systems. An ASA operating via *SEASONS* schedules and budgets to achieve a constant energy consumption, regardless of the signal. To achieve this, *SEASONS* uses time buffering. When the ASA needs to increase the sampling rate (because it has encountered significant samples), *SEASONS* collects the data, and immediately enqueues it, delaying processing and communication to maintain constant energy consumption. Doing so mitigates a period of high energy consumption typical of an ASA. When the ASA reduces the sampling rate, *SEASONS* fast tracks the processing and communication of the enqueued samples, maintaining the constant rate of energy consumption. The fast tracking uses up the the extra energy that could not be stored over a long term, avoiding wastage of surplus energy.

While time buffering resolves the incompatibility between the two approaches, it can cause arbitrarily delays in communicating the sensed data. Expired samples, data that has not been sent within application defined latency constraints Surbatovich et al. [2021], Hester et al. [2017], occur more frequently with time buffering. *SEASONS* avoids such scenarios by analyzing queue dynamics to adjust sampling rates dynamically and ensure all collected

data is processed and communicated within the latency constraint.

With *SEASONS*, energy-aware intermittent systems can employ signal-aware sampling policies to reduce the impact of low energy operation on sensor data accuracy. We find *SEASONS* improves data accuracy by an average of 31% across multiple sensor datasets and energy budgets compared to prior purely energy-aware intermittent systems.

This work makes the following contributions:

- We analyze how and why straightforward implementation of ASAs on intermittent systems fail.
- We propose practical solutions that circumvent these failures and present *SEASONS*, a novel framework that enables ASA operation on intermittent systems.

To our knowledge, we are the first to study adaptive sampling on intermittent systems, formalize the shortcomings of naive adoption, and offer a solution. We provide the first signal & energy-aware intermittent interface. Furthermore, we are the first to point out the importance of queue dynamics on an intermittent system and capitalize on it.

CHAPTER 2

BACKGROUND & MOTIVATION

We consider a common low-power embedded sensing device with the goal of providing accurate and timely data to a central server. This design is common in traditional embedded sensing Zhang et al. [2004] or even emerging intermittent setups Afanasov et al. [2020]. Purely Energy-aware Intermittent designs achieve this goal by altering sensor data quality based on available energy. A popular and common approach is uniformly reducing sample collection rate Bakar et al. [2021], Maeng and Lucia [2020]. In contrast, battery-backed signal- & energy-aware systems under the same energy constraints use an ASA to improve data accuracy by devoting energy to the most significant data samples Silva et al. [2017], Chatterjea and Havinga [2008], Alippi et al. [2007], Murad et al. [2020], Campos et al. [2017].

Figure 2.1 shows a comparison between the two approaches. The blue line shows an accelerometer signal sampled at 50Hz. The purple and orange lines show the inferred signal from the intermittent system and the ASA, respectively. Both the intermittent and battery-backed system are under energy constraints that permit collection of just 60% of possible samples; however, the ASA strategy results in 88% improvement in data accuracy compared to the intermittent result.

The improvement is due to the ASA’s selective collection of the samples. Unlike the intermittent system, the ASA equipped battery-backed system does not uniformly collect 6 samples out of 10 possible samples. Since the accelerometer signal begins with highly volatile behavior (figure 2.1 (b))—which the ASA considers significant—the ASA collects almost every possible sample during that time. Later (figure 2.1 (c)), the ASA considers the less variable signal less significant, so it collects 2 or 3 samples out of 10 possible samples. Over the entire operation window, the ASA improves accuracy under the same energy constraints. Combining this higher accuracy with the batteryless benefits of intermittent systems would be advantageous. However, a straightforward combination of both fails to achieve this higher

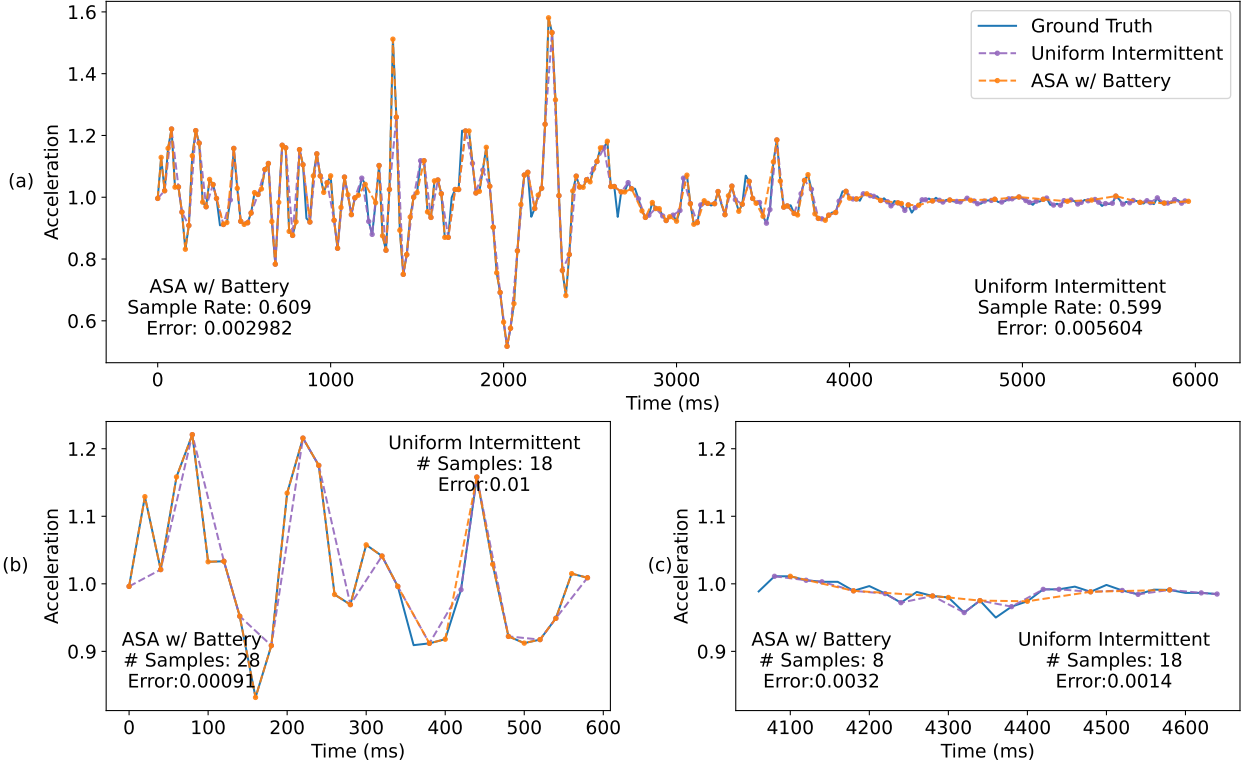


Figure 2.1: (a): Energy-Aware intermittent uniform sampling (purple) vs battery-backed ASA (orange) of accelerometer data over 6 seconds w/ 60% collection rate. The blue line (labeled ground truth) represents the real accelerometer signal sampled at 50Hz. (b): Behavior of sampling policies over a highly volatile phase. Adaptive sampling collects 10 more samples and achieves one order of magnitude lower error than uniform sampling. Since the adaptive method takes more samples, it consumes more energy than the uniform over this period. (c): Behavior of subsampling policies over a low volatility phase. The uniform strategy collects 10 more samples and achieves 2.2x lower error. Since the adaptive method takes less samples, it consumes less energy than uniform.

accuracy.

Typical operation of an ASA assumes a large energy buffer to support variable energy consumption rates. In figure 2.1 (a), the ASA strategy requires taking significantly more samples in the 0–3000 ms time range compared to the 4000–6000 ms range. The initial increase in sampling rate corresponds to a higher draw rate from the system’s energy buffer. Later, the reduced sampling rate reduces energy draw, balancing the total consumed energy across the entire operation window. An intermittent system with a buffer that can only support milliseconds of operation is unable to support the seconds of distance between this energy balancing. Ported directly to an intermittent system, the ASA cannot reliably alter the sampling rate without changing the energy consumption rate and causes power failures, wasting energy and reducing accuracy.

CHAPTER 3

SEASONS

SEASONS is a novel framework that combines the benefits of batteryless intermittent systems with signal-aware ASAs. *SEASONS* resolves the conflicting energy availability assumptions of these system by buffering in time, a concept that relies on the postponability of processing and communication tasks. Time buffering alone, however, is unaware of latency constraints and can lead to latency violations. Such violations reduce the accuracy, so *SEASONS* utilizes sample queue dynamics to anticipate and avoid such violations. This section explains these two key concepts.

3.1 Adaptive Sampling via Time Buffering

As studied by prior intermittent work, many parts of an intermittent application do not have a strict deadline Maeng and Lucia [2020]. Samples are numerous and their deadlines are hard, but processing and communication occur less regularly and their deadlines are more flexible. Thus, when feasible, a system can postpone processing or communication. We dub this operation as **time buffering**. Further, with time buffering, samples must wait in a non-volatile queue for processing and communication Hester et al. [2017], Maeng and Lucia [2020], Yildirim et al. [2018]. We call this queue the *sample queue*.

With time buffering, an intermittent system can employ an ASA without altering the energy consumption as follows. Given significant samples, the system increases sampling rate and enqueues samples until later processing (*e.g.* 0–3000 ms time range in figure 2.1); given less significant samples, the system reduces sampling rate, fast-tracks sample processing and communication to empty the sample queue (*e.g.* 4000–6000 ms time range in figure 2.1).

However, time buffering maintains no control over the number of samples in the queue. In general, the system cannot predict when a less significant window of samples will arrive to

balance the earlier window of significant samples. This unpredictably could cause the latency or queue size to grow arbitrarily long. This is unacceptable for embedded edge devices as they are resource (including memory) limited and must operate within an application-defined latency.

Respecting application latency constraints is vital in modern intermittent runtimes Hester et al. [2017], Surbatovich et al. [2021]. This concept, dubbed *freshness*, ensures that samples are processed in a timely manner, despite potential power failures and energy shortages. Once a freshness violation has been detected (*i.e.* a sample has *expired* its latency constraint), these systems drop the sample and prevent it from progressing in the application pipeline so that additional energy is not wasted on processing and communicating an expired sample.

Figure 3.1 depicts the operation of a time buffered intermittent system running and ASA with freshness enforcement under the same constraints as figure 2.1. While freshness enforcement is vital to prevent energy wastage from sending expired samples, it results in many samples being dropped before they are communicated to the server. These sample drops then reduce the overall accuracy. In figure 3.1, 13% of taken samples expire before they can be sent, leading to an a signal inference that misses important events in the signal and an error rate more that of energy-aware intermittent sampling without the ASA (figure 2.1 (a)).

3.2 Expiration Mitigation via Sample Queue Dynamics

Instead of enforcing freshness only when samples expire, we propose anticipating and preventing latency violations by analyzing sample queue dynamics. As intermittent systems operate, they enqueue and dequeue their sample queue at various rates based on energy availability and mode of operation. We refer to these rates as **sample queue dynamics**.

Prior freshness enforcement mechanisms that drop samples do not align well with time buffering—as shown in 3.1—because expired samples are opportunity costs. Expired samples




Figure 3.1: Adaptive subsampling with time buffering and freshness enforcement by dropping samples.

both (1) fail to provide data to the server and (2) consume valuable energy that could have been better spent to ensure other sample data does get to the server. We can avoid the opportunity cost by mitigating expirations, ensuring that any energy spent on sampling contributes to the overall accuracy of the data.

We enforce this mitigation by analyzing sample queue dynamics. As the intermittent system with time buffered ASA operates, enqueueing and dequeuing rates vary based on the ASA. Certain enqueueing and dequeuing rates are more likely to fill the sample queue and cause future sample expirations. For example, a period of significant samples can enqueue so many samples that no future processing and communication can dequeue in time; these unsafe rates cause expirations. Hence, if the current queue dynamics match these dangerous rates, *SEASONS* limits the ASA to slower rates to avoid sample expirations.

Time buffering is central to ensuring ASA operation on signal-aware intermittent systems without frequent power failure. But time buffering is not without issues, so analyzing sample queue dynamics to mitigate expiration keeps the time buffering under control and prevents it from becoming disadvantageous. Through this combined operation, *SEASONS* enables the accuracy of signal-awareness on intermittent systems. The next section discusses the technical aspects of this design in more depth.

CHAPTER 4

SEASONS FRAMEWORK

SEASONS requires only small changes to work with existing intermittent frameworks and adaptive sampling algorithms. Thus, for brevity we focus on *SEASONS*'s novel aspects: a wrapper for interfacing with ASAs (subsection 4.1), a scheduling policy for time buffering (subsection 4.2), and a controller for analyzing sample queue dynamics and enforcing expiration mitigation (subsection 4.3).

Figure 4.1 depicts the *SEASONS* framework on an intermittent setup. The sampling rate is determined by the ASA (1). The controller considers sampling rate (1) energy availability (2), and sample queue dynamics (3) to determine a safe sampling rate (4), balancing signal-awareness and sample expiration mitigation. Finally, the safe sampling rate is communicated to the scheduler (4), which schedules sampling, processing, and communication based on the safe sampling rate.

These additional mechanisms can add significant overhead to processing, specially given the processing limitation of low power MCUs such as MSP430 (*e.g.* no FLU) MSP. To minimize this overhead, *SEASONS* employs preprocessing tools and runtime lookup tables. Subsection 5.1.3 discusses the overhead in more detail.

4.1 ASA and the Wrapper

The ASA module accomplishes the signal-awareness aspect of *SEASONS* by determining the current sampling rate based on prior samples. For instance, the linear ASA Chatterjea and Havinga [2008] alters its rate based on differences in consecutive samples. This ASA considers big differences to be significant, so it increases the sampling rate if these differences are more than expected and reduces sampling rate if they are not. While various ASAs with their own definitions of significant exist, the majority work based on the above principle of

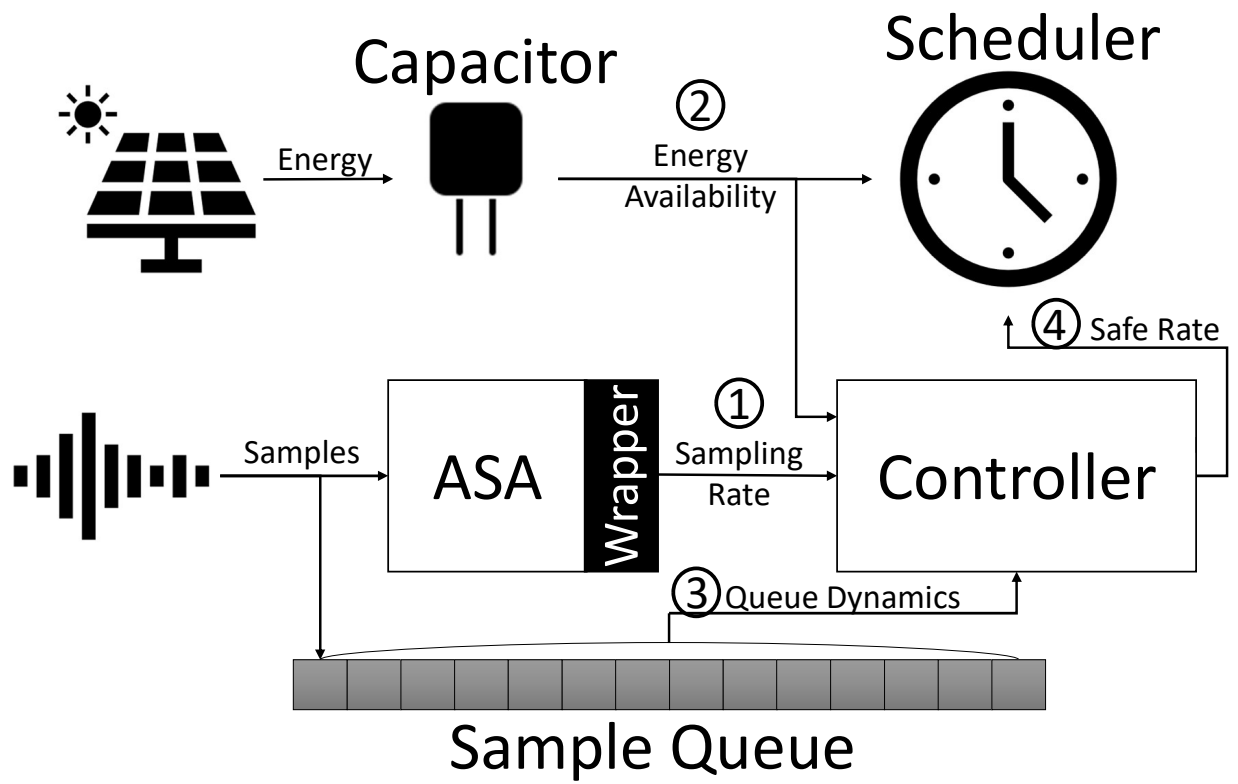


Figure 4.1: *SEASONS*'s overview

comparing recent samples to a threshold that determines significance Alippi et al. [2007], Chatterjea and Havinga [2008], Silva et al. [2017] To ease communication between system components and different ASAs, we also introduce an ASA wrapper.

4.2 Scheduler

The scheduler module manages the time buffering. Prior schedulers determined the sampling rate internally based on energy availability; *SEASONS*'s scheduler leaves the sampling rate determination to the controller. This external sampling rate determines when the scheduler should schedule sampling tasks. These tasks consume some ratio of the available harvested energy. The scheduler then uses the rest of the energy to schedule processing and communication tasks. This strategy maintains a constant energy consumption rate for an intermittent system and achieves time buffering through postponing processing and communication tasks.

4.3 Controller

The controller module is in charge of analyzing sample queue dynamics to mitigate expirations. Expirations occur when the systems has too many samples in the queue and dequeuing rate is not sufficient to process and communicate the samples by their latency deadlines. We refer to the sample queue count in this case as the *critical* count.

When the system has a critical number of samples in the queue, the controller reduces the maximum possible ASA rate to a uniform rate which matches the dequeue rate. In doing so, the system still takes significant samples, but does so at a predictable rate. This strategy matches the enqueueing and dequeuing rate, ensuring that the queue size does not increase in the critical phase and thus no expirations can occur (as the oldest sample is always dequeued). However, when taking less significant samples, the system follows an adaptive strategy, which reduces the queue size. This rate reduction remains in place until

the system moves out of the critical phase.

CHAPTER 5

EXPERIMENTAL RESULTS

5.1 Experimental Setup

We evaluate *SEASONS* with our in house simulator on a variety of configurations.¹ The simulation operates as follows. We use sensor datasets from table 5.1 as ground truth for our sensing simulations. These datasets represent a variety of events and exhibit various signal properties. Sampling policies (either uniform or an ASA) determine which entries should be taken from these datasets. The intermittent system collects the samples, processes the collected data and transmits them to a server. At each step, the simulation ensures that sufficient energy is present to accomplish a task, regardless of whether the simulation is in battery-backed or intermittent mode. Failure occurs otherwise, reducing accuracy. Latency violations result in sample expirations, which are not transmitted to the server and thus reduce accuracy. The server reconstructs the full data sequence using linear interpolation to infer missing data points and measures the mean absolute error (MAE) between the reconstructed signal and ground truth.

Table 5.1: Properties of Evaluation Datasets.

Dataset	Labels	Mean	Std. Deviation	CV
EOGFang and Shinozaki [2018]	12	101.17	113.81	1.12
EpilepsyVillar et al. [2016]	4	1.39	0.69	0.50
PasswordsPas	5	0.81	0.59	0.73
StrawberryHolland et al. [1998]	2	0.72	0.69	0.95
TiselacIenco et al. [2017]	9	806.43	333.84	0.41
TrajectoriesWilliams et al. [2006]	20	1.45	0.78	0.54
UCIAnguita et al. [2012]	12	1.03	0.19	0.19

1. All code and datasets will public released after the anonymous reviewing.

Table 5.2: Example Power Cost of Operation in an Example Edge Sensor.

Operation	Power Cost (mW)	Ratio of Total Power
Sampling	2.1	20%
Encryption	0.2	1%
BLE	8.4	79%
Total	10.7	100%

We use the the power requirements for sampling, encryption, and communication from table 5.2 over a 1 second window as guide for the power consumption metrics in our simulation. The application samples 50 accelerometer measurements using an IMU IMU, encrypts them on a TI MSP430 FR5994 MCU MSP and communicates them via a CC2640R2 BLE chip CC2. The power measurements are taken using the EnergyTrace tool Ene. Table 5.2 displays the measured values. It should be noted that power ratios represented in table 5.2 closely mirror prior deployed intermittent applications Afanasov et al. [2020], Desai et al. [2022].

We compare the following,

- *Energy-Aware Intermittent (EI)*. This represents state-of-the-art energy aware intermittent systems that use uniform sampling to meet energy constraints, *e.g.* CatNap Maeng and Lucia [2020] & REHASH Bakar et al. [2021].
- *Signal- & Energy-Aware Battery-backed (SEB)*. This represents a state-of-the-art ASA with a battery Chatterjea and Havinga [2008].
- *SEASONS no Latency Guarantee (LG)*. *SEASONS* without the analysis done by the controller. This configuration does not anticipate and mitigate sample expirations, but will not communicate expired samples).
- *SEASONS*. The approach in this paper: signal- and energy-aware intermittence while respecting user-defined latency constraints.

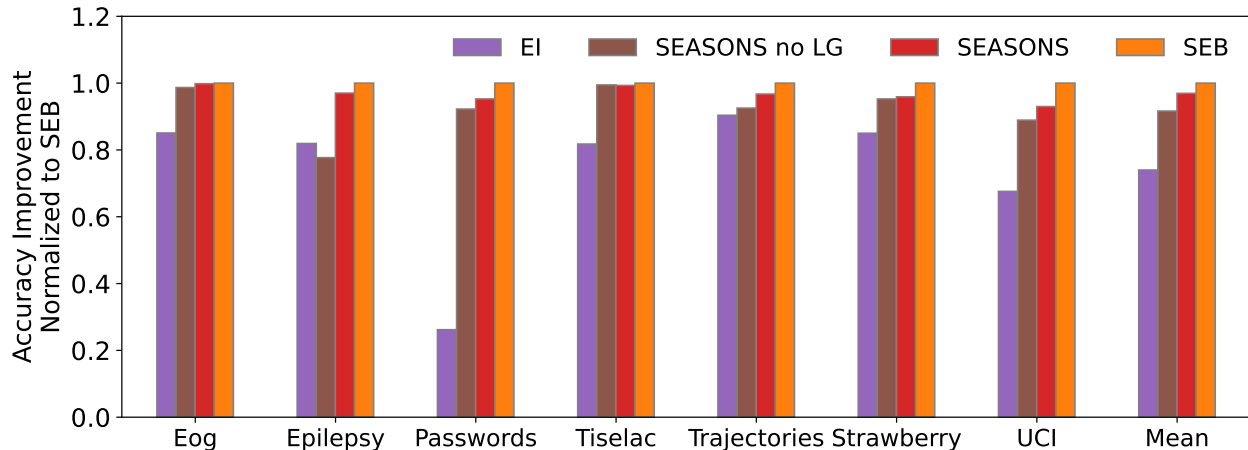


Figure 5.1: Accuracy improvement of *SEASONS* over energy-aware intermittent w/ uniform sampling. *SEASONS* improves the accuracy of intermittent systems and approaches the feasible accuracy of a battery-backed ASA.

We experiment with the above configurations over a range of latency and energy constraints to properly evaluate *SEASONS*. To compare improvement in accuracy, we measure improvement over EI normalized to SEB.

5.1.1 Accuracy Improvements w/ *SEASONS*

We simulate each configuration on the datasets from table 5.1 with 7 different energy constraints and reasonable latency constraints to understand *SEASONS*' accuracy improvements over prior work. For each run, we report the accuracy improvement in figure 5.1. Overall, *SEASONS* without latency guarantees (brown), *SEASONS* (red), and the battery-backed (orange) configuration achieve respectively 22%, 31%, and 35% improvements in accuracy over the uniform strategy employed by prior intermittent systems (purple).

Compared to the intermittent baseline, *SEASONS* without latency guarantees improves accuracy, but does not do so reliably. For instance, with the epilepsy dataset, *SEASONS* without latency guarantees results in many sample expirations and produces accuracy below existing intermittent systems. In contrast, *SEASONS* with latency guarantees reliably improves accuracy for all datasets, achieving close to existing ASAs, but without a battery.

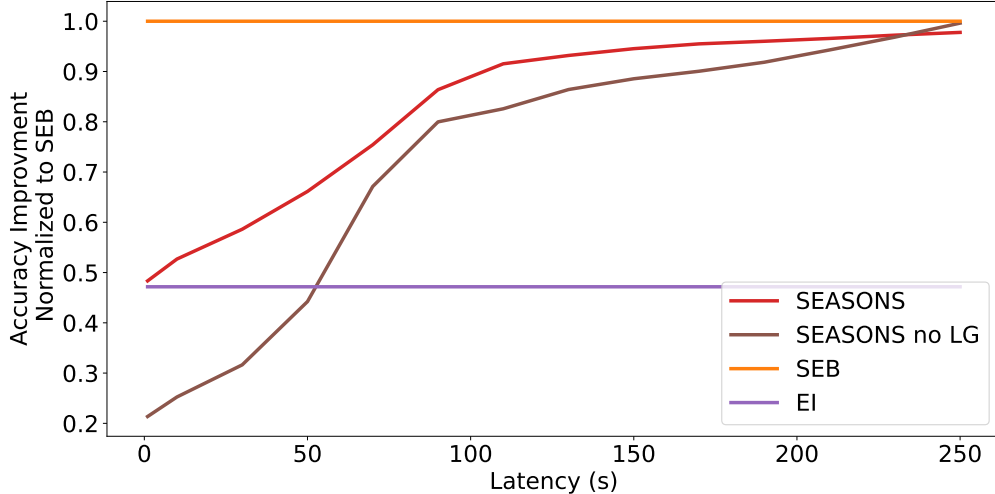


Figure 5.2: *SEASONS* Comparison between EI, SEB, & EI w/ *SEASONS* as latency constraint changes for a collection rate of 60%.

This scenario demonstrates that SEASONS correctly combines an intermittent power setup with an adaptive sampling algorithm to gain signal-aware accuracy benefits with a batteryless design.

5.1.2 Latency Sensitivity of *SEASONS*

We simulated each configuration on a range of latency constraints on the *Linear* ASA Chatterjea and Havinga [2008] to investigate the accuracy impact of the user-specified latency. For brevity, we only show accuracy improvements for the UCI dataset in figure 5.2 (other datasets exhibit a similar trend). For very short constraints, *SEASONS* achieves accuracy similar to that of an energy-aware configuration, which is reasonable given that there is little room for postponability. On the other hand, increasing latency constraints improves accuracy as *SEASONS* operates more akin to a battery powered system with an ASA. Moreover, *SEASONS* with no latency guarantees follows a similar trend, but performs much worse given tight constraints

This experiment demonstrates relaxing latency constraints improves *SEASONS*' accuracy improvement over prior work; however, *SEASONS* improves accuracy even with tight latency

constraints. Further, *SEASONS*' latency guarantees are vital in ensuring that *SEASONS* is beneficial to accuracy, even given short constraints. In summary, *SEASONS is never worse than existing approaches, and given any latency flexibility improves on existing work.*

5.1.3 Overhead of *SEASONS*

To measure the overhead of *SEASONS*, we employ a uniform policy (as in prior intermittent work) and *SEASONS* on the TI MSP430FR5994 MSP and emulate their operation on the UCI dataset. *SEASONS* retains the accuracy improvements of the simulation, but incurs an overhead of 1.6x in processing cost compared to the uniform policy. However, the 60% increase in processing power amounts to only 20 μ W, less than 0.2% of the entire system's power consumption and 2 orders of magnitude less than sampling or communication (table 5.2). *Hence, SEASONS's overhead is negligible.*

CHAPTER 6

RELATED WORK

Intermittent Systems Batteryless deployment and its benefits have been discussed extensively in prior work Maeng and Lucia [2020], Lucia and Ransford [2015], Desai et al. [2022], Colin and Lucia [2016], Gobieski et al. [2019b], Hester et al. [2017], Yildirim et al. [2018], Denby and Lucia [2020], Afanasov et al. [2020], Bakar et al. [2021]. Frequent checkpointing Balsamo et al. [2015] and task based programming Colin and Lucia [2016] are popular methods of ensuring forward progress in prior work. CatNap Maeng and Lucia [2020] & REHASH Bakar et al. [2021] are state of the art energy-aware intermittent systems that adapt to scarce energy conditions. To the best of our knowledge, no prior intermittent systems have deployed adaptive sampling because they could not support variable energy consumption.

Adaptive Sampling Algorithms ASAs employ signal awareness to balance energy consumption and sensor data accuracy on low power embedded systems. Some algorithms rely on linear or autoregressive models Alippi et al. [2007], Chatterjea and Havinga [2008], Silva et al. [2017], while others employ sophisticated methods such as Reinforcement Learning Murad et al. [2020] or Neural Networks Campos et al. [2017], Kannan and Hoffmann [2021]. To our knowledge, all prior ASAs assume a large energy buffer and cannot be directly combined with batteryless systems.

CHAPTER 7

CONCLUSION

SEASONS is the first framework to achieve signal- & energy-aware operation on a batteryless intermittent operation. Through Time Buffering, the ASA employed by *SEASONS* operates without requiring a big energy bank (*e.g.* battery). Further, *SEASONS* efficiently avoids expirations via analyzing Sample Queue dynamics. The framework operates with minimal overhead and can achieve 31% more accuracy compared to traditional energy-aware intermittent sensing. Given the negligible cost and significant improvements in accuracy, taking smarter samples in intermittent systems via *SEASONS* is worth the effort.

CHAPTER 8

FUTURE WORK

8.1 Generalized Queue Dynamics for Granular, Intermittent Application Degradation

State of the art intermittent systems degrade application quality to maintain reasonable performance in the face of variable energy conditions. Degradation has been done in small, granular steps, as demonstrated by Maeng and Lucia [2020], or via limited tiers with big energy consumption gaps by Bakar et al. [2021]. Small, granular steps can more closely match performance to energy conditions, but current state of the art solutions degrade performance without regard for application pipeline latency. In contrast, limited tiers are not energy optimal. Each tier is individually programmed, so pipeline latency will not deteriorate due to degradation.

We believe that generalizing queue dynamics to the entire intermittent application program flow can achieve granular degradation without adverse effects on latency. Queue dynamics inform the system of the application queues that are backlogged due to degradation. Using this information, the system will be able to make smarter degradation choices without deteriorating application latency. We plan to evaluate our hypothesis by comparing our idea to state of the art degradation schemes and demonstrating that 1, our solution reduces application latency and 2, our solution spends less energy on expired data.

8.2 Granular Inference on the Edge

Data transmission has a high energy cost in intermittent systems. Hence, inference on the edge is highly valuable for intermittent sensors: inference reduces the number of transmitted bytes, but maintains the quality of the transmitted information Gobieski et al. [2019a]. In

short, inference saves energy in resource constrained intermittent systems. Prior work accomplishes inference via compressed DNN models Gobieski et al. [2019a], though such models do not take variable energy conditions into account. Islam and Nirjon [2020], Montanari et al. [2020] purpose degrading neural network inference with early exit points to accommodate for changing energy conditions. However, the proposed methodologies are not granular enough for the varying energy conditions.

We believe that ensemble inference methods such as random forests are a better fit for intermittent systems. For instance, degradation of random forests can occur via running a portion of the decision trees and since there are numerous trees, random forests allow for more granular degradation. We plan to compare ensemble methods to DNNs with multiple exist points to demonstrate that 1, granular inference on the edge is preferable and 2, ensemble methods can achieve such granular inference.

REFERENCES

- Ti msp430 cc2640r2 datasheet. URL <https://www.ti.com/lit/ds/symlink/cc2640r2f.pdf>.
- Ti msp430 energytrace technology. URL <https://www.ti.com/tool/ENERGYTRACE>.
- lsm6dsox datasheet. URL <https://www.st.com/resource/en/datasheet/lsm6dsox.pdf>.
- Ti msp430 fr5994 datasheet. URL <https://www.ti.com/lit/ds/symlink/msp430fr5994.pdf>.
- Graphical password dataset.
- Mikhail Afanasov, Naveed Anwar Bhatti, Dennis Campagna, Giacomo Caslini, Fabio Massimo Centonze, Koustabh Dolui, Andrea Maioli, Erica Barone, Muhammad Hamad Alizai, Junaid Haroon Siddiqui, and Luca Mottola. Battery-less zero-maintenance embedded sensing at the mithræum of circus maximus. *SenSys '20*, 2020. doi:10.1145/3384419.3430722.
- Cesare Alippi, Giuseppe Anastasi, Cristian Galperti, Francesca Mancini, and Manuel Roveri. Adaptive sampling for energy conservation in wireless sensor networks for snow monitoring applications. In *2007 IEEE International Conference on Mobile Adhoc and Sensor Systems*, 2007. doi:10.1109/MOBHOC.2007.4428700.
- Davide Anguita, Alessandro Ghio, Luca Oneto, Xavier Parra, and Jorge L. Reyes-Ortiz. Human activity recognition on smartphones using a multiclass hardware-friendly support vector machine. In *Ambient Assisted Living and Home Care*, pages 216–223. Springer, 2012.
- Abu Bakar, Alexander G. Ross, Kasim Sinan Yildirim, and Josiah Hester. Rehash: A flexible, developer focused, heuristic adaptation platform for intermittently powered computing. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, 2021. doi:10.1145/3478077.
- Domenico Balsamo, Alex S. Weddell, Geoff V. Merrett, Bashir M. Al-Hashimi, Davide Brunelli, and Luca Benini. Hibernus: Sustaining computation during intermittent supply for energy-harvesting systems. *IEEE Embedded Systems Letters*, 2015. doi:10.1109/LES.2014.2371494.
- Victor Campos, Brendan Jou, Xavier Giro-i Nieto, Jordi Torres, and Shih-Fu Chang. Skip rnn: Learning to skip state updates in recurrent neural networks, 2017.
- Supriyo Chatterjea and Paul Havinga. An adaptive and autonomous sensor sampling frequency control scheme for energy-efficient data acquisition in wireless sensor networks. In Sotiris E. Nikolettseas, Bogdan S. Chlebus, David B. Johnson, and Bhaskar Krishnamachari, editors, *Distributed Computing in Sensor Systems*. Springer Berlin Heidelberg, 2008. ISBN 978-3-540-69170-9.

- Alexei Colin and Brandon Lucia. Chain: Tasks and channels for reliable intermittent programs. OOPSLA 2016, 2016. doi:10.1145/2983990.2983995.
- Bradley Denby and Brandon Lucia. Orbital edge computing: Nanosatellite constellations as a new class of computer system. ASPLOS '20, 2020. doi:10.1145/3373376.3378473.
- Harsh Desai, Matteo Nardello, Davide Brunelli, and Brandon Lucia. Camaroptera: A long-range image sensor with local inference for remote sensing applications. *ACM Trans. Embed. Comput. Syst.*, 2022. doi:10.1145/3510850.
- Fuming Fang and Takahiro Shinozaki. Electrooculography-based continuous eye-writing recognition system for efficient assistive communication systems. In *Ambient Assisted Living and Home Care*, 2018.
- Graham Gobieski, Brandon Lucia, and Nathan Beckmann. Intelligence beyond the edge: Inference on intermittent embedded systems. ASPLOS '19, New York, NY, USA, 2019a. Association for Computing Machinery. doi:10.1145/3297858.3304011.
- Graham Gobieski, Brandon Lucia, and Nathan Beckmann. Intelligence beyond the edge: Inference on intermittent embedded systems. ASPLOS '19, 2019b. doi:10.1145/3297858.3304011.
- Josiah Hester, Kevin Storer, and Jacob Sorber. Timely execution on intermittently powered batteryless sensors. SenSys '17, 2017. doi:10.1145/3131672.3131673.
- J K Holland, E K Kemsley, and R H Wilson. Use of fourier transform infrared spectroscopy and partial least squares regression for the detection of adulteration of strawberry purées. *Journal of the Science of Food and Agriculture*, 1998.
- Dino Ienco, Raffaele Gaetano, Claire Dupaquier, and Pierre Maurel. Land cover classification via multitemporal spatial data by deep recurrent neural networks. *IEEE Geoscience and Remote Sensing Letters*, 2017. doi:10.1109/LGRS.2017.2728698.
- Bashima Islam and Shahriar Nirjon. Zygard: Time-sensitive on-device deep inference and adaptation on intermittently-powered systems. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, 2020. doi:10.1145/3411808.
- Tejas Kannan and Henry Hoffmann. Budget rnns: Multi-capacity neural networks to improve in-sensor inference under energy budgets. In *2021 IEEE 27th Real-Time and Embedded Technology and Applications Symposium (RTAS)*, 2021. doi:10.1109/RTAS52030.2021.00020.
- Brandon Lucia and Benjamin Ransford. A simpler, safer programming and execution model for intermittent systems. PLDI '15, 2015. doi:10.1145/2737924.2737978.
- Kiwan Maeng and Brandon Lucia. Adaptive low-overhead scheduling for periodic and reactive intermittent execution. PLDI 2020, 2020. doi:10.1145/3385412.3385998.

- Alessandro Montanari, Manuja Sharma, Dainius Jenkus, Mohammed Alloulah, Lorena Qendro, and Fahim Kawsar. Eperceptive: Energy reactive embedded intelligence for battery-less sensors. *SenSys '20*, 2020. doi:10.1145/3384419.3430782.
- Abdulmajid Murad, Frank Alexander Kraemer, Kerstin Bach, and Gavin Taylor. Information-driven adaptive sensing based on deep reinforcement learning. *IoT '20*, 2020. doi:10.1145/3410992.3411001.
- João Marco C. Silva, Kalil Araujo Bispo, Paulo Carvalho, and Solange Rito Lima. Lite-sense: An adaptive sensing scheme for wsns. In *2017 IEEE Symposium on Computers and Communications (ISCC)*, 2017. doi:10.1109/ISCC.2017.8024689.
- Milijana Surbatovich, Limin Jia, and Brandon Lucia. Automatically enforcing fresh and consistent inputs in intermittent systems. *PLDI 2021*, 2021. doi:10.1145/3453483.3454081.
- Jose R Villar, Paula Vergara, Manuel Menéndez, Enrique de la Cal, Victor M González, and Javier Sedano. Generalized models for the classification of abnormal movements in daily life and its applicability to epilepsy convulsion recognition. 2016.
- Ben H. Williams, Marc Toussaint, and Amos J. Storkey. Extracting motion primitives from natural handwriting data. In *Artificial Neural Networks – ICANN 2006*, pages 634–643. Springer, 2006.
- Kasim Sinan Yildirim, Amjad Yousef Majid, Dimitris Patoukas, Koen Schaper, Przemyslaw Pawelczak, and Josiah Hester. Ink: Reactive kernel for tiny batteryless sensors. *SenSys '18*, 2018. doi:10.1145/3274783.3274837.
- Pei Zhang, Christopher M. Sadler, Stephen A. Lyon, and Margaret Martonosi. Hardware design experiences in zebranet. *SenSys '04*, 2004. doi:10.1145/1031495.1031522.