Fundamental Limits of Energy Harvesting Communications

Omur Ozel, Kaya Tutuncuoglu, Sennur Ulukus, and Aylin Yener

Omur Ozel and Sennur Ulukus are with the University of Maryland.

Kaya Tutuncuoglu and Aylin Yener are with The Pennsylvania State University.

This work was supported by NSF Grants CNS 09-64632 / CNS 09-64364, and CCF 14-22111 / CCF 14-22347.

ABSTRACT

Wireless networks composed of energy harvesting devices will introduce several transformative changes in wireless networking as we know it: energy self-sufficient, energy self-sustaining, perpetual operation; reduced use of conventional energy and accompanying carbon footprint; untethered mobility; and an ability to deploy wireless networks in hard-to-reach places such as remote rural areas, within structures, and within the human body.

Energy harvesting brings new dimensions to the wireless communication problem in the form of intermittency and randomness of available energy, which necessitates a fresh look at wireless communication protocols at the physical, medium access, and networking layers. Scheduling and optimization aspects of energy harvesting communications in the medium access and networking layers have been relatively well-understood and surveyed in the recent paper [1]. This branch of literature takes a physical layer rate-power relationship that is valid in energy harvesting conditions under large-enough batteries and long-enough durations between energy harvests so that information-theoretic asymptotes are achieved, and optimizes the transmit power over time in order to maximize the throughput.

Another branch of recent literature aims to understand the fundamental capacity limits, i.e. information-theoretic capacities, of energy harvesting links under smaller scale dynamics, considering energy harvests at the channel use level. This branch necessitates a deeper look at the coding and transmission schemes in the physical layer, and ultimately aims to develop an information theory of energy harvesting communications, akin to Shannon's development of an information theory for average power constrained communications.

In this introductory article, we survey recent results in this branch and point to open problems that could be of interest to a broad set of researchers in the fields of communication theory, information theory, signal processing, and networking. In particular, we review capacities of energy harvesting links with infinite-sized, finite-sized, and no batteries at the transmitter.

INTRODUCTION

Energy harvesting devices offer several significant advantages over conventional grid-powered and non-rechargeable battery-powered devices [2]. These advantages include energy self-sufficient, energy self-sustaining operation with lifetimes limited only by the lifetimes of their hardware. These devices utilize energy harvested from alternative natural resources such as solar, vibrational, electromagnetic, thermoelectric, etc. As such, their widespread adoption will reduce the use of conventional energy and the accompanying carbon footprint, and benefit the environment. In addition, these devices do not require conventional recharging, enabling untethered mobility, and deployment of these devices in hard-to-reach places such as remote rural areas, within structures, and within the human body.

The circuits and devices side of engineering has been contributing to the development of energy harvesting devices for decades. However, on the communications, networking, and systems side of engineering, the focus has been on energy-conscious communication system design, in the form of optimum average-power constrained communications, and energy-efficient networking. Such approaches have aimed to minimize total energy usage for a fixed performance, or maximize the performance for a fixed total energy availability. Only recently, communications subject to explicit energy harvesting conditions has garnered attention. In such systems, energy needed for data transmission becomes available (arrives at the transmitter) intermittently and in random amounts, due to the randomness in energy harvesting processes and sources. Then the goal becomes not only to minimize the overall energy consumption, but to maintain reliable communication under random and intermittent energy arrivals. Energy is an essential component of communication: in order to create a desired physical change at the receiving end of a channel, the transmitter needs to use (send) energy. In an energy harvesting system, there is uncertainty in this very basic component of communication, and the communication mechanisms need to be designed by explicitly accounting for these energy harvesting constraints.

Such explicit energy harvesting constraints have been introduced in the context of data scheduling and transmit power optimization in [3–6]; see also [1] for a review of recent results in this branch of literature. This branch of literature utilizes a concave rate-power relationship, such as the logarithmic Shannon capacity formula for the Gaussian channel, and optimizes the transmit power over time, subject to energy harvesting constraints, to maximize the throughput. Such an approach is valid under large-enough battery sizes and long-enough durations between energy harvests, which are valid abstractions in the medium access and networking layers. Such approaches are also valid approximations in the physical layer when the communication modulation and coding scheme is fixed to a potentially sub-optimum scheme. A departure from this approach and an important research direction is to determine the ultimate information-carrying capacity of an energy harvesting link under intermittent and random energy arrivals. While the capacity of an average power-constrained Gaussian channel is well known, what is the capacity of an energy harvesting link that harvests energy in the amount of E_i in channel use i with an average recharge rate of $E[E_i]$?

This question necessitates a deeper look at the explicit coding and transmission schemes in the physical layer. It also necessitates consideration of explicit interactions between energy arrivals, energy storage in the battery, and energy used by the codebook, for data transmission. In particular, the battery can be viewed as an energy queue where the incoming energy is saved for future use. The codebook, which transmits data, can be viewed as a server that serves energy out of this energy queue. On the other hand, energy arrives over time and the amount of available energy in the energy queue dictates the set of feasible codebooks. This introduces interactions between quantities that have been traditionally thought of separately, i.e. energy (which is physical) and data (which is cyber) interact explicitly at every time instant, as opposed to in some average sense as in the classical literature.

In an energy harvesting system, every transmitted symbol is instantaneously constrained by the energy available in the battery. This constraint, which is imposed at every channel use, is different than classical average-power and constant-amplitude constraints, and as such, results in a new form of channel input constraint. From an information-theoretic point of view, the amount of energy available in the battery may be viewed as a *state*, which is naturally causally known by the transmitter, but unknown by the receiver. Even when energy arrivals are independent and identically distributed (i.i.d.), code symbols cannot be generated independently due to the presence of a finite-sized battery and highly time-correlated nature of the battery state. In addition, the transmitter's own actions (past transmitted symbols) affect the future of the state. The goal of this article is to review this branch of literature at an introductory level and present open problems. In particular, we will see that the capacity of an energy harvesting link heavily depends on the size of the battery. We will review results for the cases of infinite-sized batteries, zero-sized batteries, and finite-sized batteries. We start with a brief summary of the results and developed insights in the scheduling and optimization side of the literature.

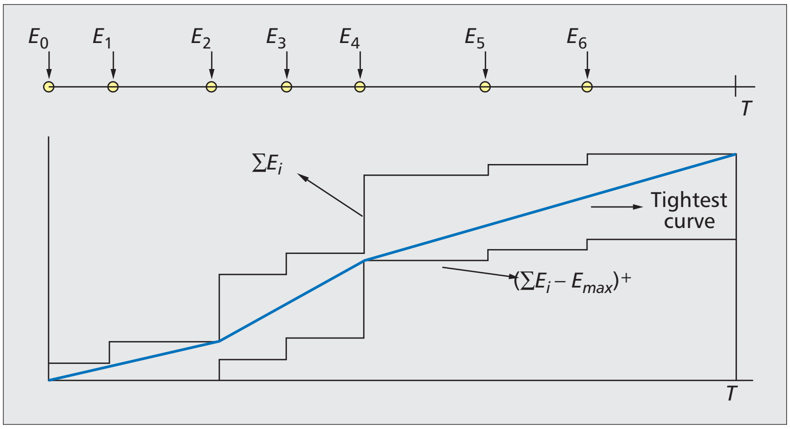

Figure 1. Optimal transmission policy is the *tightest curve* (or *shortest string*) in the *energy feasibility tunnel*.

Data Scheduling and Optimization in Energy Harvesting Communications

Data scheduling in energy harvesting communications is studied under two main categories: offline scheduling and online scheduling. Offline refers to the availability of knowledge of events, such as energy arrivals and channel fade level changes, prior to the start of data scheduling; online refers to the availability of this knowledge only causally over time, but not a priori. In this section we summarize offline scheduling; a more thorough review may be found in [1].

A key assumption in data scheduling literature [3–6] is that a rate-power relation, $r(p)$, determines the rate achieved when the transmission power is chosen as p; a typical example of $r(p)$ is the Shannon capacity formula for the Gaussian channel. Note that the Shannon capacity formula is a monotonically increasing concave function of p, and the development in this literature is valid for all monotonically increasing concave rate-power relationships. Energy arrives from an exogenous energy source over time and is saved in the battery before being utilized for data transmission. Since energy cannot be utilized before it arrives (is harvested), the transmission power has to obey the *energy causality* constraint [3], which states that the energy utilized until a time instant must be less than or equal to the energy harvested by that time instant. In addition, since the battery has a finite-size $E_{\max}$, energy may overflow if the battery does not have a sufficiently large space to accommodate it. Under offline knowledge of energy arrivals, the transmitter has to determine the transmit powers so that all incoming energies are accommodated in the battery, i.e. no energy is wasted; this constraint is called the *no-energy-overflow* constraint [4, 5].

The work in [3] and [4] determined the optimal transmission schemes for an energy harvest-

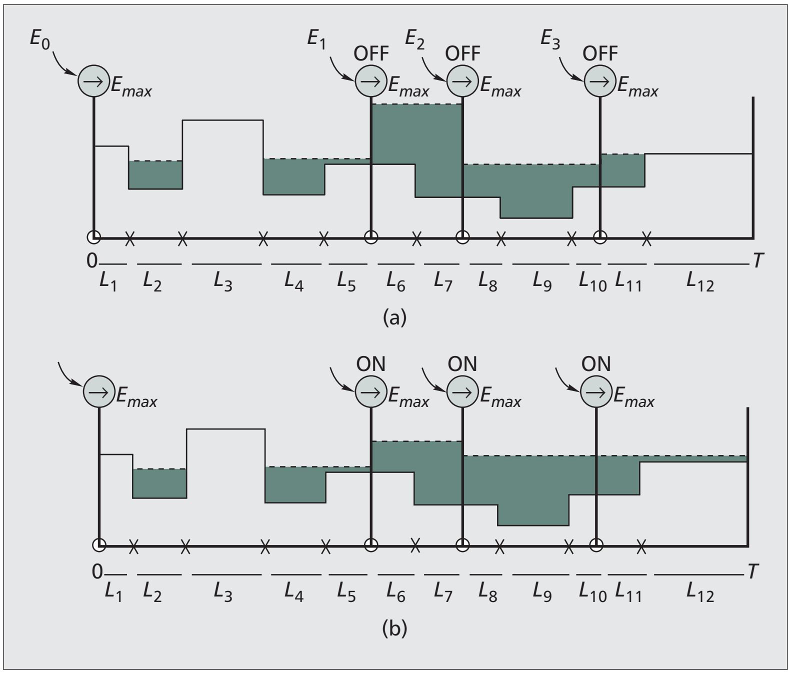

Figure 2. Directional water-filling algorithm: a) first, allocate arriving energy until next energy arrival; and b) then, allow energy to flow right until a balance is reached.

ing transmitter in a non-fading channel with infinite and finite batteries, respectively. The solution has a *shortest string* (or *tightest curve*) structure through an *energy feasibility tunnel*, as shown in Fig. 1, because the length of a curve is a convex function of its slope. The minimization of this convex function may be viewed as maximization of its negative, and the solution of the throughput maximization problem is invariant to the form of the objective function, so long as it is concave and monotone increasing. The upper curve that determines the energy feasibility tunnel represents the cumulative energy arrivals: the energy expenditure of any transmission policy should be below this curve due to *energy causality*. The lower curve in the energy feasibility tunnel is a vertically shifted version of the cumulative energy arrival curve by exactly $E_{\max}$, which is the size of the battery, since at most $E_{\max}$ units of energy can be saved in the battery. The energy expenditure of any transmission policy should be above this curve due to *no-energy-overflow*. Because of the concavity of the rate-power function, keeping the power as constant as possible subject to energy feasibility constraints is optimal. Therefore, the optimal transmit power policy has the *shortest string* structure: the power is kept constant to the extent possible subject to energy causality and no-energy-overflow conditions. The optimal transmission power sequence is the sequence of slopes of the shortest string that lies within the energy feasibility tunnel.

In [5] the solution of [3, 4] is extended to the case of fading channels. The solution is found by a *directional water-filling algorithm*, depicted in Fig. 2. This algorithm is based on a view of energy as water. The incoming energy (water) is allocated (poured) into the time interval till the next energy arrival, as shown in Fig. 2a. Then the water is allowed to flow through the taps from left to right (from the past to the future) due to energy causality (Fig. 2b). This signifies that energy can be saved and used in the future, but energy that has not arrived yet cannot be used. The extent of this energy flow is also limited by the battery limit $E_{\max}$: the amount of water that can flow through each tap is $E_{\max}$. The flow continues till the water levels are balanced. This is a generalization of the classical water-filling algorithm to the case of energy harvesting setting, which introduces a directionality (due to energy causality) and maximum-flow (due to no-energy-overflow) constraints. Note that the slopes in the tightest string interpretation for the static channel correspond to the water levels. See also [6] for a staircase water-filling algorithm.

The literature on data scheduling in energy harvesting communications, such as [3–6] and those that follow them, assumes a monotonically increasing concave rate-power relation $r(p)$ and optimizes the sequence of transmit powers subject to energy causality and no-energy-overflow constraints. However, exact capacities and the resulting rate-power relationships (if any) are not known in energy harvesting settings with dynamics occurring at the channel use level. In fact, we will see that the rate-power relationship highly depends on the size of the available battery $E_{\max}$. Starting in the next section, our goal will be to determine this relationship when possible. We will start by describing the detailed system model at the channel use level.

System Model

As shown in Fig. 3, the transmitter wishes to send a message W in n channel uses so that the receiver reliably decodes the message. $E_{\max}$ is the battery energy storage limit, i.e. the size of the battery. Channel input and output symbols may be discrete or continuous. The exogenous energy source supplies E_i units of energy in the ith channel use, where E_i is an i.i.d. sequence with an average value $E[E_i]$, which we denote as P_{avg}. When channel input X_i is transmitted in the ith channel use, the receiver gets Y_i. The underlying physical communication channel is memoryless.

S_i denotes the energy available in the battery at channel use i. The transmitter observes the available battery energy S_i and transmits a symbol X_i. At each channel use, the transmitter both harvests energy and transmits a symbol. The order of harvesting and transmission in a channel use could be in two different methods. The first method is *transmit first* where the energy of the channel symbol is constrained by the battery energy in that channel use. After sending the symbol, the transmitter harvests energy. Then incoming energy E_i is stored in the battery if there is sufficient space; otherwise, only $E_{\max}$ units of energy is stored. The second method of ordering the energy harvest is *simultaneous transmit*, where energy can be utilized for data transmission in the same channel use as it enters the system. Then remaining energy is stored in the battery if there is sufficient space; otherwise, only $E_{\max}$ units are stored. In general *transmit* *first* and *simultaneous transmit* define two different systems and they yield different achievable rates. In both cases the next battery state S_{i+1} is updated as min $\{S_i - X_i^2 + E_i, E_{max}\}$. Note that the battery state information, S_i, is available at the transmitter only.

The constraints that force the current symbol to have energy less than or equal to the currently available energy are called *energy causality* constraints on the channel input at the channel use level. These constraints follow from the fact that energy cannot be utilized before it enters the system. Note that a non-trivial time correlation arises in the transmitted input sequence X_i due to the energy causality constraints. We also note that the energy arrivals E_i or the energy available at the battery S_i could be viewed as the state of the channel, which determines the set of feasible symbols that can be transmitted.

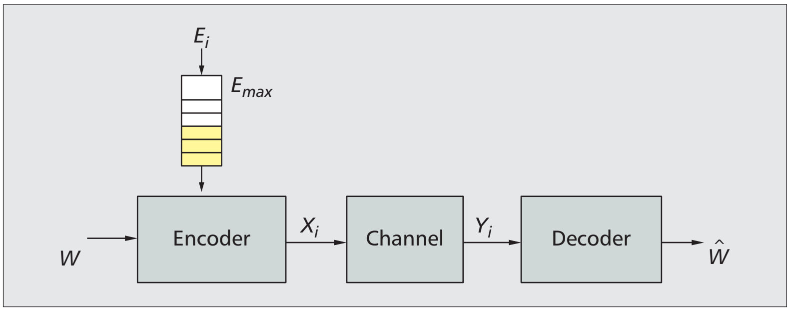

Figure 3. Information-theoretic model of point-to-point communications with an energy harvesting transmitter.

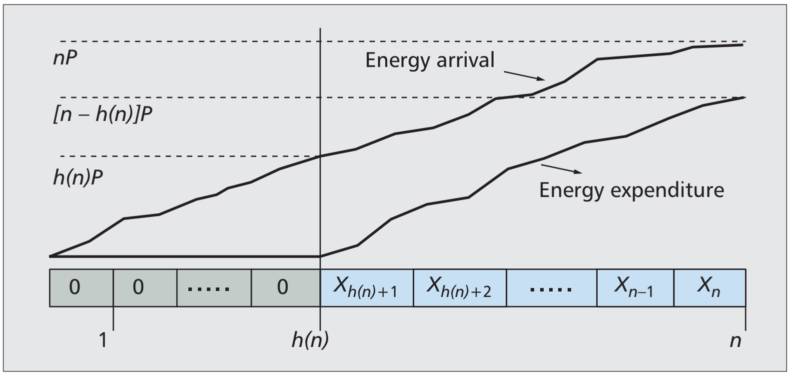

Figure 4. The code symbol energy expenditure and cumulative energy arrival curves in the save-and-transmit scheme.

Channel Capacity with Infinite Energy Storage

In this section we focus on the case when E_{max} is set to infinity in a Gaussian channel. An immediate upper bound for the channel capacity C is the corresponding Shannon capacity with average power constrained to the average recharge rate P_{avg}. This bound holds due to the fact that each codeword satisfying the energy causality constraints automatically satisfies the average power constraint. The work in [7] established that this upper bound is, indeed, achievable and provided two different achievability schemes. We now consider these schemes.

Save-and-Transmit Scheme

In the save-and-transmit scheme, data transmission is performed in two phases, as shown in Fig. 4. The saving phase consists of the first $h(n)$ channel uses; the transmission phase consists of the following $n - h(n)$ channel uses. In the saving phase the transmit symbols are set to 0. Therefore, no energy is spent and the battery is fueled with harvested energy. In the transmission phase, information carrying code symbols are sent. We consider such $h(n)$ functions that grow to infinity sublinearly with n; that is, $h(n)$ is in $o(n)$ and grows to infinity as n grows. The remaining $n - h(n)$ code symbols are the information carrying symbols, which are selected from a Gaussian distribution with zero mean and average power smaller than but arbitrarily close to P_{avg}. Since $h(n)$ is in $o(n)$, the saving phase allows a sufficient number of channel uses for the data transmission phase so that the rate hit due to the saving phase is asymptotically zero. Moreover, since $h(n)$ grows to infinity, the energy saved in the saving phase is sufficient to send the designed code symbols without any energy shortages. This scheme achieves rates arbitrarily close to the Shannon capacity with average power constraint equal to the average recharge rate P_{avg}.

Best-Effort-Transmit Scheme

An alternative single-phase scheme that attains the capacity is the best-effort-transmit scheme. This scheme runs as follows. Let $X_1, X_2, \ldots, X_n$ be a codeword and select X_i as i.i.d. Gaussian with mean zero and variance smaller than but arbitrarily close to P_{avg}. We say that the symbol is infeasible if there is not sufficient energy to send it, i.e. if the battery energy S_i is less than X_i^2. In this case, the input to the channel is set to 0, otherwise it is X_i. That is, in the transmitted codeword, some of the symbols in the actual codeword in the codebook are replaced with zeros. This causes a mismatch between the encoder and the decoder. However, the resulting mismatch is proved to be negligible in [7], and communication with rates arbitrarily close to the Shannon capacity with average power constraint equal to the average recharge rate P_{avg} is possible.

Comments on the Schemes

In the save-and-transmit scheme, the available channel uses are divided into two phases. The saving phase duration $h(n)$ is selected to have a sublinear growth to infinity. For example, $h(n)$ can be selected as $\log(n)$ or $\sqrt{n}$. This, along with the unlimited sized battery, allows averaging out the uncertainty in the available energy. Remaining $n - h(n)$ channel uses are used for channel coding with an average power constraint equal to the average recharge rate. In contrast, in the best-effort-transmit scheme, transmission starts right away and code symbols are put to the channel only when it is feasible to do so. As long as the average energy expenditure of the codewords are below the average recharge rate, any infeasibilities become negligible asymptotically and reliable decoding is possible.

It is crucial to note that both save-and-transmit and best-effort-transmit schemes need unlim-

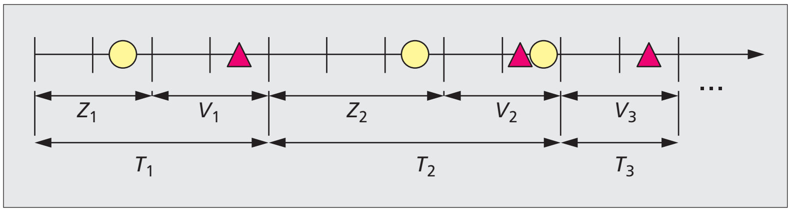

Figure 5. Graphical representation of the variables T_n, V_n and Z_n in the equivalent timing channel. An energy arrival is represented by a circle and a 1 channel symbol is represented by a triangle.

ited sized batteries. It is more obvious that the save-and-transmit scheme needs an unlimited sized battery, since the battery energy needs to go to infinity in the saving phase as the block length becomes large. The fact that the best-effort-transmit scheme also needs an unlimited sized battery is less obvious. While the best-effort-transmit scheme starts transmission right away, since average energy expenditure is less than average recharge rate, eventually the battery energy goes to infinity. In fact, this is the reason that energy shortages occur only in finitely many channel uses. Essentially, after a large enough channel use index, the battery has so much energy that no energy shortages occur.

We note that the capacity with an energy harvesting transmitter with an unlimited battery is invariant to the memory in the energy arrival process, so long as the energy arrival process is stationary and ergodic. That is, an i.i.d. energy arrival process and a non-i.i.d. energy arrival process with the same average arrival rate will yield the same capacity so long as the non-i.i.d. energy arrival process is stationary and ergodic.

Channel Capacity with Zero Energy Storage

In this section we consider the other extreme where there is no battery at the transmitter, i.e. E_{max} is zero, and at each channel use, the incoming energy is either used or it is lost. In this case, we assume *simultaneous transmit order*, as storing the incoming energy is not possible, which is required for the *transmit first order*. In this case, stochastic energy levels at the transmitter cause stochastic instantaneous amplitude constraints, which is a generalization of the static amplitude constrained Gaussian channel in [8]. Indeed, the channel could be viewed as a state-dependent channel with state information available at the transmitter only, where the state is the available energy [9]. In this case, Shannon strategies are the capacity achieving schemes. This is known due to Shannon's work [10] for the capacity of state-dependent channels with state information at the transmitter only. In particular, let us assume without loss of generality that the energy arrivals are selected from a finite alphabet with cardinality $|\mathcal{E}|$. By assigning an input T_i to be put as the channel input whenever E_i arrives, the channel capacity is the maximum mutual information between the extended channel input $T_1, \ldots, T_{|\mathcal{E}|}$ and the channel output; maximization is with respect to the joint distributions of $T_1, \ldots, T_{|\mathcal{E}|}$.

For the Gaussian channel, the optimal input distribution of the extended channel inputs $T_1, \ldots, T_{|\mathcal{E}|}$ are numerically verified to be discrete distributions in [9]. This observation is compatible with the discreteness result in [8]. However, the exact proof for the discreteness of the capacity achieving input distribution for the general case is still an open problem in the current literature.

Channel Capacity with Finite Energy Storage

We now focus on the practically relevant case where the battery size is neither zero nor infinity, but a finite number. Determining the channel capacity with a finite-sized battery, in general, is a difficult problem. In the infinite and zero energy storage cases, the channel capacity is achieved by transmission strategies that do not vary with the time index. However, this does not immediately extend to the finite-sized battery case, and the optimal transmission strategies may potentially need to use all past information of energy arrivals and signal transmissions. A complete answer to this question is still missing in the literature, but significant progress has been made. In the following we provide currently available results for this case.

General Case

In [11] Mao and Hassibi studied the channel capacity for the energy harvesting channel with finite energy storage, and found a general capacity formula in terms of Shannon strategies [10]. This formula involves maximization of the mutual information rate between a sequence of Shannon strategies and the sequence of corresponding channel outputs. Moreover, it is conjectured that the optimal sequence may be obtained by observing the current battery energy level only and ignoring the history of the battery level sequence. In addition, [11] presents a set of achievable schemes using Shannon strategies. These schemes have the desirable property that the corresponding achievable rates, i.e. the limit of the n-letter mutual information rates, can be calculated by a simple simulation-based algorithm.

Deterministic Energy Arrivals

In [12] Jog and Anantharam studied the capacity of the Gaussian channel with a finite-sized battery and deterministic energy arrivals. In this work the channel capacity is determined as the maximum n-letter mutual information rate between the input and output sequences. In contrast to the general capacity formula in [11], the n-letter mutual information rate does not involve Shannon strategies [10]. An implication of this result is that the need for Shannon strategies stems from the randomness in the energy arrivals, and the fact that the state information is available only at the transmitter. In addition, entropy-power-inequality based lower bounds are found for the capacity of this channel. Numerical results show that even with small battery sizes, the capacity approaches quickly to the capacity with an infinite-sized battery.

An Approximate Capacity Approach

In [13] Dong and Ozgur introduced a new approach to find an approximate capacity for the Gaussian channel with a finite-sized battery. The main result in this work is that the channel capacity with any finite-sized battery can be obtained within a constant gap from the channel capacity with infinite-sized battery. In order to establish this result, an achievable rate for a genie-aided receiver with energy arrival side information is derived. Then it is shown that this rate translates into a lower bound on the capacity by a novel evaluation of the side information.

Noiseless Binary Energy Harvesting Channel with Unit Energy Storage

Special cases may be helpful in the challenging problem of determining the channel capacity of the energy harvesting channel with a finite-sized battery. We now present the results of studying a noiseless binary channel with a unit-sized energy storage [14]. Note that even though the channel is noiseless, uncertainty of the receiver regarding the battery energy level at the transmitter side makes it challenging for the receiver to decode the messages of the transmitter.

In this binary energy harvesting channel, encoding and decoding can be equivalently performed over the time intervals between two consecutive 1s. This establishes an equivalence between this channel and a timing channel with additive geometric noise, where this noise state information is available at the transmitter only. In a timing channel, channel symbols are sent by the intervals between the time the packet is put in the queue and the time it is served out of the queue [15]. In our binary communication channel, a "packet" is replaced with a 1 symbol. The codebook used to transmit messages acts as a server to the 1 symbol in the energy queue, and determines its dynamics. The energy available at the energy queue is the state, which determines the set of feasible symbols that can be transmitted. Therefore, the transmitter causally observes the time a 1 symbol waits to be served in the queue.

The equivalent timing channel is represented by the input V_n, noise Z_n, and output T_n, as shown in Fig. 5. The output T_n is the time interval between two 1 symbols at the output, and the input V_n is the number of channel uses a 1 symbol is released from the queue as a channel symbol after Z_n units of waiting. Accordingly, T_n is equal to $V_n + Z_n$. The receiver observes T_n perfectly as the channel is noiseless.

Since Z_n is available at the transmitter before determining V_n, this is a state-dependent channel with state information available at the transmitter only, where the state is the noise Z_n. Hence, the channel capacity in this case can be expressed with a single-letter Shannon strategy at the transmitter [14]. In particular, the capacity is the maximum mutual information between an auxiliary random variable U (independent of the state Z) and the channel output T per average number of channel uses in the classical channel $E[T]$. In this case the channel input is given as a deterministic function V of the auxiliary variable U and the noise Z. In general, a single-letter capacity expression is more desirable compared to an n-letter expression. However, the infinite cardinality requirement of the auxiliary random variable U and determination of the function V cause difficulties in evaluating the capacity.

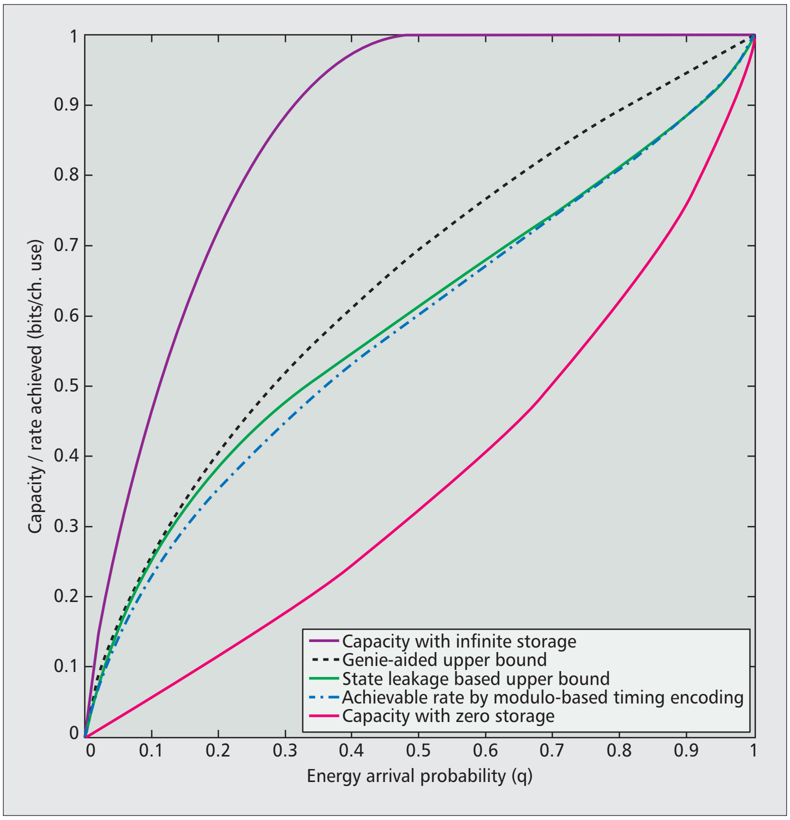

Figure 6. Comparison of upper bounds and achievable rates in [14].

The work in [14] proposes modulo-based timing encoding schemes. Moreover, [14] provides an upper bound for the channel capacity based on the state information leakage to the receiver side. Numerical studies show that this upper bound outperforms an older genie-aided upper bound and is quite tight. We present the comparison of upper bounds and achievable rates in Fig. 6.

Conclusions and Future Directions

In communications with energy harvesting devices, the energy source is intermittent and random. In addition, a battery is available to save the excess energy and use it later. The uncertainty of the available energy at the transmitter, which is also partially controlled by the transmitter's own actions, creates unprecedented constraints on the channel inputs, and this, together with the unavailability of the energy state information at the receiver side, makes the determination of the ultimate capacity limits of such channels a timely and challenging problem in modern wireless communication theory. In this article we have presented partial results for the general finite-sized battery case and com-

> *The uncertainty of the available energy at the transmitter, which is also partially controlled by the transmitter's own actions, creates unprecedented constraints on the channel inputs, and this, together with the unavailability of the energy state information at the receiver side, makes the determination of the ultimate capacity limits of such channels a timely and challenging problem.*

plete results for the special cases of very large sized batteries and very small sized batteries.

We have observed that the capacity critically depends on the size of the battery. For instance, for the Gaussian channel, while for an infinite-sized battery case the capacity can be obtained by using classical Gaussian codebooks together with save-and-transmit and best-effort-transmit schemes, for the zero-sized battery case the capacity is achieved with time-varying discrete channel inputs. For the finite-sized battery case, we have presented results that give partial answers to the case of general channel and general energy arrivals, and lower and upper bounds for deterministic energy arrivals in Gaussian channels, and constant gap results for general energy arrivals in Gaussian channels. For a simple noiseless binary channel with unit energy storage, we have shown an interesting correspondence to the timing channel and an exact single-letter capacity expression whose evaluation required determination of an auxiliary random variable.

This literature is in its infancy and there are several open research directions. From an information theory point of view, the capacity for a general channel with a general finite-sized battery is still open. In addition, capacities of the channels with various side information availability conditions are also open. For instance, if the receiver also harvests energy from the same source or a similar source, then the receiver will have highly correlated side information regarding energy arrivals at the transmitter, and the capacity and achievable schemes in this case are open problems. From a coding theory point of view, explicit coding schemes are open problems. From a wireless communication theory point of view, practical adaptive coding and modulation schemes, and finite block length scenarios, are open directions. From a networking point of view, capacity regions of multi-user versions of these problems (multiple access, broadcast, interference, relay networks) are interesting future research directions. From a signal processing perspective, how the received signals should be processed, and how energy expenditure for processing and other essential components should be accounted for, are important open problems.

References

[1] D. Gunduz *et al.*, "Designing Intelligent Energy Harvesting Communication Systems," *IEEE Commun. Mag.*, vol. 52, Jan. 2014, pp. 210–16.

[2] V. Raghunathan, S. Ganeriwal, and M. Srivastava, "Emerging Techniques for Long-Lived Wireless Sensor Networks," *IEEE Commun. Mag.*, vol. 44, Apr. 2006, pp. 108–14.

[3] J. Yang and S. Ulukus, "Optimal Packet Scheduling in an Energy Harvesting Communication System," *IEEE Trans. Commun.*, vol. 60, Jan. 2012, pp. 220–30.

[4] K. Tutuncuoglu and A. Yener, "Optimum Transmission Policies for Battery Limited Energy Harvesting Nodes," *IEEE Trans. Wireless Commun.*, vol. 11, Mar. 2012, pp. 1180–89.

[5] O. Ozel *et al.*, "Transmission with Energy Harvesting Nodes in Fading Wireless Channels: Optimal Policies," *IEEE JSAC*, vol. 29, Sept. 2011, pp. 1732–43.

[6] C. K. Ho and R. Zhang, "Optimal Energy Allocation for Wireless Communications with Energy Harvesting Constraints," *IEEE Trans. Signal Proc.*, vol. 60, Sept. 2012, pp. 4808–18.

[7] O. Ozel and S. Ulukus, "Achieving AWGN Capacity under Stochastic Energy Harvesting," *IEEE Trans. Inf. Theory*, vol. 58, Oct. 2012, pp. 6471–83.

[8] J. G. Smith, "The Information Capacity of Amplitude and Variance-Constrained Scalar Gaussian Channels," *Information and Control*, vol. 18, Apr. 1971, pp. 203–19.

[9] O. Ozel and S. Ulukus, "AWGN Channel Under Time-Varying Amplitude Constraints with Causal Information at the Transmitter," *Asilomar Conf.*, Nov. 2011.

[10] C. E. Shannon, "Channels with Side Information at the Transmitter," *IBM J. Research and Development*, vol. 2, no. 4, 1958, pp. 289–93.

[11] W. Mao and B. Hassibi, "On the Capacity of a Communication System with Energy Harvesting and a Limited Battery," *IEEE ISIT*, July 2013.

[12] V. Jog and V. Anantharam, "An energy Harvesting AWGN Channel with a Finite Battery," *IEEE ISIT*, June 2014.

[13] Y. Dong and A. Ozgur, "Approximate Capacity of Energy Harvesting Communication with a Finite Battery," *IEEE ISIT*, June 2014.

[14] K. Tutuncuoglu *et al.*, "Improved Capacity Bounds for the Binary Energy Harvesting Channel," *IEEE ISIT*, June 2014.

[15] V. Anantharam and S. Verdu, "Bits Through Queues," *IEEE Trans. Info. Theory*, vol. 42, Jan. 1996, pp. 4–18.

Biographies

OMUR OZEL [S'08, M'15] (ozel@berkeley.edu) received a Ph.D. degree in electrical and computer engineering (ECE) from the University of Maryland (UMD) College Park in 2014. He is currently a postdoctoral scholar in the Electrical Engineering and Computer Sciences Department at the University of California Berkeley. His doctoral research was awarded a Distinguished Dissertation Fellowship by the ECE Department at UMD. His research interests lie in the intersection of wireless communications, information theory, and networking.

KAYA TUTUNCUOGLU [S'08] (kaya@psu.edu) received two B.S. degrees in electrical engineering and in physics from the Middle East Technical University, Ankara, Turkey, in 2008 and 2009, respectively. He is currently pursuing the Ph.D. degree in the Department of Electrical Engineering, Pennsylvania State University, University Park, PA. His research interests include green communications, energy harvesting, and resource allocation for wireless networks. He received the AT&T Graduate Fellowship in 2012, and the IEEE Marconi Prize Paper Award in Wireless Communications in 2014.

SENNUR ULUKUS [S'90, M'98] (ulukus@umd.edu) is a professor of electrical and computer engineering at the University of Maryland at College Park, where she also holds a joint appointment with the Institute for Systems Research (ISR). Prior to joining UMD, she was a senior technical staff member at AT&T Labs-Research. She received her Ph.D. from WINLAB, Rutgers University, and B.S. and M.S. degrees from Bilkent University. Her research interests are in wireless communication theory and networking, and network information theory. She received the 2003 IEEE Marconi Prize Paper Award in Wireless Communications, a 2005 NSF CAREER Award, the 2010-2011 ISR Outstanding Systems Engineering Faculty Award, and the 2012 George Corcoran Education Award.

AYLIN YENER [S'91, M'00, SM'13, F'14] (yener@ee.psu.edu) has been a professor of electrical engineering at The Pennsylvania State University, University Park, PA since 2010; she joined the faculty as an assistant professor in 2002. During the academic year 2008-2009 she was a visiting associate professor in the Department of Electrical Engineering, Stanford University, CA. Her research interests are in information theory, communication theory, and network science, with recent emphasis on energy harvesting wireless communications, information security, and foundations of complex networks.