Evaluation of Operation Policies for Energy Harvesting Sensor Nodes with Variable Data Traffic

Leonardo Badia and Giulia Mansutti
University of Padova, Dept. of Information Engineering, via Gradenigo 6B, 35131 Padova, Italy
email: {badia,mansutti}@dei.unipd.it

(Invited Paper)

***Abstract*—We discuss a mathematical characterization of communication devices with energy harvesting capabilities. We consider a terminal powered both by an external source with time-varying energy supply, and a battery with finite storage, and we focus on their interaction. To this end, we make use of existing stochastic models characterizing the operating policies that regulate data transmission. We extend such models by considering data with various level of importance generated by the device with variable rate at the edge of energetic sustainability. We show that under these conditions the usual gap between efficient operating policies and simple greedy transmission becomes even more relevant, thus justifying the need for an efficient management of the energy resource in battery-constrained devices.**

***Index Terms*—Battery management systems; wireless sensor networks; green design; renewable energy sources; stochastic processes.**

I. Introduction

During the latest years, several technological improvements have brought consumer electronics to unprecedented levels of penetration in everyday's life. Smaller and smaller microcircuits and highly advanced transmission systems have caused the development of pervasive (even wearable) and cheap devices capable of sensing and communicating, for both personal exchange of information and tasks such as environmental control or health monitoring [1].

Such an ubiquitous deployment of WSNs must rely on their ability of autonomic and unsupervised operation over time. At the same time, the advantage offered by wireless sensors to avoid the need for cables implies that the nodes are normally battery-powered, which puts serious stress on the battery capacity as a limiting factor [2].

A (partial) solution in this sense can be the exploitation of external power sources through energy harvesting mechanisms for solar, motion, piezo-electric, heat, aeolian, or other renewable sources [3], [4]. The use of these energy inputs in electronic devices is generally advisable to reduce pollution. For battery-powered devices, such as wireless sensors, it is also promising to enable longer operation.

Actually, energy harvesting mechanisms have advanced to the point that future sensors are even thought to be battery-free and just rely on such sources [5]. However, present-day energy harvesting techniques are too erratic to guarantee reliable operation and therefore the use of a battery is unavoidable, especially in critical data sensing. At the same time, mobile batteries may be inadequate to ensure a sufficiently long lifetime of the devices without a recharging mechanism. Thus, we argue that both energy harvesting mechanisms and batteries are useful to provide the required power supply for data sensing, processing, and transmission; however, they need a proper coordination for their usage.

Energetic sustainability can be the bottleneck for the autonomic operation of wireless networks, and it is important to correctly understand the interplay between the energy source and the battery. Especially, even in the presence of energy scavenging, instantaneous usage of battery power must be aware of its current availability [6], [7].

Rechargeable batteries accumulate and release energy when it is more convenient, and are therefore able to filter the often unpredictable and erratic process of energy scavenging. The energy storage element can be seen as a buffer, and the combination of arrival and service processes of data with those of arrival and consumption of energy can be studied via queueing theory and dynamic optimization [8].

Naturally, some difference with classic data queues arise when extending the analysis to energy harvesting. Energy is generated from the harvesting process and is drawn for sensing and communication purposes, whenever needed. Its usage depends on an energy-aware *operation policy*, which can be thought as an algorithm for the management of these energy buffers, depending also on the how important the data are and how much energy is still present in the buffer. Also, while data queue management should usually operate in a stable regime, so as to avoid *data overflow* conditions, the energy buffer should instead mostly avoid *energy outage* situations, i.e., depletion of charge in the battery, which would make it impossible for the sensor to operate. Also *energy overflow* is to be avoided, since it would imply that the harvested battery charge is not fully exploited as it could [7].

We show how the problem of optimizing the operating policy, so as to avoid energy overflow or outage, is not trivial. However, simple policies that exploit reasonable criteria and are easy to implement in wireless sensors, are able to approach (up to 98%) optimal efficiency of the policy. On the other hand, the lack of energy-aware policies that properly harness the underlying energy harvesting process can lead to significant inefficiencies. In our investigated case, an aggressive policy that does not consider at all the harvesting process, and just transmits whenever possible, is shown to have significantly lower efficiency, because it does not exploit the buffering effect offered by the battery. Indeed, its performance is shown to be poor, regardless of the maximum battery capacity. Also, such a policy is shown to be often unable to transmit due to energy outages, that plague the device under heavy data loads.

A more energy-conscious policy that just uses a balanced consumption of the battery (i.e., does not transmit more data than the energy arrival process), is instead asymptotically efficient, i.e., reaches maximum throughput but only for very high values of the maximum battery capacity, and does not achieve full efficiency in realistic cases; thus, depending on the maximum battery capacity, it may lead to significant losses.

Moreover, batteries involve more complex mechanisms than just storing and drawing energy on-demand and without side effects. We discuss how complications may arise due to correlation in the energy generation process, and degradation of the battery [9]. These aspects require a carefully planned operating policy for the battery, and the goal of identifying a low cost implementation for WSNs is a challenging task.

The rest of this paper is organized as follows. In Section II we show a queueing model for devices with energy harvesting capability. Section III reviews possible operating policies for such devices, explaining how energy-awareness can be implemented in practice. In Section IV we show some numerical results, and in Section V we comment on possible directions to extend this study. Finally, we conclude in Section VI.

II. System Model

Consider a slotted-time system, where time slots are taken to be unitary, so that, for $k \in \mathbb{Z}$, we denote time interval $[k, k+1)$ as slot k. Widely employed models [3], [7] to represent an energy harvesting device as a queueing buffer state that energy is stored in the battery in the form of discrete atoms called *energy quanta*, whose absolute value depends on the scenario under consideration. The energy level at time k is denoted by e_k and takes values in the discrete set $\mathcal{E} = \{0, 1, \ldots, e_{\max}\}$, where $e_{\max} \geq 1$ is the battery *capacity*.

At each time slot, the following events may happen: transmission of data with the consumption of q_k energy quanta, and arrival of b_k energy quanta scavenged from the environment. Starting from the initial condition e_0, the evolution of e_k is

$$e_{k+1} = \min\left\{[e_k - q_k]^+ + b_k, e_{\max}\right\}, \quad k \geq 0, \qquad (1)$$

where $[\cdot]^+ \triangleq \max\{\cdot, 0\}$. Since the energy generation is erratic, b_k is random; in this paper, we model it as an independently identically distributed (i.i.d.) variable, taking values in $\mathcal{B} = \{0, 1, \ldots, b_{\max}\}$ with mean $\eta{=}\mathbb{E}[b_k]$. We denote η as the *average harvesting rate*.

Also, q_k is the discrete number of energy quanta describing the energy amount used by the sensor at time k to perform an operation (e.g., transmitting a packet). The *control space* is $\mathcal{Q} = \{0, \ldots, q_{\max}\}$, for some $0 < q_{\max} \leq e_{\max}$, so that $q_k \in \mathcal{Q}, \forall k$. The parameter $q_{\max}$ reflects a physical constraint on the maximum amount of energy that can be drawn from the buffer at any given time.

Given the energy level e_k and the decision q_k, the following two phenomena may occur due to (1). One, indicated as *energy outage*, describes that the energy buffer becomes empty if $q_k > e_k$, since the node runs out of energy before the completion of the executed task. In other words, an energy outage is caused by an attempt to draw more energy from the battery than what is available. The other one, called *energy overflow*, happens whenever $b_k > e_{\max} - [e_k - q_k]^+$, which means that the energy buffer is unable to store all of the harvested energy b_k. This is a consequence of the limited capacity of the energy buffer. Both events pose limitations to the full utilization of the energy scavenging mechanism as a way to provide autonomous operation for the sensor.

At time k, the amount of energy q_k to be drawn from the energy buffer is decided upon by an operating policy μ. Formally, μ is a probability measure on the action space $\mathcal{Q}$, parameterized by the state of the system, which consists first and foremost of the energy level e_k but also includes past events of energy arrivals, overflows, and outages. We will relax this last assumption in the following.

For the sake of simplicity, we limit the analysis to the case in which every arriving packet always consume exactly one quantum of energy for transmission [7]. Thus, q_k can be either 0 or 1, depending on whether the transmitter is active or not. Actually, this assumption only marginally affects the analysis, which can be re-scaled by including different amounts of consumed energy, but the derivation of the system steady-state probabilities becomes much easier. At the same time, we also consider to have L types of packets that may arrive at the transmitter, and therefore different rewards according to the type of packet.

These differences are characterized by considering a set $\mathcal{V}$ of cardinality L, which contains the *values* of the different packets that may arrive at the transmitter. We also assume, similar to [7], that the process that determines the values of the packet is uncorrelated with the arrival of energy quanta.

We define the reward function $G : \mathcal{V} \times \mathcal{E} \mapsto \mathbb{R}^+$ as

$$G(e_k, v_k) = \begin{cases} 0 & e_k = 0 \\ v_k \mu(e_k, v_k) & e_k \geq 1, \end{cases} . \qquad (2)$$

If $e_k = 0$ the reward is 0, which models the inability of the sensor to complete a task in case of energy outage. Otherwise, a reward v_k is accrued with probability $\mu(e_k, v_k)$, i.e., whenever the operating policy dictates that the packet is transmitted, in which case also an energy quantum is spent.

Under this definition of the system model, our goal is to efficiently manage the energy in the device, despite the unpredictability of the source and the limitedness of the battery capacity. If the device did not have these limitations, it would be able to always transmit data whenever needed. Since instead the battery does not sometimes have energy to support the transmission, eventually the achieved throughput will be lower than the ideal case without battery limitations.

Thus, we evaluate the efficiency of the device, considered as the ratio between the throughput achieved and its theoretical maximum, which would be achieved if the device were always able to transmit when needed. We also assume that the battery behavior and the arrival process cannot be manipulated, the only action A_k that can be applied at every time slot k is the choice on whether to transmit or not, depending on the importance of the data and the level of the battery charge.

III. Operating Policies

We define as the *operating policy* of the device, a criterion $\mu(e, v)$ that determines the action (i.e., to transmit a packet or not) given that the importance of the data to transmit has value v and the battery level is e. The pair (e, v) describes the system state. We assume that data, which are not immediately transmitted, must be discarded; therefore, v is a value specific to the present time slot. Also, we allow for the possibility that no data packet arrives in a time slot, in which case $v = 0$.

For a policy μ, we can evaluate its performance through the *expected reward*, computed as the average value of the packets that the device is able to successfully transmit at the steady-state [7]. In other words, we give a quantification of the service that the device is able to provide under policy μ, when time $k \to \infty$. Also, note that for the sake of simplicity we limit our attention to the set of simpler policies where action A_k is just a function of e_k and v_k, i.e., the packet value and the energy level at time k. Thus, we neglect past events of energy arrivals, overflows, and outages. It would actually be possible to define more cumbersome policies where also past history of the state of charge matters. We argue that our simplifying choice is enough to give most of the insights. At the same time, it seems reasonable to assume that existing devices follow similar approaches, because the inherently simple circuitry of wireless sensor nodes does not allow to implement more complex policies.

We assume that $A_k \in \{0, 1\}$ and actions 0 and 1 corresponds to no transmission and transmission, respectively. Thus, we can set $\mu(e, v)$ as the probability that the device will transmit when the system state is (e, v), in which case $A_k = 1$; similarly, $A_k = 0$ with probability $1 - \mu(e, v)$.

Given initial state S_0, the *expected reward* for policy μ should be evaluated through a complex computation given by

$$\bar{G}(\mu, S_0) = \lim_{K\to\infty} \inf \frac{1}{K} \cdot \mathbb{E}\left[\sum_{k=0}^{K-1} A_k V_k \middle| S_0\right] \tag{3}$$

where the expectation is computed on the variables $\{A_k, V_k\}$, and A_k depends on μ, as specified before.

However, according to [7], it is possible to simplify the math, since the optimal policy has a threshold structure, i.e., it is

$$\begin{cases} \mu(e, v) = 1 & v \geq v_{th}(e) \\ \mu(e, v) = 0 & v < v_{th}(e) \end{cases} \tag{4}$$

where $v_{th}(e)$ is a proper threshold value defined depending on the energy level $e \in \mathcal{E}$, $e \neq 0$. In other words, the device shall transmit with probability 1 if the value of the data and the energy level are both high enough.

Naturally, the exact derivation of the optimal threshold may be quite complicated to derive, as discussed in [7]. Here, we are more interested in choosing simple threshold policies, whose definitions follow the theory also derived in that paper, and see how they behave. It is worthwhile noting that, because of these simplifications in the threshold setup, the achieved reward will not be optimal. Still, it is possible to compare the goodness of a policy in quite direct quantitative terms, by measuring the actual average reward achieved. A more efficient policy is in fact able to utilize the same amount of received energy to transmit more valuable data. Also, we investigate the role played by the erratic behavior of the energy source and the data arrival process, which lead to an on-line optimization process that may or may not be robust with respect to the theoretical optimum.

For this comparison, we introduce three different policies: an "aggressive" policy that just involves transmission whenever energy is available in the buffer, a threshold-based "balanced" policy, and another "near-optimal" threshold-based policy inspired by [7], which is shown to get the best performance of the lot, by simply applying just few basic criteria, namely to avoid outage and overflow of the battery as much as possible, all with low-complexity procedure.

The *aggressive* policy is just a greedy strategy that myopically transmits any packet regardless of its value, simply provided that the energy stored in the buffer allows for it. Thus, $\mu(e, v) = 0$ only in the case of an empty buffer. Formally,

$$\mu(e, v) = \chi\{e > 0\} \qquad \text{with } \chi\{\cdot\} \text{ a characteristic function.}$$

Strictly speaking, also this policy is threshold based, but with a trivial threshold $v_{th}(e) = 0$ for all $e \neq 0$. It is immediate to realize that such an aggressive policy increases the occurrence of energy outage events and leads to a frequently empty battery. Thus, upon reception of an important packet, there may be a significant chance that it will not be transmitted.

The *balanced* policy also uses a threshold criterion, where the threshold is chosen so that the average consumption of energy from the battery is identical to the arrival rate η. Thus, the consumption is statistically balanced, and always lead to a (marginally) stable system. However, the actual stability of the system depends on the arrival process of energy; being it erratic, unexpected outages are always possible.

More complex numerical derivations can be used to determine the thresholds for the balanced policy. For the sake of simplicity, in the numerical results of this paper we always considered an energy arrival rate η equal to the sum of the most important kinds of packets, which means that these packets always get to be transmitted, whenever possible, whereas packets of lower importance are discarded. Formally, if v_{y} is the separating value for the packets we want to transmit from those we do not, we obtain

$$\begin{cases} \mu(e, v) = 1 & v \geq v_{\mathrm{y}} \\ \mu(e, v) = 0 & v < v_{\mathrm{y}}. \end{cases} \tag{5}$$

Under this policy, consumption and generation of energy quanta have the same rate, which makes the policy much less aggressive than the previous one. However, also this policy may be disadvantageous when a packet with low importance (under v_{y}) is received when the energy level e is very high, and maybe also an energy quantum simultaneously arrives. The policy dictates to discard this packet, even though it could have been reliably transmitted. Also, it is intuitive to realize that, compared to the aggressive policy, this balanced policy drives the operating point of system much closer to the overflow event, since it is a conservative policy that only transmit when data are deemed worthy of the energy consumption.

Finally, we discuss a *near-optimal*, which is a threshold-based policy that virtually represents a compromise between the aggressiveness and the conservative balance of the previous two policies, and therefore tries to avoid overflow and outage at the same time. It is inspired from the policy proposed in [7], where its theoretical properties have been proved. In particular, its name derives from achieving a performance sufficiently close to optimality; however, a truly optimal policy would possibly require a complex search and dimensioning of the transmission threshold, whereas this simplified policy just uses simple threshold values that are easier to program into computationally constrained small sensors.

Assuming $e_{max} \geq L$, this policy works conservatively, i.e., by transmitting only packets of the most important type L when $e = 1$. This avoids energy outage, unless the packet has very high importance value, so that it is preferable to transmit it anyway. When the energy level is equal to $e_{\max}$, all packets are transmitted. Otherwise, if the packet is of intermediate importance, i.e., $0<v<L$, and the energy level is also strictly between 1 and $e_{\max}$, the packet is transmitted with a threshold approach akin to that used under the balanced policy. For our numerical evaluations, $L = 3$ (actually, there are 4 options, also counting the case of no arrivals discussed later) and y $= 2$, thus the near-optimal policy is

$$\begin{cases} \mu(1, v) = 1 & v \geq v_3 \\ \mu(e, v) = 1 & v \geq v_2, \quad 2 \leq e < e_{max} \\ \mu(e_{max}, v) = 1 & v \geq v_1 \end{cases} \tag{6}$$

IV. Numerical Results

We consider a scenario with i.i.d. individual arrivals of a packet, which need to be immediately transmitted at the device. The model also includes the case where no packet arrives during a certain time slot, which is represented by a packet value $v_0 = 0$, and which happens with probability p_0. We consider that a packet with value v_j arrives with probability p_j, and $L = 3$ types of packets, with different importance, are considered; thus, j can take values between 0 and 3. More specifically, the values of the packets are $v_1 = 1$, $v_2 = 2$, and $v_3 = 4$. For better readability of data, we always consider $p_1 = p_2 = 0.3$, to reduce the number of variables.

Similar to data packets, also energy quanta arrivals are represented by i.i.d. random variables; in particular we assume that the arrival in a given slot is Bernoulli distributed with rate η, meaning that either a packet arrives with probability η or no packet arrives with probability $1 - \eta$. In the numerical evaluations, we always set $\eta = p_2 + p_3$. That is, we consider packets with value 2 or 4 to be "important" data that need to be transmitted whenever possible. Packets with value 1 are instead less important data that could be avoided. However, since the energy source is erratic, it may happen that the device still transmits them. Our condition on η means that transmission of all the data is not energetically sustainable, yet there is (marginally) enough energy to support transmission of only the important data.

The aforementioned three policies have been implemented as Markov chains, taking the system state as comprising the energy level of the battery e and the value of the packet under transmission v_i. In our case, during each time slot at most one energy quantum can arrive and/or be used by the transceiver unit. For this reason, the system can be seen as a Quasi-Birth-and-Death (QBD) process [10].

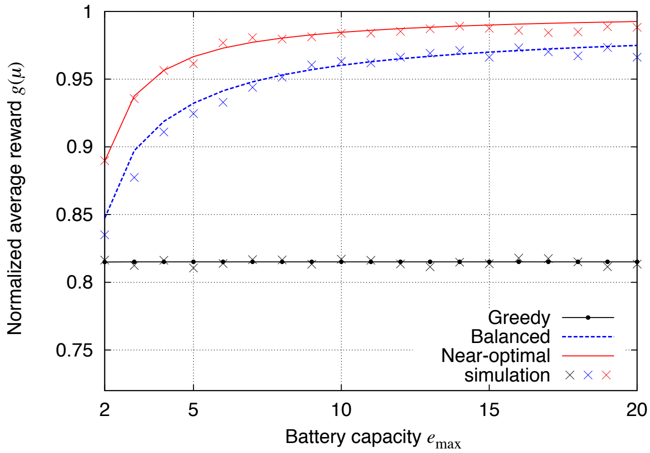

Fig. 1: Efficiency (normalized average reward) $g(\mu)$ as a function of $e_{\max}$. Very frequent data arrival ($p_0 = 0.01$).

It is easy to prove that this system admits steady-state probabilities of being in a given state (e, v), which we denote as $\pi(e, v)$. In other words, at the steady state, the energy in the battery is e and the current value of the packet arrived at the sensor is v with probability $\pi(e, v)$. From the steady-state probabilities, it is immediate to compute the average system reward according to (2). Actually, for better readability of the results, we plot the following quantity:

$$g(\mu) = \sum_{e=1}^{e_{\max}} \frac{\sum_{i=1}^{3} v_i \cdot \pi(e, v_i) \cdot \mu(e, v_i)}{\sum_{i=2}^{3} v_i \cdot p_i} \tag{7}$$

where the numerator is the average reward collected by adopting policy μ, and the denominator represent the theoretical maximum reward achievable by our system, i.e., with sure transmission of only packets with value 2 and 4. This result in a normalized value between 0 and 1, which corresponds to the highest system efficiency.

We first consider a case where data packets arrive almost always, i.e., $p_0 = 0.01$; consequently, there is also high probability of arrival for important data, i.e., $p_3 = 0.39$. The results achieved for the three policies are reported in Fig. 1. Simulation results are also plotted for comparison.

As visible from the plot, the aggressive policy is quite ineffective in achieving high efficiency, and does not benefit from a larger value of $e_{\max}$, which is reasonable, since any energy in the battery is immediately used. The greedy approach of this policy results in a loss of efficiency from the optimum of about 18%.

The balanced and the near-optimal policies are both shown instead to be asymptotically optimal, meaning that they reach 100% efficiency for $e_{\max} \to \infty$. However, the balanced policy reaches this asymptote with a slower rate, and for most values shown in the plot is still far from full efficiency. For $e_{\max} = 5$ the gap is about 6%, for $e_{\max} = 10$ it is still higher than

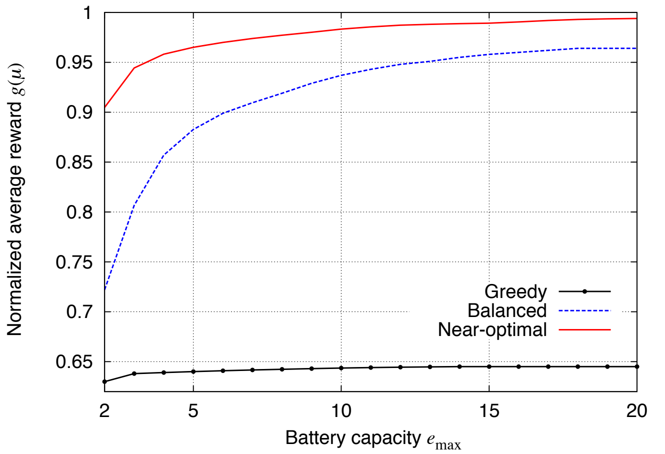

Fig. 2: Efficiency (normalized average reward) as a function of $e_{\max}$. Less frequent data arrival ($p_0 = 0.15$).

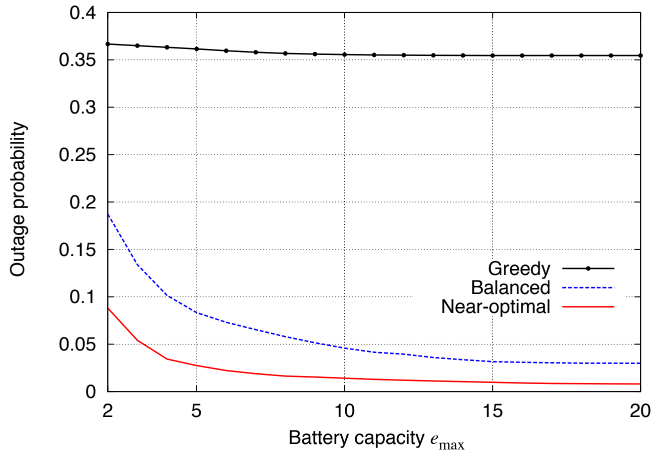

Fig. 3: Outage (empty battery) probability as a function of $e_{\max}$. Less frequent data arrival ($p_0 = 0.15$).

3%. Conversely, the near-optimal policies reaches for the same values of $e_{\max}$ an efficiency of 96.7% and 98.5%, respectively.

We now consider a higher value of $p_0 = 0.15$, and therefore, also a lower value of $p_3 = 0.25$. In the following graphs, simulation points are not shown since they again exhibit excellent agreement as in the previous case. This scenario of data arrival is considered in Figs. 2–3. In particular, Fig. 2 shows the normalized reward, analogous to Fig. 1. The trend is similar, but the gap between simplistic policies and efficient energy management widens. The greedy policy, in particular, is shown to be highly inefficient. Also the balanced policy becomes worse, since, with respect to the previous scenario, highly valuable packets are less frequent (lower value of p_3); the balanced policy acts too conservatively and does not transmit packets with lower value, thus wasting some opportunities to accrue reward. Instead, the low-complexity policy confirms its asymptotically good behavior and satisfactory performance, since the average reward is around 98% already for $e_{\max} \approx 10$.

Fig. 3 shows instead the outage probability, evaluated as the steady-state case where $e = 0$, i.e., $\sum_v \pi(0, v)$. It is shown that adopting the greedy policy leads to very frequent energy outages and, remarkably, the value of $e_{\max}$ has no impact on the reward or the outage, since, as soon as there is energy in the battery, it is used to transmit packets regardless of their value. The balanced and near-complexity policies are more energy-concerned, so they avoid outage more often. However, the balanced policy does so only because it only transmits important packets. In our scenario, where their arrival happens with the same rate as the energy generation, still a non-negligible outage rate is present, and its decrease is very slow even in the presence of a high battery capacity $e_{\max}$. Indeed, the figure shows a floor around 3% for the balanced policy that decreases very slowly. Conversely, the near-optimal policy achieves a very low outage rate, which asymptotically vanishes when the battery capacity is large, but it is already acceptable (below 1.5%) for $e_{\max} \geq 10$.

To sum up, energy awareness is an important element in the design of the operating policy. As discussed in the following section, these improvements, which still involve some percent of the efficiency, may become dramatically higher if other factors, such as correlated energy arrivals and battery inefficiencies and failures, are taken into account.

V. Possible Extensions

The analysis and the results presented in the previous sections considered an i.i.d. arrival process of packets and especially energy quanta. Actually, correlation is naturally present in the process of energy scavenging, e.g., through solar or wind generators. It is even very relevant if we just assume that mobile devices are being recharged at power sockets. Indeed, it has been shown by [13] that correlation of the energy generation process sort of amplifies the aforementioned loss of efficiency of energy-unaware operating policies, and requires to keep the state of the harvester into account. According to this study, a loss of efficiency between 10% and 20%, depending on the process, may be expected if correlation of the energy generation process is not properly taken into account.

Another important aspect is that the entire analysis discussed earlier relies on a perfect estimate of the energy in the battery. Actually, this may not be trivial at all. As argued in [7], the estimation of the charge level for electrochemical batteries is not a trivial task: an estimation bias up to 30% may be present, due for example to temperature differences. Also, online internal estimation procedures of the energy status rely on computationally heavy procedures, and therefore are inherently power-consuming themselves [11]. Thus, gaining precise knowledge of the value e_k may be unreliable or too expensive. Fortunately, according to the results shown in that paper, a coarse quantization of the energy levels may still work with good precision if it is properly accounted for when designing the operating policy of the device. The key rule of a good operating policy, that is, to avoid energy outage and overflow, is still applicable even in the absence of a precise knowledge of the battery level, but just knowing whether the energy in the battery is either low (at risk of outage) or high (at risk of overflow).

Similarly, a degradation effect in the battery storage is also present, which leads to instantaneous leakage and a longterm degradation of the maximum capacity of the battery [9]. In particular, degradation due to deep discharge is a relevant problem for consumer electronics, whose often utilize Li-ion batteries. Typically, the lifetime of such components, quantified as the number of charge and discharge cycles that can be supported, may change even of orders of magnitude [12]. These degradation phenomena can be rather easily included in the model discussed earlier, at the price of some reasonable assumptions on the memory of such a process, and at the price of causing an increase in the number of system states. However, properly designed operating policies, which not only avoid energy outage and overflow, but also account for leakages and try avoiding deep discharges and subsequent battery capacity degradation, can lead to near-optimal usage of the devices and at the same time can significantly prolong the lifetime of its battery [7].

All these elements are required to be kept into account in the development of operating policies for WSNs, also considering the need for low complexity and distributed implementations, since they are supposedly implemented in common consumer electronics. Note that we took a single-node perspective in this paper, but these problems become even more relevant when a network dimension is added. In this sense, theoretical studies can be extremely useful to prove general properties of policies for energy management, which may serve as guidelines for development in real devices.

VI. Conclusions

We investigated the performance of different operation policies for an energy harvesting device which operates under erratic energy sources. We showed that aggressive transmission policies may cause considerable performance losses. Moreover, even a balanced policy, which simply transmits on average with the same amount of the energy arrivals, may be inadequate; a performance loss of this policy is observable, which increases with decreasing buffer capacity. Close-to-optimal performance may be instead obtained via a simple policy which just aims at avoiding energy outages and overflows through a simple threshold setup.

At the same time, we also argued that the general problem of managing the harvested energy in the wireless sensor implies several challenges, depending on the mathematical description of the energy arrival process, the knowledge of the state of charge of the battery and the arrival process, the characterization of battery leakage, and the degradation of the battery lifetime due to too intense discharges.

The goal of future research is to combine all these issues and address them carefully, also including the study of low-cost and low-complexity implementations in relatively cheap and widespread everyday's devices.

References

[1] M. A. Hanson, H. C. Powell Jr., A. T. Barth, K. Ringgenberg, B. H. Calhoun, J. H. Aylor, J. Lach, "Body area sensor networks: challenges and opportunities", *IEEE Computer*, vol. 42, no. 1, pp. 58-65, Jan. 2009.

[2] M. Gatzianas, L. Georgiadis, and L. Tassiulas, "Control of wireless networks with rechargeable batteries," *IEEE Trans. Wireless Commun.*, vol. 9, pp. 581–593, Feb. 2010.

[3] N. Jaggi, K. Kar, and A. Krishnamurthy, "Rechargeable sensor activation under temporally correlated events," *Springer Wireless Networks (WINET)*, vol. 15, pp. 619–635, Jul. 2009.

[4] C. Luo and H. F. Hofmann, "Wideband energy harvesting for piezoelectric devices with linear resonant behaviour," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 58, no. 7, pp. 1294–1301, Jul. 2011.

[5] Y. K. Ramadass and A. P. Chandrakasan, "A battery-less thermoelectric energy harvesting interface circuit with 35 mV startup voltage," *IEEE J. Solid-State Circuits*, vol. 46, no. 1, pp. 333–341, Jan. 2011.

[6] J. A. Paradiso and T. Starner, "Energy scavenging for mobile and wireless electronics," *IEEE Pervasive Computing*, vol. 4, pp. 18–27, Jan. 2005.

[7] N. Michelusi, K. Stamatiou, and M. Zorzi, "On optimal transmission policies for energy harvesting devices," *Proc. ITA*, pp. 249–254, 2012.

[8] D. Bertsekas, *Dynamic programming and optimal control*. Athena Scientific, Belmont, Massachusetts, 2005.

[9] B. Devillers and D. Gunduz, "Energy harvesting communication system with battery constraint and leakage," *Proc. IEEE GLOBECOM Workshops*, Dec. 2011, pp. 383–388.

[10] M. F. Neuts. *Matrix geometric solutions in stochastic models: An algorithmic approach*. Johns Hopkins University Press, Baltimore 1981.

[11] M. Charkhgard, M. Farrokhi, "State-of-charge estimation for lithium-ion batteries using neural networks and EKF," *IEEE Trans. Ind. El.*, vol. 57, no. 12, pp. 4178–4187, Dec. 2010.

[12] D. Linden and T. B. Reddy. *Handbook of Batteries*. 3rd ed., McGraw Hill, 2002.

[13] N. Michelusi, L. Badia, R. Carli, K. Stamatious, and M. Zorzi, "Correlated energy generation and imperfect state-of-charge knowledge in energy harvesting devices," *Proc. IWCMC*, pp. 401–406, 2012.