Energy-Neutral Source-Channel Coding in Energy-Harvesting Wireless Sensors

P. Castiglione*, O. Simeone†, E. Erkip‡, and T. Zemen*
*Forschungszentrum Telekommunikation Wien, Austria
†CWCSPR, ECE Dept, NJIT, New Jersey, USA
‡Dept. of ECE, Polytechnic Inst. of NYU, New York, USA

***Abstract*—This work addresses the problem of energy allocation over source processing and transmission for a single energy-harvesting sensor. An optimal class of policies is identified that simultaneously guarantees a maximal average distortion and the stability of the queue connecting source and channel encoders, whenever this is feasible by any other strategy. This class of policies performs an independent resource optimization for the source and channel encoders. Analog transmission techniques as well as suboptimal strategies that do not use the energy buffer (battery) or use it only for adapting either source or channel encoder energy allocation are also studied.**

I. Introduction

Wireless sensor networks (WSNs) are typically designed under the assumptions that communication resources are limited by the energy available in the battery and that the most significant source of energy expenditure is radio transmission. However, modern cyber-physical systems are expected to operate over a virtually infinite lifetime (see, e.g., [1]). This can only be achieved by overcoming the limitations of battery-powered sensors and allowing the sensors to *harvest* the energy needed for their operation from the environment, e.g., in the form of solar, vibrational or radio energy [2], [3]. The regime of operation in which the system operates in a fully self-powered fashion is referred to as *energy neutral* [4]. Moreover, when sensors are tasked with acquiring complex measures, such as long time sequences of given phenomena of interest, and when transmission takes place over small distances, the energy cost of running the source acquisition system (sensing, sampling, compression) may be comparable with that of radio transmission [5], [6].

Based on the discussion above, in this paper, we address the problem of energy management for a WSN in which sensors are powered via energy harvesting and in which source acquisition and radio transmission have comparable energy requirements. We focus on a system with a single sensor communicating to a single receiver, as shown in Fig. 1, in order to concentrate on the main aspects of the problem. The sensor is equipped with a battery in which the harvested energy is stored. In each time slot, the sensor acquires a time sequence for the phenomenon of interest, which is characterized by a measurement signal-to-noise ratio (SNR) and autocorrelation, and stores the resulting bits, after possible compression, into a data queue. At the same time, it transmits a number of bits from the data queue to the fusion center over a fading channel with an instantaneous channel SNR. Based on the statistics of the energy harvesting process, and based on the current states of the measurement quality, of channel SNR, and of the data queue, the energy management unit must perform energy allocation between source acquisition and data transmission so as to optimally balance competing requirements such as distortion of the reconstruction at the receiver, queue stability and delay. This optimization problem is the main subject of this work.

This work was supported by the Austria Science Fund (FWF) through grant NFN SISE (S106). The Telecommunications Research Center Vienna (FTW) is supported by the Austrian Government and the City of Vienna within the competence center program COMET.

The model at hand is inspired by the work in [5], [6] and [7]. In [7], the energy-harvesting sensor allocates power to data transmission over different channel SNRs, since the bit arrival process is assumed to be given and not subject to optimization. This is unlike our work in which a key problem is that of allocating resources between transmission and source compression in order to guarantee given constraints such as distortion and queue stability. The problem of energy allocation between source compression and transmission was instead first studied in [5], [6], but in power-limited systems with no energy-harvesting capabilities.

The paper is organized as follows. In Sect. II, a model for an energy-harvesting sensor operating over a time-varying channel is provided and the problem set-up is presented. Sec. III contains the main contribution of this work. A class of policies that is able to stabilize the data queue and to satisfy a maximum average distortion constraint as long as this is possible is identified. Numerical results are presented in order to compare this class of policies to several suboptimal strategies as well as to analog transmission. The results are briefly summarized in Sec. IV.

II. System Model

In this section, we introduce the system model, main assumptions and problem definition.

We consider a system in which a single sensor communicates with a single receiver as depicted in Fig. 1. In most of the paper, we assume that the sensor performs separate source and channel coding, as described in the following. A different approach is considered in Sec. III-C.

Time is slotted. The energy $E_k \in \mathbb{R}_+$ harvested in time-slot k is stored in an "energy buffer", also referred to as

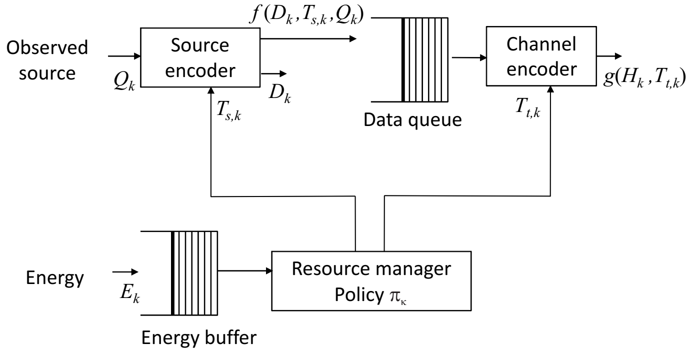

Figure 1. An energy-harvesting sensor composed of a cascade of a source and a channel encoder powered by a resource manager that allocates the energy available in the buffer (e.g., battery or capacitor).

battery, with infinite size. For convenience, the energy E_k is normalized to the number N of channel discrete-time symbols available for communication in each time slot, also referred to as channel uses. The energy arrival E_k is assumed to be a stationary ergodic process. The probability density function (pdf) of E_k is $p_E(e)$. The energy $\tilde{E}_{k+1}$ available for use at slot $k+1$ is the residual energy from the previous slot plus the energy arrival at time-slot $k+1$. This evolves as

$$\tilde{E}_{k+1} = \left[\tilde{E}_k - (T_{s,k} + T_{t,k})\right]^+ + E_{k+1}, \qquad (1)$$

where $T_{s,k}$ and $T_{t,k}$ account for the energy spent in slot k per channel use for source acquisition and data transmission, respectively, as discussed below. Notice that the energy arriving at time slot $k+1$ is immediately available for use in that slot.

The sensor measures M samples of a given source during each slot. The quality of such observation in slot k depends on a parameter $Q_k \in \mathcal{Q}$, which is assumed to be a stationary ergodic process over the time slots k. For instance, the sensor may perform measurements of the phenomenon of interest whose SNR Q_k changes across blocks k due to source movement or environmental factors affecting the measurement quality. The set $\mathcal{Q}$ is assumed to be discrete and finite, and the (stationary) probability mass function (pmf) for Q_k is given by $\Pr(q) = \Pr(Q_k = q)$, for $q \in \mathcal{Q}$. The sensor acquires the source in a lossy fashion. The loss, due to sampling, analog-to-digital conversion and compression, is characterized by distortion $D_k \in \mathbb{R}^+$, as measured with respect to some distortion metric such as the mean square error (MSE).

The number of bits generated by the source encoder at the sensor at slot k is $X_k = f(D_k, T_{s,k}, Q_k)$, where f is a given function of the distortion level D_k, of the energy per channel use allocated to the source encoder $T_{s,k}$ and on the observation state Q_k. The resulting bit stream is buffered in a first-input-first-output (FIFO) data queue with queue length $\tilde{X}_k$. The function $f(D_k, T_{s,k}, Q_k)$ is assumed to be separately continuous convex and non-increasing in D_k and $T_{s,k}$. For simplicity, we will denote such functions also as $f^q(D_k, T_{s,k}) = f(D_k, T_{s,k}, Q_k = q)$. An example for function f will be provided below in Sec. II-A.

The fading channel between sensor and destination is characterized by a process H_k, assumed to be stationary ergodic, where $H_k \in \mathcal{H}$, with set $\mathcal{H}$ being discrete and finite in order to ease the numerical evaluations. We assume a slowly time-variant scenario. The pmf of H_k is given by $\Pr(h) = \Pr(H_k = h)$, for $h \in \mathcal{H}$. The *channel encoder* uses the channel N times per slot, and the transmission requires $T_{t,k}$ energy per channel use. A maximum number $g(H_k, T_{t,k})$ of bits per slot can be delivered successfully to the destination. The channel rate function $g(H_k, T_{t,k})$ is assumed to be continuous, concave, non-decreasing in $T_{t,k}$, and $g(H_k, 0) = 0$. We also use the notation $g^h(T_{t,k}) = g(H_k = h, T_{t,k})$. An example is the Shannon capacity on the complex additive white Gaussian noise (AWGN) channel $g^h(T_{t,k}) = N \times \log(1 + hT_{t,k})$ [8]. We remark that adopting the Shannon capacity implies the use of rate-adaptive schemes with sufficiently long codewords so that the block error probability becomes negligible. For the given function $g(H_k, T_{t,k})$, the channel encoder takes $\min[\tilde{X}_k, g(H_k, T_{t,k})]$ bits from the data buffer, using the selected transmission energy $T_{t,k}$. Note that we do not consider the effects of channel errors nor the costs of channel encoding/decoding and of channel state information feedback, which are beyond the scope of the present paper and subject to future work.

Based on the discussion above, the data queue evolves as

$$\tilde{X}_{k+1} = \left[\tilde{X}_k - g(H_k, T_{t,k})\right]^+ + f(D_k, T_{s,k}, Q_k). \qquad (2)$$

To illustrate the trade-offs involved in the energy allocation between $T_{t,k}$ and $T_{s,k}$, we remark that, by providing more energy $T_{s,k}$ to the source encoder, one is able, for the same distortion level D_k, to reduce the number $f(D_k, T_{s,k}, Q_k)$ of bits to be stored the data buffer. At the same time, less energy $T_{t,k}$ is left for transmission, so that the data buffer is emptied at a lower rate $g(H_k, T_{t,k})$. Viceversa, one could use less energy to the source encoder, thus producing more bits $f(D_k, T_{s,k}, Q_k)$, so that more energy would be available to empty the data buffer.

A. *Rate-Distortion-Energy Trade-Off*

Consider the observation model $R_{k,i} = \sqrt{Q_k}U_{k,i} + Z_{k,i}$, where M samples of the random process $U_{k,i}$, for $i \in \{1, \ldots, M\}$, are measured during the slot k and each measurement $R_{k,i}$ is affected by Additive White Gaussian Noise (AWGN) Z_k with unitary variance. Parameter Q_k represents the observation SNR in slot k. From [5], an approximated and analytically tractable model for $f^q(D_k, T_{s,k})$ is

$$f^q(D_k, T_{s,k}) = \frac{N}{b} \times f_1^q(D_k) \times f_2(T_{s,k}), \qquad (3)$$

where $b = N/M$ is the bandwidth ratio and $f_2(T_{s,k}) = \zeta \times \max[(bT_{s,k}/T_s^{max})^{-1/\eta}, 1]$ models the rate-energy trade-off at the source encoder. The parameter $\zeta > 1$ is related to the efficiency of the encoder, the coefficient $1 \leq \eta \leq 3$ is specified by the the given processor [9] and parameter T_s^{max} upper bounds the energy $T_{s,k}$ that can be used by the source encoder. Function $f_1^q(D_k)$ is a classical rate-distortion function [8]. For the model described in this example, assuming that the source is independent identically distributed (i.i.d) in time with

$U_{k,i} \sim \mathcal{N}(0, d_{max})$, the rate-distortion trade-off is given by

$$f_1^q(D_k) = \left(\log \frac{d_{max} - d_{mmse}}{D_k - d_{mmse}}\right)^+, \quad (4)$$
$$\text{where } d_{mmse} = \left(\frac{1}{d_{max}} + q\right)^{-1}$$

is the estimation minimum MSE (MMSE) for the estimate of $U_{k,i}$ given $R_{k,i}$ [10]. Notice that the distortion D_k is upper bounded by d_{max} and lower bounded by d_{mmse}.

B. Problem Definition

At each time slot k, a *resource manager* must determine the distortion D_k and the energies $T_{s,k}$ and $T_{t,k}$ to be allocated to the source and channel encoder, respectively. The decision is taken according to a policy $\pi := \{\pi_k\}_{k\geq 1}$, where $\pi_k := \{D_k(S^k), T_{s,k}(S^k), T_{t,k}(S^k)\}$ determines parameters $(D_k, T_{s,k}, T_{t,k})$ as a function of the present and past states $S^k = \{S_1, \ldots, S_k\}$ of the system, where the $S_i = \{\tilde{E}_i, \tilde{X}_i, Q_i, H_i\}$ accounts for the state of the available energy $\tilde{E}_i$, for the data buffer $\tilde{X}_i$, for the the source observation state Q_i and the channel state H_i. We define the set of all policies as Π.

III. Stability Under a Distortion Constraint

In this section, we adopt as performance criterion the stability of the data queue connecting source and channel encoders. We also impose the constraint that the policy guarantees the following condition on the long-term distortion:

$$\lim_{n\to\infty} \frac{1}{n}\sum_{k=1}^{n} D_k \leq \bar{D} \quad (5)$$

for a fixed maximum average distortion level $\bar{D}$ tolerated by the system. We define a policy as $\bar{D}$*-feasible* if it guarantees the stability of the data queue connecting source and channel encoders under the average distortion constraint (5). Recall that stability of the data queue holds if the distribution of $\tilde{X}_k$ is asymptotically *stationary* and *proper*, i.e., $\Pr(\tilde{X}_k = \infty) \to 0$ [11].

A. Distortion-Optimal Energy-Neutral Class of Policies

For a given distortion $\bar{D}$, our objective in this section is to identify a class of policies that is able to stabilize the data queue and satisfy the distortion constraint (5) as long as this is possible. We refer to this class of policies as *distortion-optimal energy-neutral*. Notice that this definition generalizes that of "throughput optimal" policies [12] considered in related works such as [7], where only the stability constraint is imposed. By definition, a *distortion-optimal energy-neutral class* of policies $\Pi^{do} \subseteq \Pi$ contains at least one $\bar{D}$-feasible policy. For instance, the set Π of all policies is clearly distortion-optimal energy-neutral. However, this is a rather unsatisfying solution to the problem. In fact, it does not help in any way to identify a $\bar{D}$-feasible policy for a given system setup. Instead, we want to identify a smaller class Π^{do}, which is parametrized in a way that makes it easy to evaluate a $\bar{D}$-feasible policy. The propositions below identify a distortion-optimal energy-neutral class of policies for the separate source and channel encoders model depicted in Fig. 1 and described in Sec. II.

Proposition 1. *For a given distortion $\bar{D}$, a necessary condition for the existence of a $\bar{D}$-feasible policy is the existence of a set of parameters $D^q \geq 0$, $T_s^q \geq 0$ for $q \in \mathcal{Q}$, $T_t^h \geq 0$ for $h \in \mathcal{H}$, and $0 < \alpha < 1$ such that*

$$\sum_q \Pr(q) f^q(D^q, T_s^q) < \sum_h \Pr(h) g^h(T_t^h), \quad (6)$$

$$\sum_q \Pr(q) D^q \leq \bar{D}, \quad (7)$$

$$\sum_q \Pr(q) T_s^q \leq (1-\alpha)\,\mathbb{E}[E_k], \quad (8)$$

and

$$\sum_h \Pr(h) T_t^h \leq \alpha \mathbb{E}[E_k]. \quad (9)$$

Remark 1. Parameters D^q, T_s^q , T_t^h and α, whose existence is necessary for the existence of a $\bar{D}$-feasible policy according to Proposition 1, have a simple interpretation. In particular, T_s^q , D^q can be read as the average energy and distortion that the source encoder selects when the observation state is $Q_k = q$, whereas T_t^h can be seen as the average energy that channel encoder draws from the available energy for transmission when the channel state is $H_k = h$. Moreover, condition (6) is necessary for the stability of the data queue, condition (7) is necessary to satisfy the constraint (5), and conditions (8) and (9) are necessary for energy neutrality. This interpretation will be used below to derive a class of distortion-optimal energy-neutral policies.

Proof: (Sketch) The policy π must be asymptotically stationary for queue (2) to be asymptotically stationary. Under this assumption, the necessary condition for the distribution of $\tilde{X}_k$ to be asymptotically proper is $\mathbb{E}_\pi[f(D_k, T_{s,k}, Q_k)] < \mathbb{E}_\pi[g(H_k, T_{t,k})]$. Defining $D^q = \mathbb{E}_\pi[D_k \mid Q_k = q]$, $T_s^q = \mathbb{E}_\pi[T_{s,k} \mid Q_k = q]$, and $T_t^h = \mathbb{E}_\pi[T_{t,k} \mid H_k = h]$, the condition (6) is then proved by using Jensen inequality at both sides. As for (8) and (9), for a stationary ergodic policy π, the energy harvesting process (1) establishes that $\mathbb{E}_\pi[T_{s,k}] + \mathbb{E}_\pi[T_{t,k}] \leq \mathbb{E}[E_k]$. Given the definitions and the inequality above, (8)-(9) are proved, having set $\alpha = \mathbb{E}_\pi[T_{t,k}]/\mathbb{E}[E_k]$. To conclude, for (7), we observe that the distortion constraint (5) is satisfied. ■

We now look for a distortion-optimal energy-neutral class of policies. To this end, based on Proposition 1, it is enough to exhibit a class of policies such that it contains a $\bar{D}$-feasible policy as long as the necessary conditions (6)-(9) are satisfied for some set of parameters D^q, T_s^q , T_t^h and α. Proposition 1 suggests that it is possible to find $\bar{D}$-feasible policies that select D_k and $T_{s,k}$ based on the observation state Q_k only, whereas the selection of $T_{t,k}$ depends on the channel state H_k only. Based on this consideration, let us define the class of policies Π^{do}

$$\Pi^{do}\begin{cases}D_k = D^q & \text{for } Q_k = q\\ T_{s,k} = \min\left[(1-\alpha)\,\tilde{E}_k - \epsilon, T_s^q\right] & \text{for } Q_k = q\\ T_{t,k} = \min\left[\alpha\tilde{E}_k - \epsilon, T_t^h\right] & \text{for } H_k = h\end{cases}, \tag{10}$$

where $D^q \geq 0$, $T_s^q \geq 0$ for $q \in \mathcal{Q}$, $T_t^h \geq 0$ for $h \in \mathcal{H}$, and $0 < \alpha < 1$ are fixed design parameters.

Proposition 2. *A policy in Π^{do} is $\bar{D}$-feasible if conditions (6) and (7) hold, along with*

$$\sum_q \Pr(q) T_s^q \leq (1-\alpha)\,\mathbb{E}\left[E_k\right] - \epsilon \tag{11}$$

and

$$\sum_q \Pr(q) T_t^h \leq \alpha\mathbb{E}\left[E_k\right] - \epsilon. \tag{12}$$

Remark 2. The sufficient conditions in Proposition 2 for the policies in Π^{do} to be $\bar{D}$-feasible coincide, for $\epsilon \to 0$, with the necessary conditions derived in Proposition 1. Therefore Π^{do} contains a $\bar{D}$-feasible policy any time the necessary conditions of Proposition 1 hold. As discussed above, this implies that the *set* Π^{do} *is a distortion-optimal energy-neutral class*. Moreover, it should be noted that the class Π^{do}, given (10), is parametrized by a small number of parameters and the policies in Π^{do} perform separate resource allocation optimizations for the source and channel encoders. In particular, the energy allocated to the source encoder $T_{s,k}$ only depends on the observation state Q_k, and not on the channel quality H_k, whereas the energy $T_{t,k}$ for the channel encoder only depends on H_k, and not on Q_k. The energy allocation between the two encoders is governed by a single parameter $0 < \alpha < 1$. This entails that, once this parameter is fixed, and thus the energy budget available at the two encoders is fixed, resource allocation at the two encoders can be done separately without loss of optimality.

Proof: (Sketch) For $0 < \alpha < 1$, in the system (1) we must have $\mathbb{E}\left[T_{s,k} + T_{t,k}\right] < \mathbb{E}\left[E_k\right]$, so that the energy harvested is larger than the energy consumed on average and the energy queue is not stable [11, Ch.3] (see also [7] for the same argument). This leads to the asymptotically infinite size of the stored energy, as the buffer capacity is assumed infinite. Therefore we have, $T_{s,k}\left(Q_k = q\right) \to T_s^q$ and $T_{t,k}\left(H_k = h\right) \to T_t^h$ from (10) and $\sum_q \Pr(q) f^q\left(D^q, T_s^q\right) < \sum_h \Pr(h) g^h\left(T_t^h\right)$ becomes the sufficient condition for the stability of the queue $\tilde{X}_k$ [11, Ch.3]. For a stationary ergodic D^q such that $\sum_q \Pr(q) D^q \leq \bar{D}$ the class of policies Π^{do} satisfies the constraint (5). ■

Remark 3. A problem of interest is to find the minimal distortion $\bar{D}$ for which the set of distortion-optimal energy-neutral policies Π^{do} is not empty. In other words, assessing the minimal distortion that can be supported without causing the data queue to be unstable. Given the separate nature of the source and channel energy allocations, it can be seen that one should optimize both terms in (6) separately, once the optimal value for α has been found. In particular, when $g\left(H_k, T_{t,k}\right)$ is the Shannon capacity, the policy $T_{t,k}$ that minimizes $\bar{D}$ is the *water-filling* [8].

B. Suboptimal Classes of Policies

In Sec. III-A, a distortion-optimal energy-neutral class Π^{do} has been identified. This class of policies, as made clear by the proof of Proposition 2 requires infinite energy storage capabilities at the sensor node. Let us instead consider the class of greedy policies Π^{sub1} that do not use the energy buffer but allocates all the energy arrival E_k to source and channel coding according to a fraction $0 \leq \alpha^{q,h} \leq 1$ that depends on both source $Q_k = q$ and channel $H_k = h$ states:

$$\Pi^{sub1}\begin{cases}D_k = D^{q,h}\\ T_{s,k} = \alpha^{q,h} E_k & \text{for } Q_k = q \text{ and } H_k = h\\ T_{t,k} = (1-\alpha^{q,h}) E_k\end{cases}, \tag{13}$$

where the distortion $D^{q,h} \geq 0$ also depends on both source and channel states. Notice that this is unlike the class of distortion-optimal energy-neutral policies (10) in which, as explained in Remark 2, energy allocation is done independently for source (only based on Q_k) and channel decoder (only based on H_k). Here, parameters $\alpha^{q,h}, D^{q,h}$ are selected on the basis of both channel and source states Q_k and H_k to partially compensate for the loss due to the greedy approach. For further reference, we also consider the subclass of policies $\Pi^{sub2} := \Pi^{sub1}|_{\alpha^{q,h}=\alpha \text{ for all } (q,h)}$, for which the power allocation is not adapted to the channel and observation states.

Proposition 3. *Policies in Π^{sub1} are $\bar{D}$-feasible if the following conditions hold:*

$$\sum_q \sum_h \Pr(q)\Pr(h)\mathbb{E}\left[f^q\left(D^{q,h}, \alpha^{q,h} E_k\right)\right] < \tag{14}$$

$$< \sum_q \sum_h \Pr(q)\Pr(h)\mathbb{E}\left[g^h\left((1-\alpha^{q,h}) E_k\right)\right]$$

and

$$\sum_q \sum_h \Pr(q)\Pr(h) D^{q,h} \leq \bar{D}, \tag{15}$$

where the expectation $\mathbb{E}$ in (14) is over the energy harvesting process E_k.

Remark 4. In general, the set of policies Π^{sub1} is not guaranteed to be a distortion-optimal energy-neutral class, since the necessary conditions of Proposition 1 could hold where the sufficient conditions of Proposition 3 do not. This is also confirmed via numerical simulations in Sec. III-D. However, for constant observation and channel states, i.e., $H_k = h_0$ and $Q_k = q_0$ for all k, and for f and g linear in $T_{s,k}$ and $T_{t,k}$, respectively, the class Π^{sub1} is distortion-optimal energy-neutral. In fact, under these assumptions, the sufficient condition (14) becomes $f^{q_0}\left(D^{q_0,h_0}, \alpha^{q_0,h_0}\mathbb{E}\left[E_k\right]\right) < g^{h_0}\left(\left(1-\alpha^{q_0,h_0}\right)\mathbb{E}\left[E_k\right]\right)$. Defining $T_s^q = \alpha^{q_0,h_0}\mathbb{E}\left[E_k\right]$, $T_t^h = \left(1-\alpha^{q_0,h_0}\right)\mathbb{E}\left[E_k\right]$ and $D^q = D^{q_0,h_0}$, conditions (14) and (15) correspond to (6) and (7). Thus, the class of policies Π^{sub1} is distortion-optimal energy-neutral.

Proof: Proof follows from the proofs of Propositions 2. ■

The greedy policies introduced above do not make use of the energy buffer at all, whereas the distortion-optimal energy-neutral class of policies Π^{do} does. For comparison purposes, it is interesting to consider hybrid policies that differ from those in Π^{do} as the energy buffer is used only either for compression or for transmission. The first policies Π^{hyb1} require an energy buffer for the channel encoder only in order to adapt the transmission power to the channel state, i.e., $T_{t,k} = \min[\alpha\tilde{E}_k - \epsilon, T_t^h]$ for $H_k = h$. The energy allocated to the source encoder is instead independent of the observation state, i.e., $T_{s,k} = (1-\alpha)\, E_k$. Viceversa, the second policies Π^{hyb2} are adapted to the observation state instead of the channel state, and require an energy buffer only for the source encoder, i.e., $T_{s,k} = \min[(1-\alpha)\, \tilde{E}_k - \epsilon, T_s^q]$ for $Q_k = q$ and $T_{t,k} = \alpha E_k$.

Proposition 4. *Policies in Π^{hyb1} are $\bar{D}$-feasible if conditions (7) and (12) hold, along with $\sum_q \Pr(q)\mathbb{E}\left[f^q\left(D^q, (1-\alpha)\, E_k\right)\right] < \sum_h \Pr(h) g^h\left(T_t^h\right)$. Similarly, policies Π^{hyb2} are $\bar{D}$-feasible if conditions (7) and (11) hold, along with $\sum_q \Pr(q) f^q\left(D^q, T_s^q\right) < < \sum_h \Pr(h)\mathbb{E}\left[g^h\left(\alpha E_k\right)\right]$.*

Proof: Proof follows from the proofs of Propositions 2. ■

C. Analog Transmission

In this section, we consider for performance comparison an alternative class of strategies in which the sampled source is transmitted directly via analog modulation (see, e.g., [13]). In other words, a block of source samples is scaled and transmitted in one slot, so as to consume $T_{t,k}$ transmission energy per channel use during slot k. Energy $T_{t,k}$ is selected as $T_{t,k} = \min[\tilde{E}_k - \epsilon, T_t^{q,h}]$ for $Q_k = q$ and $H_k = h$ for given parameters $T_t^{q,h} \geq 0$, so that it depends on the current source and channel states. If the bandwidth ratio $b = N/M$ is larger than one, i.e., there are more channel uses than source samples, the extra $N - M$ source samples are unused. Instead, if $b < 1$, then a fraction $1 - b$ of source samples is not transmitted. Notice that this class of strategies does not fall in the category depicted in Fig. 1 and discussed above, since it does not have separate encoders.

We assume the observation model of Example 1 with an i.i.d. source $U_{k,i} \sim \mathcal{N}(0, d_{max})$ and an AWGN channel with SNR H_k. For a bandwidth ratio of b, it is not difficult to obtain that the MMSE at the receiver is given by

$$d_{mmse}\left(T_{t,k}, Q_k, H_k\right) = \left(\frac{bT_{t,k}Q_kH_k}{bT_{t,k}Q_k + Q_k + 1} + \frac{1}{d_{max}}\right)^{-1} \text{ if } b \geq 1, \text{ otherwise}$$

$$b \times \left(\frac{T_{t,k}Q_kH_k}{T_{t,k}Q_k + Q_k + 1} + \frac{1}{d_{max}}\right)^{-1} + (1-b)\, d_{max} \qquad (16)$$

Notice that, for fairness, the average energy used for the transmission of one sample is $bT_{t,k}$ if $b \geq 1$. Also, notice that if $b < 1$, the maximum distortion d_{max} is accrued on the fraction $(1-b)$ of samples that are not transmitted. For simplicity, we assume that analog transmission has negligible power spent for source acquisition, i.e., $T_{s,k} = 0$, though this is not entirely correct given that even in this case there is a need for sensing, sampling and analog-to-digital conversions. Nonetheless, these power consumption terms are also neglected in Examples 1 and 2. Under this assumption, we have the following.

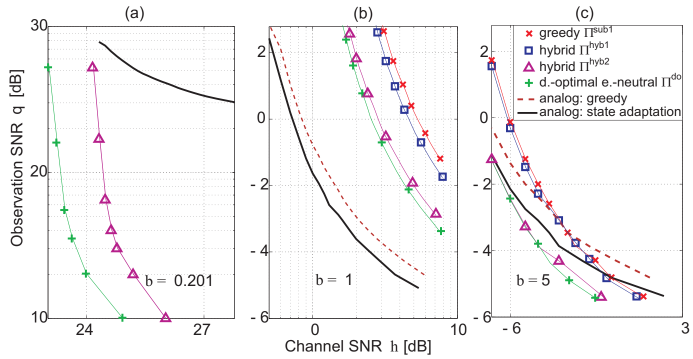

Figure 2. Achievable regions (regions above the curves) for the digital policies (lines with markers) and for the analog policies (only lines): (a) bandwidth ratio $b = 0.201$; (b) bandwidth ratio $b = 1$; (c) bandwidth ratio $b = 5$ ($E_k \sim \mathcal{U}(0, 2)$, with mean 1 Joule/channel use; $\bar{D} = 0.8$; compression model (3), with $T_s^{max} = 1$ Joule/source sample, $\zeta = 1$, $\eta = 1.5$, maximum distortion d_{max}=1).

Proposition 5. *Analog transmission satisfies the distortion constraint (5) if the following conditions are satisfied:*

$$\sum_q \sum_h \Pr(q)\Pr(h) d_{mmse}\left(T_t^{q,h}, Q = q, H = h\right) \leq \bar{D}, \qquad (17)$$

and

$$\sum_q \sum_h \Pr(q)\Pr(h) T_t^{q,h} \leq \mathbb{E}\left[E_k\right]. \qquad (18)$$

Proof: Follows similar to Proposition 2. ■

We can also consider a greedy policy $T_{t,k} = E_k$, for which the power allocation is not adapted to the channel and observation states and energy storage is not required.

D. Numerical Results

In this section we compare numerically the performance of the optimal and suboptimal source-channel coding policies presented so far, along with analog transmission strategies.

Consider first a scenario where the observation and channel states are constant, i.e., $Q_k = q$ and $H_k = h$ for all k. The energy arrival E_k has mean 1 Joule/channel use and uniform pdf between 0 and 2 Joule/channel use. We consider model (3) with $T_s^{max} = 1$ Joule/source sample, efficiency parameters $\zeta = 1$ and $\eta = 1.5$ for the source encoder, and the complex AWGN channel Shannon capacity $g^h\left(T_{t,k}\right) = N \times \log(1 + hT_{t,k})$. In Fig. 2, we identify the values of source and channel SNRs (q, h) for which different policies are able to stabilize the data queue and guarantee average distortion $\bar{D} = 0.8$. We refer to these regions as "achievable regions". Achievable regions are given in Fig. 2 by the area above the corresponding lines. We use standard tools of convex optimization for their numerical evaluation.

In Fig. 2-b, we can further observe that the achievable regions of the distortion-optimal energy-neutral class (10) are significantly larger than those of the greedy policies (13) due to possibility to store energy and thus allocate resources more effectively. Moreover, by considering also the hybrid policies Π^{hyb1} and Π^{hyb2} , we can see that most of the gains are obtained, in this example, by exploiting the energy buffer in order to allocate energy over time to the source encoder, whereas the gains accrued by using the battery for data transmission are less significant. This is observed by noticing that the achievable region of the class Π^{do} is close to that obtained by hybrid policies Π^{hyb2}, but much larger than that obtained by hybrid class of policies Π^{hyb1}. The relative comparison between the two hybrid policies, and thus between the use of the battery for source or channel encoding, depends on the functions $f\left(D_k, T_{s,k}, Q_k\right)$ and $g\left(H_k, T_{t,k}\right)$. For instance, setting a lower T_s^{max} would change the presented results by penalizing more the strategies that are not using the energy buffer for channel transmission.

We now consider the performance of analog transmission. As it is well known from rate-distortion theory [10], for bandwidth ratio $b = 1$ and the given (Gaussian) source and channel models, analog transmission is rate-distortion optimal. Separate source-channel coding is also optimal (for any $b > 0$) if compression is assumed not to consume any energy. Here, instead, the achievable region of analog transmission is expected to be larger than that of strategies that employ separate source-channel coding, as source encoding energy costs are taken into account in the given model (3). This is confirmed by Fig. 2-b. However, for sufficiently larger or smaller bandwidth ratios, the extra energy spent for compression is not enough to overcome the rate-distortion gains attained by separate source-channel coding versus analog transmission. This is apparent from Fig. 2-a and Fig. 2-c where the analog transmission is outperformed.

IV. Conclusions

We studied energy management for a system consisting of a single sensor whose task is that of reporting the measure of a phenomenon to a receiver. The main problem is that of allocating energy between the source and the channel encoders based on the current amount of available energy, state of the data queue, quality of the measurement and of the wireless channel. We looked for a distortion-optimal energy-neutral subset of all policies, that contains at least one policy able to stabilize the data queue and to satisfy a maximum average distortion constraint. We found that optimal policies according to this criterion operate a separate energy allocation of source and channel encoder. Overall, our results, which also include further comparisons with a number of suboptimal policies, shed light on the challenges and design issues that characterize modern cyber-physical systems.

Acknowledgments

The authors would like to thank Prof. Petar Popovski of Aalborg University for suggesting the analysis of the analog transmission schemes.

References

[1] "It's a smart world," *The Economist*, Nov. 6 2010.
[2] J. A. Paradiso and T. Starner, "Energy scavenging for mobile and wireless electronics," *IEEE Pervasive Computing*, vol. 4, pp. 18–27, Jan.-Mar. 2005.
[3] R. Rajkumar, L. Insup, S. Lui, and J. Stankovic, "Cyber-physical systems: The next computing revolution," in *Proc. ACM/IEEE Design Automation Conf. (DAC)*, June 13-18 2010.
[4] A. Kansal, J. Hsu, S. Zahedi, and M. B. Srivastava, "Power management in energy harvesting sensor networks," *ACM Trans. Embed. Comput. Syst.*, vol. 6, Sep. 2007.
[5] Z. He and D. Wu, "Resource allocation and performance analysis of wireless video sensors," *IEEE Trans. on Circ. and Syst. for Video Technol.*, vol. 16, pp. 590 – 599, May 2006.
[6] X. Lu, E. Erkip, Y. Wang, and D. Goodman, "Power efficient multimedia communication over wireless channels," *IEEE J. on Sel. Areas in Commun.*, vol. 21, pp. 1738 – 1751, Dec. 2003.
[7] V. Sharma, U. Mukherji, V. Joseph, and S. Gupta, "Optimal energy management policies for energy harvesting sensor nodes," *IEEE Trans. Wireless Commun.*, vol. 9, pp. 1326–1336, Apr. 2010.
[8] T. M. Cover and J. A. Thomas, *Elements of information theory*. Wiley.
[9] T. Burd and R. Broderson, "Processor design for portable systems," *J. VLSI Signal Process.*, vol. 13, pp. 203–222, Aug. 1996.
[10] T. Berger, *Rate distortion theory*. Prentice Hall, 1971.
[11] A. A. Borovkov, *Stochastic processes in queueing theory*. Springer-Verlag, 1976.
[12] L. Tassiulas and A. Ephremides, "Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks," *IEEE Trans. Autom. Control*, vol. 37, pp. 1936–1948, Dec. 1992.
[13] B. R. M. Gastpar and M. Vetterli, "To code, or not to code: lossy source–channel communication revisited," *IEEE Trans. Inf. Theory*, vol. 49, pp. 1147–1158, May 2003.