

An Analysis of Design Parameters for Energy Management of Wireless Sensor Devices

Asbjørn Engmark Espe*  and Sondre Ninive Andersen† 
Dept. of Engineering Cybernetics,
Norwegian University of Science and Technology,
Trondheim, Norway
Email: {*asbjorn.e.espe, †sondre.n.andersen}@ntnu.no

Pierluigi Salvo Rossi 
Dept. of Electronic Systems,
Norwegian University of Science and Technology,
Trondheim, Norway
Email: pierluigi.salvorossi@ntnu.no

Frank Alexander Kraemer 
Dept. of Information Security and Communication Technology,
Norwegian University of Science and Technology,
Trondheim, Norway
Email: kraemer@ntnu.no

Geir Mathisen
Dept. of Engineering Cybernetics,
Norwegian University of Science and Technology,
Trondheim, Norway
Email: geir.mathisen@ntnu.no

Abstract—Using a general model of a wireless sensor in a resource-scarce environment, we investigate in this paper how internal and external parameters affect utility. Three metrics are proposed to capture the application-specific nature of utility, and in an effective manner highlight the trade-offs encountered. We consider a simple energy management policy for the wireless sensor with two parameters: a value-of-information threshold and a buffering capacity. Employing the utility metrics, we examine how this policy may be tailored to different environments by tuning of its parameters. Furthermore, we conduct an in-depth study of the parameter sensitivity in an example environment, and show how such analyses may be used to inform sensor design. In conclusion, we show that even a simple policy will exhibit distinctly different modes of operation depending on parametrisation, and emphasise the importance of a sufficient understanding of the parameter space in which the system resides.

Index Terms—Energy management, intelligent sensors, energy harvesting, condition monitoring, wireless communication, scheduling algorithms.

I. INTRODUCTION

Ensuring the safety and reliability of transportation infrastructure is of paramount importance, as failure to properly assess the integrity of such structures can lead to the need for expensive repairs, or in the worst case to catastrophic collapse and loss of life. As early as 2001, the BRIME project (*Bridge Management in Europe*) [1] reported that a substantial number of bridges in the EU presented deficiencies—with percentages for France, Germany, the UK, and Norway reported to be 39%, 37%, 30%, and 26%, respectively. In August 2018, the effects of improper maintenance were demonstrated with the collapse of the Ponte Morandi highway bridge in northern Italy, raising further concerns about the state of infrastructure all across Europe [2].

In an effort to detect structural degradation and preemptively correct failures, a significant amount of resources is routinely expended by road and rail authorities to survey the condition of infrastructure. Condition monitoring using wireless sensors has become an increasingly attractive approach to

supplement these maintenance schemes, as the required sensor hardware has become cheaper, more power-efficient, and more accurate. For ease of deployment, these systems are often battery powered, and therefore have a limited life-time. By employing energy harvesting, the life-time may be increased substantially. The resulting energy-constrained environment naturally presents the challenge of how to maximise the quality-of-service within a tight energy budget. Typically part of the design process, such optimisation can help inform the dimensioning of the system to meet desired behaviour requirements.

Sensor systems will typically be designed to satisfy a certain requirement specification. In addition to functional requirements, a set of performance goals are usually formulated—for instance in terms of sample rate, latency, or system uptime. A common technique used in the literature when analysing systems based on such aggregate metrics is to use an analytic or stochastic approach to model how parameter changes impact the behaviour of the system [3]–[9]. This kind of analysis is powerful and can provide a thorough understanding of the system under analysis, as it attempts to reveal the underlying mechanisms that governs system behaviour. However, in the design phase of a given system, formulating an analytical model using such aggregate measures can be challenging. And once a system is modelled and analysed, incorporating new parameters may require replacing the entire model.

In this paper, we employ an alternative approach where the system is modelled directly—and as realistically as desired—after which we perform time-domain simulations across the parameter space to assess its performance. This has the disadvantage that it provides a less direct insight into the underlying, abstract relationships, since any conclusions will have to be drawn from simulated empirical data. Nonetheless, it allows for much finer control over the exact actions and decisions that the system performs while accounting for the passage of time. Moreover, this modelling approach is more useful as a tool for design of IoT systems, as it allows for rapid

prototyping, gradual inclusion of parameters and metrics, and targeted, parametric simulations to determine weaknesses and optimal operating points.

We introduce a model for the data pipeline of a general wireless sensor and describe the actions the sensor may take. Three metrics of utility are formulated and used to gauge the sensor's performance in terms of quantity, urgency, and quality. Using this model, we consider a simple energy management policy with two key parameters: a value-of-information threshold and a buffering capacity. In the context of the utility metrics, the resulting system is analysed in a range of situations and parametrisations, exemplifying both general design considerations to be aware of, and specific limitations with the policy used. At the current stage of research, our work is a preliminary study of this approach to energy management design. However, the approach has the potential to be widely employed as part of sensor node design.

II. RELATED WORKS

1) *Vibration analysis of bridges*: Our instrumentation scenario uses vibration analysis to estimate the structural health of bridges. This has been shown to be a valid approach, for instance by using low-cost accelerometers to record vibration and a damped harmonic oscillation model to estimate the bridge's modal parameters [10].

In more general terms, there are three main approaches of modal analysis of bridges through measurements of vibration.

- 1) Structure excitation by a controlled, artificial force and measurement of the corresponding response [11].
- 2) Measurements of both ambient excitation forces and the corresponding response [12].
- 3) Measurement of the bridge's response exclusively [10], [13].

To enable ubiquitous monitoring of the vast body of structures in unsatisfactory condition, only 2) and 3) are practically and economically viable solutions. Omidalizarandi *et al.* [13] argue that given only measurements of bridge responses for a sufficient length of time, a robust and automatic vibration analysis procedure can provide adequate results. They describe how a portion of the measurement data may be selected to capture the *ring-down* phase of an induced vibration. This is the phase in which the structure freely vibrates without any external forces and thereby exhibits exponential decay in vibration amplitude. Measurements of this decay allows estimation of the structure's modal parameters.

2) *Energy management*: There is a substantial body of work in the literature proposing and examining energy management policies for specific systems, or narrowly defined models. Among recent solutions, we find statistical quality-of-service-driven energy control policies developed by Gao *et al.* [6] to maximise certain energy harvesting efficiency metrics. Pan *et al.* [7] analysed simultaneous wireless information and power transfer, and derived closed-form expressions of the statistical properties of the system, validating these expressions through Monte Carlo simulations. Patil *et al.* demonstrated in [8] how an optimal transmission policy could be formulated

for systems where transmission opportunities are externally determined, and showed that the decision of whether or not to transmit can be made by comparing the value of the information to a threshold dependent on battery state. A stochastic energy management scheme was formulated by Ahmed *et al.* [14] to handle the uncertainty inherent in solar energy harvesting. Hanschke and Renner [15] proposed an algorithm to schedule interdependent, atomic tasks to meet timing and energy constraints, given an energy harvest prediction. Finally, Draskovic and Thiele introduced in [16] a finite-horizon approach with the aim of maximising the lifetime of a node while guaranteeing a minimal energy use.

There is significantly less research regarding more general techniques for assessing arbitrary wireless sensor systems in terms of how design parameters affect utility metrics. Niyato *et al.* [3] formulated an analytical, probabilistic node model based on a multidimensional Markov chain to examine the likelihood of packets being dropped or blocked, using simulation to validate the model. Sharma *et al.* [4] studied a model of a wireless sensor with an energy and a data buffer, and analytically obtained policies that were either optimal with regards to throughput or with regards to latency. An optimisation strategy was presented by Moser *et al.* in [5] for systems employing energy harvesting, taking into account the uncertainty of future harvesting yields. Lastly, the inherent trade-off between energy consumption and measurement accuracy was investigated by Kraemer *et al.* [9], who showed how this can be controlled by the sampling interval.

A common theme for the solutions in the literature is the usage of abstract or stochastic models in an attempt to derive analytic, closed-form relationships. While providing great insight, such methods are of limited applicability during the design phase of wireless systems due to the complexity of system modelling. We attempt to present a more useful approach to this end.

III. SYSTEM MODELLING AND SIMULATION

We consider the the scenario of a wireless sensor placed on a road or railway bridge. The sensor is to provide information concerning the bridge's structural integrity by sampling ambient vibration, and to transmit this information to a central unit wirelessly. The sensor is fitted with a small solar panel to harvest energy with which to replenish its energy reserves.

Vibration ring-down events are generated by the environment whenever the bridge is excited by a vehicle and the induced vibrations are allowed to freely decay. Such events can be used to estimate the bridge's modal parameters, and thus render a certain image of its structural health. The sensor can sample these ring-down events and assess the degree to which the signal envelope follows the expected exponential decay. Aligning with the concept of value-of-information (VoI) [17], we assume that this assessment results in a score expressing the potential value of the observation. In a more energy-intensive task, the sampled vibration event is processed by a DSP subsystem to estimate the bridge's modal parameters, along with a more confident estimate of VoI. Subsequently,

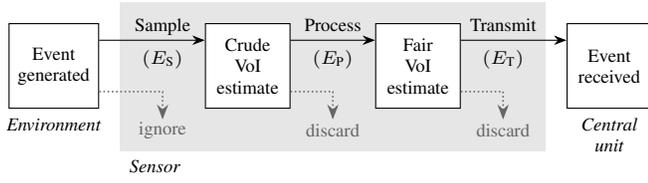


Fig. 1: A general model of a wireless sensor's operation from event generation by the environment to transmission to the central unit. Each task consumes some energy E_S , E_P , or E_T .

the sensor may transmit these modal parameters to the central unit.

This behaviour pattern can be generalised to many different situations. Conceptually, many Internet-of-Things (IoT) devices are tasked with condition monitoring by regularly reporting measurements of some phenomenon to a central unit by wireless transmission. Such devices are almost exclusively powered by a limited energy accumulator (e.g., a battery or supercapacitor), and in many cases employ some form of energy harvesting. The IoT device is naturally restricted in terms of energy usage, and will therefore operate according to some form of energy management policy to determine what to do at each point in time. As such, the sensor's decision pattern should be planned and evaluated to ensure a satisfactory strategy.

A. System model

The general operation of our wireless sensor can be modelled as shown in Figure 1. The environment of the sensor will sporadically generate vibration events, which the sensor can consume a small amount energy to (1) *sample*, or be forced to ignore as a result of insufficient energy. Sampling an event produces a packet of raw measurement data and a crude estimate of its informational value. The sensor can spend an additional amount of energy to (2) *process* a sampled event in more detail to prepare it for transmission, and get a better sense of the event's VoI. Ultimately, the sensor can expend a large amount of energy to (3) *transmit* a processed event to the central unit. At any point in this process, the sensor may also discard the event.

1) *Energy harvesting*: The energy harvesting yield of the device at a given point in time depends on the overall efficiency of the energy harvesting system and the underlying energy harvesting potential. Solar energy is employed in this work, meaning that the effective power P harvested by the system may be formulated as

$$P = A\eta G, \quad (1)$$

where A is the physical area of the photovoltaic (PV) panel and η is a measure of the system's energy harvesting efficiency. G denotes the Global Horizontal Irradiance (GHI), for which a real-world data set [18] is employed. This data set contains per-minute averages of GHI at NTNU Gløshaugen in Trondheim, Norway, which typically peaks around 800 W m^{-2} for a clear summer day.

The overall energy harvesting efficiency η captures a number of efficiency coefficients η_i , such as PV panel efficiency, accumulator charging circuit efficiency, and factors depending on placement and surroundings.

$$\eta = \prod \eta_i \quad (2)$$

In a typical bridge monitoring scenario it may not be feasible to place the system in direct sunlight. This leads to a reduction in yield, as the PV panel will mostly harvest ambient sunlight. For such a system, a reasonable overall efficiency is $\eta = 10^{-3}$.

2) *Events*: Events of the phenomena to be measured are sporadically generated by the environment, with measurable data series corresponding to these events available to be sampled by the system. The event generation is modelled as a Poisson point process with expected rate λ , such that the interarrival times between events are exponentially distributed. Each event represents a vibration *ring-down*, and is assumed to be endowed with an intrinsic informational value v that the system may estimate. In our model, the VoI of an event is sampled from a standard normal distribution, $v \sim \mathcal{N}(0, 1)$.

3) *Sensor*: The device is modelled as a wireless sensor with an energy accumulator state E_{acc} . The sensor is capable of performing three tasks: sampling, processing, and transmission—each requiring a certain amount of energy, E_S , E_P , and E_T , respectively. The tasks are assumed to take a negligible time to complete. We model the accumulator state as a function of the harvested energy and the sensor's energy consumption,

$$E_{\text{acc}}(t) = \int_0^t P(\tau) - \sum_{i \in \mathcal{T}} E_i \delta(\tau - t_i) d\tau, \quad (3)$$

where $\delta(t)$ is the Dirac delta function, and \mathcal{T} is the set of tasks that the sensor performs. For a task i , t_i denotes the time at which the task was performed and $E_i \in \{E_S, E_P, E_T\}$ the consumed energy.

The environment sporadically offers events to the sensor, which determines how to handle each event by some policy. The sensor may estimate the true informational value v of a given event, with \hat{v}_s and \hat{v}_p being the estimates after sampling (prior to processing) and after processing, respectively. These estimates have the distribution shown in (4), being modelled as a sum of the true value v and a zero-mean, normally distributed estimation error.

$$\hat{v}_s \sim \mathcal{N}(v, \sigma_s^2) \quad \hat{v}_p \sim \mathcal{N}(v, \sigma_p^2) \quad (4)$$

4) *Energy management policy*: An important design consideration is the policy employed by the wireless sensor to carry out its functions. In this paper, we examine what is presumably the most natural and straightforward strategy that could be conceived for this purpose. The policy is *myopic* in nature, and will attempt to sample every event as long as it has sufficient energy to do so. Moreover, the system will process and transmit a sampled event at the earliest opportunity permitted by its energy harvesting yield. The policy is VoI-aware, with a threshold θ that lets it avoid wasting energy on low-value data; an event will be discarded if either \hat{v}_s

```

1: procedure POLICY( $\theta, n$ )
2:   if  $ev \leftarrow \text{TRYSAMPLE}$  and  $ev.value \geq \theta$  then
3:     INSERTSORTED( $B, ev$ )
4:   end if
5:   while  $ev \leftarrow \text{GETNEXT}(B)$  do
6:     if  $ev$  is processed then
7:       if  $success \leftarrow \text{TRYTRANSMIT}(ev)$  then
8:         REMOVE( $B, ev$ )
9:       end if
10:    else
11:      TRYPROCESS( $ev$ )
12:    end if
13:    SORT( $B$ )
14:    FILTER( $B, \theta$ )
15:  end while
16:  TRUNCATE( $B, n$ )
17: end procedure

```

Fig. 2: A simple energy management policy for wireless sensors, with buffer capacity n and VoI threshold θ . B represents the sensor’s long-term memory: a buffer preserved between each run and sorted based on estimates of VoI. The TRY-procedures will attempt the corresponding action and consume energy if successful.

or \hat{v}_p is less than θ . The finite long-term memory of IoT devices is captured by a buffer capacity n , which permits the sensor to store events for which it does not currently have the required energy to process or transmit. When sufficient energy is harvested at a later stage, the sensor will handle the buffered events prioritised by their value estimate. Pseudocode for the policy is provided in Figure 2, parametrised by the value threshold θ and buffer capacity n .

B. Utility metrics

The end goal of most IoT designs is to maximise the *utility* of the system within its energy budget. Quantifying the utility of a system, however, can be challenging due to its application-specific nature. In this paper, the following three metrics are identified to capture a broad array of common prioritisations, which may be weighted appropriately depending on the desired behaviour.

- f_T : Ratio of generated events that are transmitted
- $\bar{\tau}$: Average latency from event generation to transmission
- \bar{v} : Average VoI of transmitted events

Respectively, these metrics can be said to represent the quantity, urgency, and quality provided by a given system.

IV. RESULTS AND DISCUSSION

The simulation framework is configured by a number of parameters; these are listed in Table I, grouped into physical and policy parameters. Unless otherwise stated, the following simulations employ the default parameter values listed in the table. These default values represent a typical wireless sensor, using a piezoelectric vibration sensor, a microcontroller with low-power sleep modes, a wireless protocol such as LoRa, and

TABLE I: Simulation parameters and their default values.

Symbol	Default	Description
G	– W m^{-2}	Global horizontal irradiance
A	2.2 cm^2	Photovoltaic panel area
η	0.001	Overall energy harvesting efficiency
λ	1.0 h^{-1}	Expected event rate
E_S	5 mJ	Sampling energy consumption
E_P	50 mJ	Processing energy consumption
E_T	500 mJ	Transmission energy consumption
σ_S	0.3	Standard error of v after sampling
σ_P	0.1	Standard error of v after processing
θ	$-\infty$	Value-of-information threshold
n	5	Buffer capacity

a small photovoltaic panel. The energy harvesting potential G uses 92 days of data from the summer of 2020, June 1 to August 31, obtained from the real-world data set [18].

A. General behaviour

Figure 3 illustrates the general behaviour of the sensor over two consecutive summer days—one clear and one partially cloudy. The top half of the plot shows the instantaneous energy harvesting potential, as well as the accumulator level. The bottom half shows the events that are generated by the environment, and how the system responds to and handles these events. The energy management policy is greedy and short-sighted, and will therefore use all available energy given that there are a sufficient number of events. This is illustrated by the fact that after sundown, when the energy harvesting yield is zero, the sensor expends almost all accumulated energy by continuing to sample, process, and transmit events. At certain points in the timeline, some energy goes to waste as the system samples or processes data which is later discarded either as a result of the threshold θ , or due to it being displaced by a more valuable packet in the finite buffer of capacity n .

We may categorise the general behaviour of the sensor at various points in time into three distinct *modes* of operation: (1) an energy-starved mode, (2) an energy-constrained mode, and (3) an energy-abundant mode.

1) *Energy-starved mode*: The first of these modes describes a state in which the sensor is severely curbed by lack of energy, and does not have sufficient energy to do anything beyond merely sampling events. If the energy reserves are virtually empty, the sensor may not even be able to sample events. Depending on the environment, this mode is mainly encountered at night when the remaining energy in the accumulator has been exhausted and the sensor may not recharge itself due to the absence of solar energy. We observe this behaviour after the last packet is transmitted in Figure 3, from 43 to 51 hours.

2) *Energy-constrained mode*: The second mode of operation is the energy-constrained mode, wherein the sensor transmits a nonzero number of events, but is forced to delay and prioritise which events to transmit due to a limited amount of available energy. We can observe this energy-constrained

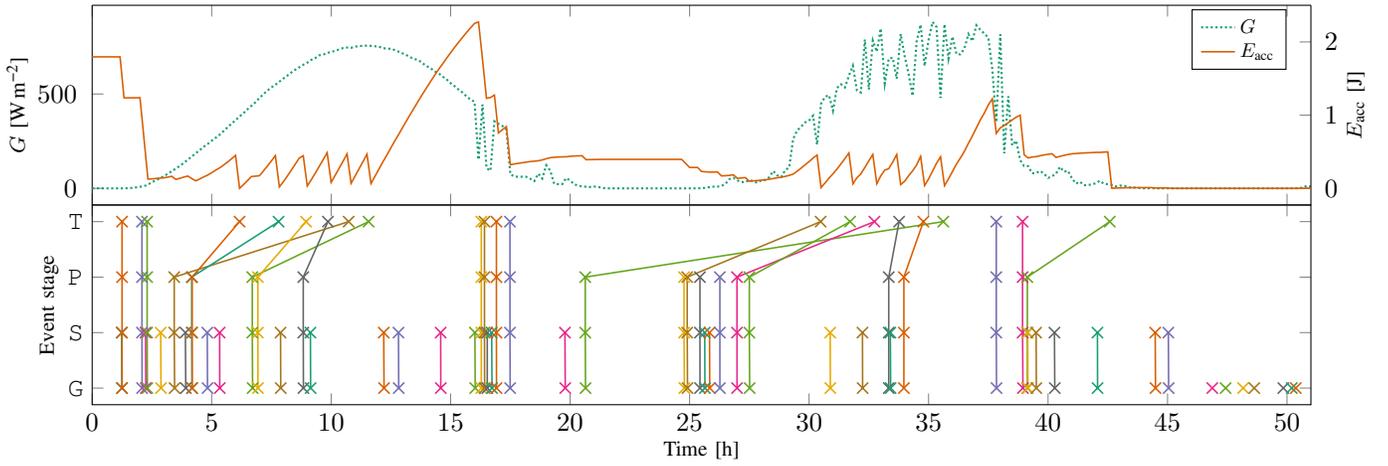


Fig. 3: A two-day excerpt (June 27 to 28, 2020) of a default-parameter simulation showing a typical timeline of the energy management policy with threshold $\theta = 0$. The different event stages are shown as G (generation), S (sampling), P (processing), and T (transmission).

mode from 18 to 37 hours in Figure 3. The prioritisation behaviour is demonstrated by the out-of-order transmissions from 30 to 37 hours. The order is dictated by the VoI estimates of the buffered events; an out-of-order transmission occurs whenever the sensor estimates the value of a later event to be higher than that of an earlier one, and thus prioritises the higher-value event in processing and transmission.

3) *Energy-abundant mode*: Finally, the energy-abundant mode is characterised by the sensor expending less energy than it harvests. In this situation, the sensor will be able to immediately sample, process, and transmit every event which meets its threshold θ . The mode is shown in Figure 3 between 12 and 18 hours. While this last mode of operation is mainly seen around peak daylight hours, it can also be caused by a low number of events occurring.

These modes of operation are transient in nature. However, we may formulate a mathematical description of environments in which they are likely to appear. Defining \bar{G} as the average of G over the horizon of interest, we can introduce the average available energy per event, E_α , as shown in (5).

$$E_\alpha = \frac{A\eta\bar{G}}{\lambda} \quad (5)$$

We may now formulate theoretical bounds describing the long-term behaviour of the sensor.

$$E_S < E_\alpha < E_S + E_P + E_T \quad (6)$$

If satisfied, (6) describes a system that—on average—harvests more energy per event than that consumed by sampling, but less than that required to fully sample, process, and transmit an event. This would result in energy-constrained operation over the long-term. The lower and upper bounds of this inequality designate limits beyond which the sensor would continuously operate in either energy-starved mode or energy-abundant mode, respectively.

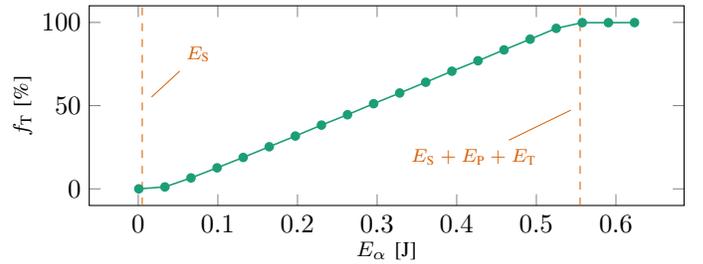


Fig. 4: The transmission ratio f_T shown as a function of E_α , for $\lambda = 1.0 \text{ h}^{-1}$ and increasing A .

Figure 4 demonstrates how the transmission ratio, f_T , is closely tied to E_α and the inequality (6). We observe that if $E_\alpha < E_S$, no events are transmitted ($f_T = 0\%$), while if $E_\alpha > E_S + E_P + E_T$, all events are transmitted ($f_T = 100\%$). Between these two extremes is a region wherein f_T is linearly dependent on E_α . In this scenario, the sensor is long-term operating in the energy-constrained mode, but may alternate between operating in all three modes in the short-term. The default environment given by Table I and shown in Figure 3 is of this nature.

B. Policy parametrisation

The energy management policy is parametrised as (θ, n) , with a value threshold θ and a buffer capacity n . In this section, we will investigate how each of these parameters affects the three main metrics in isolation, for a selection of PV panel sizes.

1) *Value threshold*: Figure 5 shows each metric as a function of θ . In the upper plot, f_T assumes a sigmoid shape for the energy-abundant scenario ($A = 10.0 \text{ cm}^2$). Furthermore, we identify a configuration-specific value θ_{crit} such that all other configurations ultimately enter the energy-abundant mode and conform to this sigmoid if $\theta > \theta_{\text{crit}}$. The sigmoid shape itself is a natural consequence of the normal distribution of v .

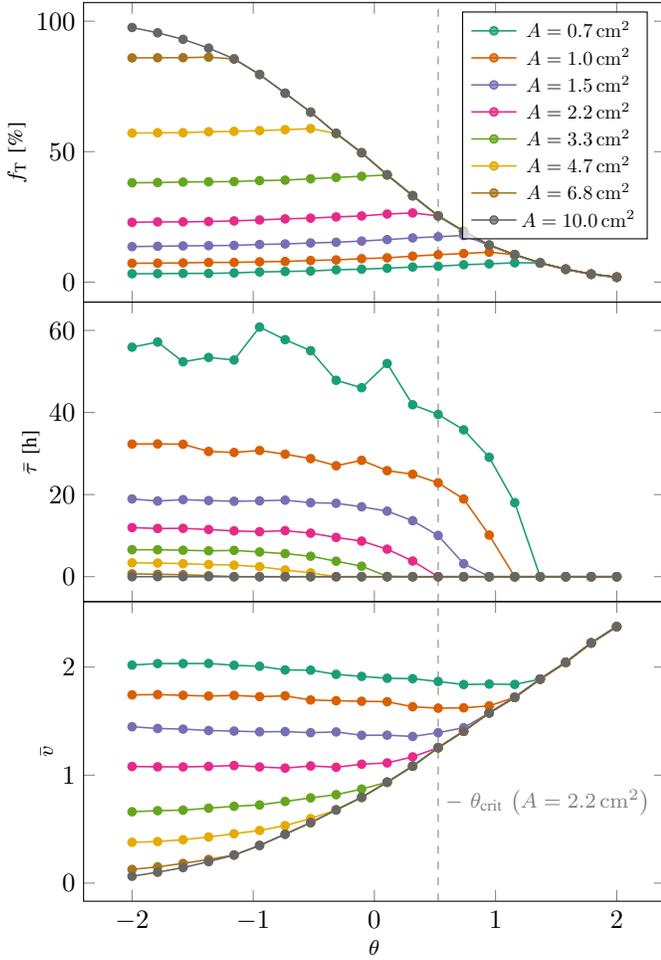


Fig. 5: The transmission ratio (top), average latency (middle), and average VoI (bottom), shown as functions of value threshold, θ , for different PV panel areas A . The critical value θ_{crit} is indicated for the default parametrisation $A = 2.2 \text{ cm}^2$.

Below θ_{crit} , we observe that f_T increases slightly with θ . This is caused by less energy being spent processing low-quality events, since an increasing number of events are discarded right after sampling due to their estimated low value-of-information.

The middle plot shows how the average latency is fairly insensitive to changes in threshold for $\theta \ll \theta_{\text{crit}}$, as the latency is mostly driven by the amount of time it takes for the system to have enough energy to transmit an event. This also explains why the plot is dominated by noise for smaller values of A ; the low f_T makes the system very sensitive to the randomness in the VoI estimates. We recognise the same critical value θ_{crit} for each A : once θ is increased enough, the sensor enters the energy-abundant mode in which $\bar{\tau} = 0$.

We observe in the lower plot that \bar{v} generally increases with θ . A smaller area A leaves the sensor forced to prioritise more aggressively, also increasing \bar{v} . As with the two previous plots, all configurations reach the energy-abundant mode and conform to the same line once $\theta > \theta_{\text{crit}}$. We note that

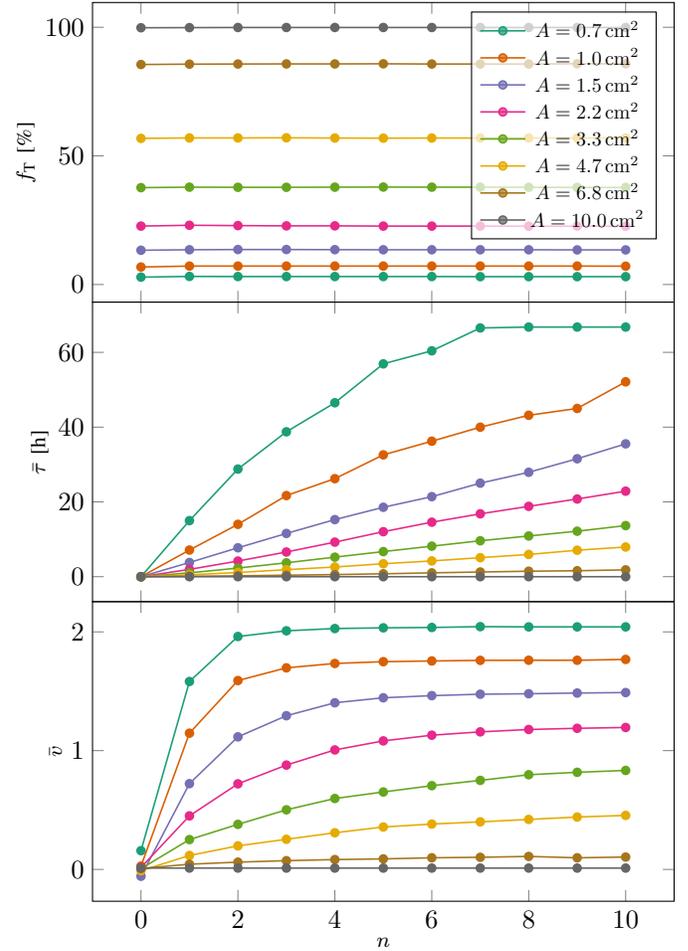


Fig. 6: The transmission ratio (top), average latency (middle), and average VoI (bottom), shown as functions of buffer capacity, n , for different PV panel areas A .

some of the near-horizontal lines slope downwards, an effect especially visible for the most energy-scarce configuration ($A = 0.7 \text{ cm}^2$). This presumably has the same cause as the slight upward slope seen in the upper plot—a slightly higher number of events is transmitted, diluting the \bar{v} score. Another explanation is that the policy may discard some events right after sampling based on quite uncertain estimates of value ($\sigma_s = 0.3$), which could lead to some high-value events being discarded.

2) *Buffer capacity*: Figure 6 illustrates how each metric responds to changes in buffer capacity. While f_T predictably increases with A , the upper plot emphasises the invariance of the transmission ratio with respect to n . This is a natural consequence of the greediness of the policy; the device will at all points in time attempt to transmit as many events as possible, which means that it is either limited by available energy or event rate.

The middle plot demonstrates how the average latency increases with n . Due to event prioritisation, an energy management policy with a large capacity will have a greater

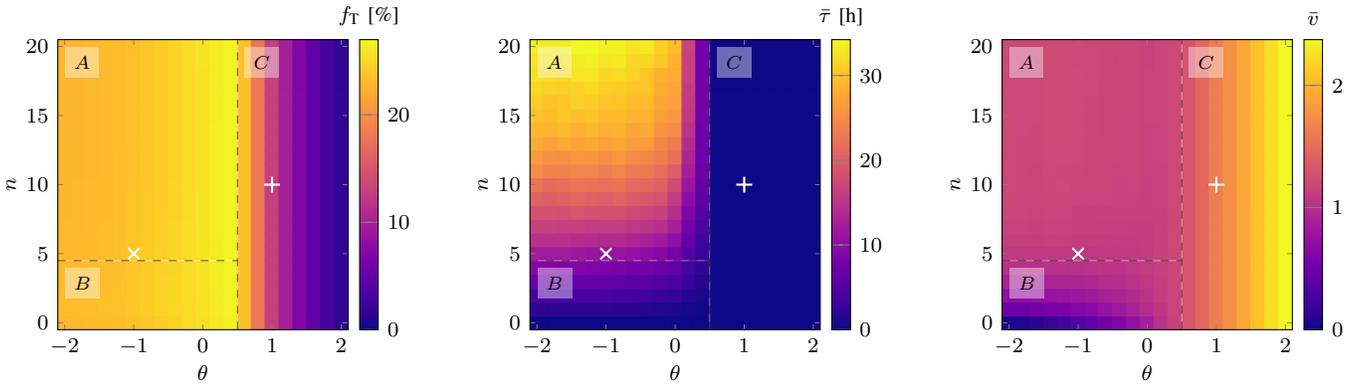


Fig. 7: The heatmaps show how the policy parametrization affects the three key metrics: transmission ratio (left), average latency (middle), and average VoI (right). Two example parametrizations are highlighted.

predisposition towards storing events for longer and sending them out-of-order, generally increasing $\bar{\tau}$. The policy will transmit the best event at the earliest possible opportunity, as such a larger A leads to a lower average latency.

As shown in the lower plot, an improved ability to prioritise events causes \bar{v} to increase asymptotically with n towards a saturation limit. Configurations with $n = 0$ may not prioritise, resulting in $\bar{v} \approx \mathbf{E}[v] = 0$. The parametrizations with a smaller A tend to give a higher \bar{v} since the events transmitted are much fewer and more aggressively prioritised.

Comparing the middle and lower plot in Figure 6, there is a clear trade-off between \bar{v} and $\bar{\tau}$ for a given value of A . In fact, we observe that no buffer capacity is Pareto-dominated by any other— f_T is constant, and both \bar{v} and $\bar{\tau}$ are strictly increasing with n .

C. Parameter analysis

We can draw certain design conclusions from the simulations and analyses shown in this work. From an analytical perspective, E_α can be useful measure to determine if the sensor is operating in—or close to—the energy-abundant mode. Furthermore, assuming the sensor has the ability to buffer and prioritise packets, Figure 5 shows that if high latency can be tolerated, a reduction in the PV panel area may lead to higher-value data—although less of it.

1) *Parameter space regions*: Figure 7 depicts how the three utility metrics are affected across the two-dimensional parameter space of buffer capacity and value threshold. We observe that the parameter space may be partitioned into three regions with different trade-offs and considerations: A , with $\theta < \theta_{\text{crit}} \wedge n \geq 5$; B , with $\theta < \theta_{\text{crit}} \wedge n < 5$; and C , with $\theta > \theta_{\text{crit}}$.

- **Region A**: With f_T and \bar{v} virtually constant, only $\bar{\tau}$ is sensitive to the parametrization. As higher latency is generally undesired, it is apparent that both reducing n and increasing θ results in better performance.
- **Region B**: In this region, f_T is almost constant while a clear trade-off dependent on both n and θ exists between $\bar{\tau}$ and \bar{v} . This trade-off becomes less pronounced with increasing θ .

- **Region C**: The system is permitted to operate in the energy-abundant mode, meaning the design parameters have little effect. The apparent trade-off between \bar{v} and f_T is a result of harsher thresholding—a higher θ will only increase the amount of unspent energy.

2) *Design examples*: Returning to the specific scenario of bridge monitoring, we demonstrate in the following how parameter analyses may be performed, why such analyses can be useful in the design of an energy management policy, and how the specific application affects prioritisation of the utility metrics.

We begin by considering a hypothetical design with $(\theta, n) = (1, 10)$, marked by ‘+’ in Figure 7. The design resides in region C , and is therefore operating in the energy-abundant mode. During testing, the designer might perceive this abundance of energy as a signal to spend more energy processing additional vibration measurements to transmit, and therefore reduces the threshold θ . If done without caution, the system could unexpectedly move to region A with qualitatively different behaviour. For instance, the sensor might suddenly be starved of energy during the night, forcing it to ignore high-value events such as urgent vibration measurements tied to imminent structural failure. A design in region C may also erroneously indicate that both buffering and thresholding are unnecessary. The designer might in this case miss an opportunity to obtain better-or-equal performance with cheaper hardware (for instance a smaller photovoltaic panel), and will therefore unnecessarily negatively distort the cost-benefit analysis. Furthermore, once deployed, this hypothetical design could harvest less energy or experience a higher-than-expected number of events, violating the assumption of energy abundance.

As a different example, we consider a device parametrised by $(\theta, n) = (-1, 5)$, marked by ‘x’ in Figure 7. This system resides on the boundary of region A , with $\bar{\tau}$ the only metric significantly affected by the parametrization. The local insensitivity to θ at this point could suggest that thresholding has no effect and that buffering is the main driver of latency, when—in fact—a threshold increase could be quite beneficial

due to a substantial reduction in latency once θ approaches θ_{crit} . The trivial approach of reducing n to reduce latency, however, will inadvertently also cause a great reduction in \bar{v} once the system enters region B . This would be an undesirable outcome for a bridge monitoring system, since a higher data quality is arguably more important than a lower latency in this context.

D. Limitations

Due to certain simplifying assumptions, the simulations presented in this work have a few limitations. Firstly, only GHI measurements from the summer months were used, thereby disregarding seasonal differences in energy availability. Secondly, other effects could be modelled such as energy leakage from standby power or an non-ideal accumulator. Thirdly, there are quite a few parameters that were not investigated in-depth, such as the task energy requirements (E_S , E_P , E_T) and the uncertainty of the value estimates (σ_s , σ_p). Lastly, the event rate, λ , is assumed in this work to be time-invariant, although it would presumably bear some correlation with time-of-day in our bridge monitoring scenario.

V. CONCLUSIONS

Motivated by a desire to introduce cheap and energy-efficient instrumentation for bridge health monitoring, we have in this paper studied design parameters for energy management of wireless sensors. The challenge of energy management is not specific to this scenario, and we therefore formulated a general model of an energy-harvesting wireless sensor in a resource-scarce environment. Using a special type of vibration event to measure bridge health, we proposed a simple energy management policy for the wireless sensor to decide when to sample, process, and transmit measurements. The policy featured two key parameters: a value-of-information threshold θ , and a buffering capacity n .

We introduced a simulation framework to analyse the utility rendered by the sensor, and its dependence on environmental and design parameters. Three utility metrics were chosen to represent the performance of the system: the ratio of events that were transmitted, f_T , the average latency from sampling to transmission, $\bar{\tau}$, and the average value-of-information of transmitted events, \bar{v} . We then conducted a parameter sensitivity analysis with a detailed evaluation of how internal and external parameters affect each utility metric in the context of bridge monitoring.

With the simple energy management policy chosen for this work, we identified three distinct modes of operation: energy-starved, energy-constrained, and energy-abundant. Resulting in qualitatively different behaviours, these modes highlight how even a simple energy management policy gives rise to complex dynamics. It is therefore of utmost importance to have a solid understanding of the parameter space when designing systems for such resource-scarce environments, as the system may end up operating in an entirely different mode after deployment than was observed during the design phase.

In future work, we aim to demonstrate how the approach introduced here may be applied in real-world scenarios, and substantiate our results with data from systems *in situ*. Furthermore, since the main strength of our method is how it could be employed to dissect the performance of any policy or wireless sensor system working within the confines of the model presented, it would be interesting to conduct a review of different energy management policies and wireless sensor designs, and investigate how their parametrisations affect their efficiency in different environments.

REFERENCES

- [1] R. J. Woodward *et al.*, "BRIME Deliverable D14 Final Report," Tech. Rep., 2001, (accessed on 25th of August, 2021). [Online]. Available: <https://trimis.ec.europa.eu/sites/default/files/project/documents/brimerep.pdf>
- [2] J. Rymcza, "Causes of the Collapse of the Polcevera Viaduct in Genoa, Italy," *App. Sciences*, vol. 11, no. 17, p. 8098, Aug. 2021.
- [3] D. Niyato, E. Hossain, and A. Fallahi, "Sleep and Wakeup Strategies in Solar-Powered Wireless Sensor/Mesh Networks: Performance Analysis and Optimization," *IEEE Trans. Mobile Comput.*, vol. 6, no. 2, pp. 221–236, Feb. 2007.
- [4] V. Sharma, U. Mukherji, V. Joseph, and S. Gupta, "Optimal Energy Management Policies for Energy Harvesting Sensor Nodes," *IEEE Trans. Wireless Commun.*, vol. 9, no. 4, pp. 1326–1336, Apr. 2010.
- [5] C. Moser, L. Thiele, D. Brunelli, and L. Benini, "Adaptive Power Management for Environmentally Powered Systems," *IEEE Trans. Comput.*, vol. 59, no. 4, pp. 478–491, Apr. 2010.
- [6] Y. Gao, W. Cheng, and H. Zhang, "Statistical-QoS Guaranteed Energy Efficiency Optimization for Energy Harvesting Wireless Sensor Networks," *Sensors*, vol. 17, no. 9, p. 1933, Aug. 2017.
- [7] G. Pan, H. Lei, Y. Yuan, and Z. Ding, "Performance Analysis and Optimization for SWIPT Wireless Sensor Networks," *IEEE Trans. Commun.*, vol. 65, no. 5, pp. 2291–2302, May 2017.
- [8] K. Patil, K. D. Turck, and D. Fiems, "A Two-Queue Model for Optimising the Value of Information in Energy-Harvesting Sensor Networks," *Perform. Eval.*, vol. 119, pp. 27–42, Mar. 2018.
- [9] F. A. Kraemer, F. Alawad, and I. M. V. Bosch, "Energy-Accuracy Tradeoff for Efficient Noise Monitoring and Prediction in Working Environments," in *Proc. of the 9th Int. Conf. on the Internet of Things*, Bilbao, Spain, 2019.
- [10] F. Neitzel, B. Resnik, S. Weisbrich, and A. Friedrich, "Vibration Monitoring of Bridges," *Rep. on Geodesy*, vol. 38, no. 4, pp. 393–422, Mar. 2011.
- [11] P. Guillaume, T. D. Troyer, C. Devriendt, and G. D. Sitter, "OMAX—A Combined Experimental-Operational Modal Analysis Approach," in *Proc. of Int. Conf. on Noise and Vibration Eng.*, Leuven, Belgium, 2006, pp. 2985–2996.
- [12] W. Heylen, S. Lammens, and P. Sas, *Modal Analysis Theory and Testing*. Leuven, Belgium: Katholieke Universiteit Leuven, 1997.
- [13] M. Omidalizarandi, R. Herrmann, B. Kargoll, S. Marx, J.-A. Paffenholz, and I. Neumann, "A Validated Robust and Automatic Procedure for Vibration Analysis of Bridge Structures Using MEMS Accelerometers," *J. of Appl. Geodesy*, vol. 14, no. 3, pp. 327–354, Jun. 2020.
- [14] R. Ahmed, B. Buchli, S. Draskovic, L. Sigrist, P. Kumar, and L. Thiele, "Optimal Power Management with Guaranteed Minimum Energy Utilization for Solar Energy Harvesting Systems," *ACM Trans. on Embedded Comput. Syst.*, vol. 18, no. 4, pp. 1–26, Jun. 2019.
- [15] L. Hanschke and C. Renner, "Scheduling Recurring and Dependent Tasks in EH-WSNs," *Sustain. Comput.: Informa. and Syst.*, vol. 27, p. 100409, Sep. 2020.
- [16] S. Draskovic and L. Thiele, "Optimal Power Management for Energy Harvesting Systems with A Backup Power Source," in *10th Mediterranean Conf. on Embedded Comput.*, Budva, Montenegro, 2021, pp. 1–9.
- [17] C. Bisdikian, L. M. Kaplan, and M. B. Srivastava, "On the quality and value of information in sensor networks," *ACM Trans. on Sens. Netw.*, vol. 9, no. 4, pp. 1–26, Jul. 2013.
- [18] MET Norway, "Frost API," 2019, accessed on: Nov. 11, 2021. [Online]. Available: <https://frost.met.no/>