2017 IEEE 56th Annual Conference on Decision and Control (CDC)
December 12-15, 2017, Melbourne, Australia

Battery Management for Control Systems with Energy Harvesting Sensors

Nicholas J. Watkins, Konstantinos Gatsis, Cameron Nowzari, and George J. Pappas

***Abstract*— In this paper, we study the problem of computing the minimum battery capacity required to stabilize a scalar plant communicating with an energy harvesting sensor over a wireless communication channel. We prove that a particular greedy battery management policy suffices to stabilize the plant, and demonstrate that stability of the system under the greedy policy can be checked by a linear program. Moreover, we show that a critical battery capacity exists, below which no policy can stabilize the system, which itself can be computed by solving a sequence of linear programs which grows logarithmically with respect to the maximum allowed storage capacity. The first of these results address an open question pertaining to the stability of energy harvesting control systems. The last allows us to efficiently compute the smallest battery capacity required to stabilize a given system, which addresses a problem of importance when device size or cost are significant concerns.**

I. Introduction

Energy harvesting technology is playing a significant role in the development of future smart infrastructures by instrumenting low-energy sensing devices at remote locations. In recent years, much progress has been made in understanding how to use energy harvesting devices in networking and communications applications [1]–[3]. However, there is scarcely any literature detailing how energy harvesting sensors can be used in *control* applications, where the closed-loop system's dynamical behavior is of importance. A key desired property of such applications is system stability, which has to be taken into account when designing low-cost devices with limited harvesting and storage capacity.

To date, most works considering the analysis of energy harvesting control systems come coupled with regularity assumptions on the underlying system which guarantee closed-loop stability, or do not explicitly discuss system stability. For example, consider [4], one of the first analytical works to consider sensing with energy harvesting devices. In this paper, the structure of optimal sensing strategies which minimize symmetric distortion costs incurred when estimating the state of finite-state Markov chains and linear-Gaussian systems with orthogonal state propagation matrices is studied. In this case, stability of the estimator for Markov chains is guaranteed by the stability of the process, and stability of the linear-Gaussian system is not explicitly discussed. More recent works have provided conservative stability conditions, which apply to systems with restricted energy harvesting processes [5]–[8]. Instead of directly assuming the open-loop systems to be stable, these works assume that at every time increment in the process, enough energy arrives so that the sensor can communicate reliably enough with the plant to guarantee stability of the plant or an estimator. Since this assumption directly implies that a positive amount of energy is harvested by the sensor at all times with a positive probability, these results do not apply in general circumstances. Indeed, systems which harvest energy through stochastic sources such as wind or sunshine may experience long periods of time in which no energy is received [2]. The only work to date known by the authors which follows an alternate approach to providing a guarantee of stability is [9], which finds conditions under which the covariance error of a Kalman filter of an unstable plant is stable. Despite providing tight stability conditions, the analysis only applies to systems with stochastic $0/1$ energy harvesting sources with independent, identically distributed increments. Since energy harvesting sources are often complex stochastic processes, a more general theory is needed.

We begin developing such a theory in this paper. In particular, we consider a model for an energy harvesting sensor sending feedback signals to an open-loop unstable plant over a communication channel with stochastic packet drops. Our study provides four key contributions:

1) A state-independent greedy battery management policy which is guaranteed to stabilize the system if *any* battery management policy stabilizes the system;
2) An efficient numerical test (linear program) for determining if a given energy harvesting control system is stable under the greedy battery management policy;
3) A proof that a critical battery capacity exists, below which no policy can stabilize the system;
4) A tractable search method for identifying the critical battery capacity required for the existence of a stabilizing battery management policy for a given dynamical system and energy harvesting process.

The first three results are technical, but important. They address open questions regarding the stability of energy harvesting control systems. The last result is an application of the first three and is practically important as battery size plays a critical role in device size and cost (see, e.g. [10]). If we are to develop energy harvesting technology which is economical, we must understand how to properly select the storage capacity of the batteries employed.

The paper is organized as follows. The architecture of the system we study is presented in Section II, along with a formal problem statement. The main results of our paper are contained in Section III. Section IV contains an example energy harvesting control system, and a simulation of its behavior. Section V concludes the paper. Note that some proofs have been deferred for later publication, in favor of limiting the size of this manuscript.

Notation: We denote by $\mathbb{Z}_{\geq 0}$ the set of non-negative integers, and for each $k \in \mathbb{Z}_{\geq 0}$, we denote by $[k]_0$ the set

N.J. Watkins, K. Gatsis, and G.J. Pappas are with the Department of Electrical and Systems Engineering, University of Pennsylvania, Pennsylvania, PA 19104, USA, {nwatk,kgatsis,pappasg}@upenn.edu; C. Nowzari is with the Department of Electrical and Computer Engineering, George Mason University, Fairfax, VA 22030, USA, cnowzari@gmu.edu.

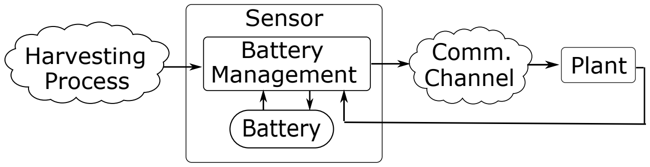

Fig. 1: Energy is supplied to an energy harvesting sensor via an external energy harvesting process which is then either immediately used for providing a feedback signal to a plant via a wireless channel or stored for later use in a finite-capacity battery.

of non-negative integers $\{0, 1, ..., k\}$. We denote by $[\![k]\!]_a^b$ the projection of k into the interval $[a, b]$. •

II. Problem Statement

A visual representation of the architecture of the system we study is given in Figure 1. The system models a setting in which an energy-harvesting sensor communicates over a stochastic communication channel to stabilize the evolution of a plant. The sensor stores energy in the battery, and restores its charge via a stochastic energy harvesting process. Our role in designing this system is in determining a principled method for deciding precisely when it is necessary to provide feedback over the communication channel in order to generate stable closed-loop behavior, while adhering to strict energy causality constraints. As such, we focus on determining precisely when battery management policies which are causal with respect to the battery state, the energy harvesting process, and plant state exist which stabilize the evolution of the plant state. We now detail mathematical models for each component of the system, and provide a formal problem statement.

A. Plant Dynamics

The architecture is deployed to control a scalar linear system with the dynamics

$$x(t+1) = \begin{cases} a_c x(t), & \gamma(t) = 1; \\ a_o x(t), & \gamma(t) = 0; \end{cases} \quad (1)$$

where the random variable $\gamma(t)$ indicates whether or not the plant has received a feedback signal at time t, a_c is a real constant which describes the nominal evolution of the state $x(t)$ when the system operates in closed loop (i.e. has successfully received a feedback signal), a_o is a real constant which gives the nominal evolution of the state variable $x(t)$ in open loop (i.e. when the system has not received a feedback signal). Note that the system model given here is intentionally simple. Not only is it the case that many applications for low-cost energy harvesting sensors will be to simple systems, which may be approximated well by scalar systems, but also many of the results can be extended to the case in which the state dynamics (1) are high dimensional, and subject to stochastic disturbances. However, the analysis is more involved as the details become more significant, and so we leave such cases for publication in future work. For purposes of focusing on the most interesting case, we assume throughout the paper that $0 < |a_c| < 1 < |a_o|$. In every other case, the stability of the system is not in question: if both $|a_c|$ and $|a_o|$ are strictly in the unit interval, the system is stable under any switching process $\{\gamma(t)\}$; if both $|a_c|$ and $|a_o|$ are larger than 1, then the system is unstable under any switching process $\{\gamma(t)\}$.

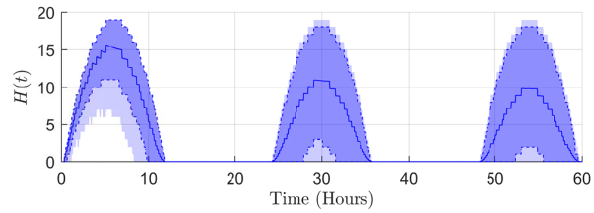

Fig. 2: A plot of a 10k sample Monte Carlo simulation of a hidden Markov model of solar intensity. The 80% confidence region is given in dark blue; the 98% confidence region given in light blue.

B. Communication Channel

In order to model channel imperfections and the decision process involved in determining when to transmit a signal, we model the distribution of $\gamma(t)$ as itself being a function of the amount of energy committed by the sensor to transmitting the feedback signal at time t. We interact with the behavior of $\{\gamma(t)\}$ by selecting the sensor transmission energy $\{\mathcal{E}(t)\}$ at each time t, where the selection may in general be stochastic, in which case we design its distribution. The probability that the plant successfully receives the communication conditioned on a particular transmission energy ε is given by

$$\mathbb{P}(\gamma(t) = 1 \,|\, \mathcal{E}(t) = \varepsilon) = \begin{cases} \lambda, & \varepsilon \geq \bar{\varepsilon}, \\ 0, & \texttt{otherwise}, \end{cases} \quad (2)$$

where $\bar{\varepsilon}$ is an energy threshold above which the transmission is successful with probability λ, and below which all transmissions are unsuccessful. This model well approximates stochastic channels, such as the sigmoid models often considered in practice [11]–[13].

C. Harvesting Processes

We assume the energy harvesting process, i.e. the process which details how much energy is received by the sensor from the environment at each time, takes values in some finite set of integers $\mathcal{H}$, and evolves as an independent discrete-time, discrete-space hidden Markov process. For practical purposes, we assume $\mathcal{H}$ to be bounded above by some finite constant H_{max}. Formally, we can decompose $\{H(t)\}$ into three fundamental components: a discrete, finite set $\mathcal{L}$ of latent process states, an $(|\mathcal{L}| \times |\mathcal{L}|)$-dimensional probability transition matrix L, and a deterministic function $h : \mathcal{L} \mapsto \mathcal{H}$ which maps an element ℓ from the latent space of $\{H(t)\}$ to the range space of $\{H(t)\}$. Since we may take $\mathcal{L} = \{1, 2, \ldots, |\mathcal{L}|\}$ without loss of generality, we characterize harvesting processes $\{H(t)\}$ in this paper by specifying the pair (h, L).

We make no explicit assumptions about the ergodicity or time-invariance of $\{H(t)\}$, and as such our model incorporates as special cases models which range from simple (e.g., with each $H(t)$ taking the value of some fixed constant $h \in \mathcal{H}$) to complicated (e.g., a periodic hidden Markov model). This level of abstraction allows us to incorporate models for a wide variety of sources into the same framework. For instance, we can think of systems subjected to a regular charging cycle as being modeled by a harvesting process which is essentially deterministic (such as the inductive charging used for *in vivo* biomedical sensors [14]), as well as systems subjected to highly stochastic, time-varying charging (such as hidden Markov models used to model the evolution of solar intensity over time, as in Figure 2), with each having an energy source well-modeled by a hidden Markov model.

D. Battery Dynamics

We assume the battery storage process $\{B(t)\}$ to be bounded above by a finite battery capacity constant B_{cap}, which is a parameter to be designed in this paper, where any energy available at time t which is not used used to transmit a feedback signal and is above the battery capacity level is lost due to overflow. This model guarantees that each element $B(t)$ is in the set $[B_{\text{cap}}]_0$, and that the battery capacity process obeys the nonlinear, stochastic update equation

$$B(t+1) = [\![B(t) + H(t) - \mathcal{E}(t)]\!]_0^{B_{\text{cap}}}, \tag{3}$$

where $\mathcal{E}(t)$ is the selected sensor transmission energy (see Section II-B), and $H(t)$ is the amount of energy harvested at time t (see Section II-C). This battery evolution model is common in energy harvesting literature (see, e.g., [5]). To ensure the sensor always only uses energy available to it at each time, we assume that the energy usage process $\{\mathcal{E}(t)\}$ is subject to the energy availability constraint

$$\mathcal{E}(t) \leq B(t) + H(t), \tag{4}$$

which guarantees that the system never uses energy in excess of the current stored energy, and the amount which has been harvested in the current time increment. Note that this constraint implies directly that $\mathcal{E}(t)$ takes values on the set $[H_{\max} + B_{\text{cap}}]_0$ at all times.

E. Battery Management

In this paper, our control over the system's evolution arises through the design of a *battery management policy*, i.e., a stochastic decision rule that the sensor uses to determine when and how to dedicate energy to providing feedback to the plant. We can think of battery management policies as conditional probability distributions, and denote them as

$$\begin{aligned} u(\varepsilon|\{B(\tau),&H(\tau), x(\tau)\}_{\tau\leq t}) \\ &\triangleq \mathbb{P}(\mathcal{E}(t) = \varepsilon \,|\, \{B(\tau), H(\tau), x(\tau)\}_{\tau\leq t}), \end{aligned} \tag{5}$$

where we note that the conditioning implies that the policy u is causal with respect to the process $\{(B(t), H(t), x(t))\}$. As can be seen by (5), the space of battery management policies we consider is very large, and indeed captures most battery management policies one may wish to use in practice.

F. Problem Statement

We wish to determine when a battery management policy u exists such that the evolution of the plant is stable, so we introduce our notion of stability.

Definition 1 (Mean-Square Exponential Stability) The process $\{x(t)\}$ is mean-square exponentially stable if and only if there exists some finite constants $\alpha \geq 1$, and $\xi \in (0,1)$ such that

$$\mathbb{E}(x(t))^2 \leq \alpha\xi^t \mathbb{E}(x(0))^2, \tag{6}$$

holds for all $t \in \mathbb{Z}_{\geq 0}$, and any square integrable $x(0)$. •

Note that while there are other variants of mean-square stability which appear in literature (e.g. asymptotic mean-square stability), we will refer to exponential mean-square stability as simply mean-square stability throughout the paper. As is shown in Section III, the system we study defines a class of Markov jump linear systems (MJLS), for which exponential mean-square stability and asymptotic mean-square stability are equivalent [15].

To make our technical statements concise, we refer to an *energy harvesting control system*, defined as follows:

Definition 2 (Energy Harvesting Control System) An energy harvesting control system (EHCS) is formally defined as the 7-tuple $(a_c, a_o, h, L, \lambda, \bar{\varepsilon}, B_{\text{cap}})$, which encodes the closed-loop dynamics, open-loop dynamics, energy harvesting process, packet reception probability, transmission energy, and battery capacity of the system, respectively. •

From the preceding discussion, we see that the object $(a_c, a_o, h, L, \lambda, \bar{\varepsilon}, B_{\text{cap}})$, contains all parameters of the model proposed. The last definition we require is a formal notion of stabilizability for EHCSs, which we give in the following:

Definition 3 (Energy Harvesting Control System Stabilizability) An EHCS $(a_c, a_o, h, L, \lambda, \bar{\varepsilon}, B_{\text{cap}})$ is said to be stabilizable if and only if there exists a causal battery management policy u as defined by (5) such that state process $\{x(t)\}$ is mean-square stable under u. •

In this language, our problem is to determine the smallest nonnegative integer B_{cap} such that the EHCS $(a_c, a_o, h, L, \lambda, \bar{\varepsilon}, B_{\text{cap}})$ is stabilizable. In addressing this problem, we overcome the following challenges: i) identifying an appropriate battery management policy; ii) determining a means for evaluating the stability of the system; and iii) developing an efficient search mechanism for finding the minimum stabilizing battery capacity.

III. Minimum Battery Sizing for Stabilizability

In this section, we provide the main results of the paper. We first establish that a stabilizing battery management policy exists if and only if a particular deterministic greedy battery management policy stabilizes the system. Then, we develop an efficient numerical test for determining whether or not the greedy battery management policy stabilizes the system. Finally, we provide a description of a binary search procedure for identifying the critical battery capacity after performing a logarithmic number of stability tests.

A. A Stabilizing Battery Management Policy

We introduce the greedy energy management policy that is described mathematically by the conditional probability distribution u_g, defined as

$$u_g(\varepsilon|B(t), H(t)) \triangleq \begin{cases} 1, & \varepsilon = \bar{\varepsilon}, B(t) + H(t) \geq \bar{\varepsilon}; \\ 1, & \varepsilon = 0, B(t) + H(t) < \bar{\varepsilon}; \\ 0, & \texttt{otherwise}. \end{cases} \tag{7}$$

The battery management policy u_g applies exactly $\bar{\varepsilon}$ units of energy to transmitting a feedback signal at precisely those moments at which the sensor has enough energy available to do so. Despite its simplicity - note that it is a deterministic function of the battery and harvesting states at each time - we prove that it is the *only* policy which needs to be investigated to establish the stabilizability of an EHCS.

Theorem 1 (Existence of Stabilizing Policy) *Consider an EHCS* $(a_c, a_o, h, L, \lambda, \bar{\varepsilon}, B_{\text{cap}})$. *The EHCS is stabilizable if and only if the process* $\{x(t)\}$ *is mean-square stable under the greedy battery management policy* u_g *defined by* (7).

We detail next the essential features of the argument supporting Theorem 1. Interestingly, most of the weight of the proof can be shifted onto proving a pathwise stochastic dominance inequality between the greedy policy and any other stochastic policy with respect to the state process $\{x(t)\}$. To demonstrate this, we need to formally define a sample space for the process. For the remainder of the paper, we define the sample space as

$$\Omega \triangleq \{(f_\omega, g_\omega) \mid f_\omega : \mathbb{Z}_{\geq 0} \mapsto \mathcal{L} \times [0,1], g_\omega : \mathbb{Z}_{\geq 0} \mapsto \{0,1\}\},$$

where $f_\omega(t)$ is a vector function containing the evolution of $\{L(t)\}$ in its first component $f_{1\omega}(t)$, and numbers for the randomization required to determine particular actions from a stochastic energy management policy u in its second component $f_{2\omega}(t)$, and $g_\omega(\ell)$ is a function indicating whether or not the ℓ'th feedback attempt reaches the plant successfully. Note that we have defined the sample space to consist of *pairs* of functions so as to be able to index time t and the number of communication attempts ℓ made by the system separately; this technical detail is important.

Intuitively, the function f_w contains all of the randomization needed to model the processes which are indexed naturally with respect to time. From it, we may fully determine the evolution of the harvesting process $\{H(t)\}$, as described in Section II-C, and the randomization required to implement a stochastic battery management policy as described in Section II-E. The function g_ω contains the randomization needed to model the communication channel according to Section II-B. From it, we can determined whether or not the ℓ'th time the transmission energy process exceeds $\bar{\varepsilon}$ - that is, the ℓ'th communication *attempt* - results in a loop closure.

The only subtle point required in verifying that Ω is a proper sample space for the process is in confirming that all process variables are fully determined by the selection of a particular sample path ω. Since all of the process variables at a particular time t can either be determined directly from ω or $\mathcal{E}(t)$, we briefly discuss how one may compute the channel energy at every time from a selected ω. After observing any sequence of events through time t, and under any fixed policy u, we may partition $[0,1]$ into a collection of disjoint intervals $\{\mathcal{I}_{(u,t)}(\varepsilon)\}$ such that the Lebesgue measure of $\mathcal{I}_{(u,t)}(\varepsilon)$ is equal to the probability that $\mathcal{E}(t) = \varepsilon$. By associating to $f_{2\omega}(t)$ the probability measure of a sequence of independent, identically distributed (i.i.d.) uniform random variables on $[0,1]$, we may take $\mathcal{E}(t) = \varepsilon$ for whichever ε satisfies $f_{2\omega}(t) \in \mathcal{I}_{(u,t)}(\varepsilon)$ and have that $\mathcal{E}(t)$ follows the correct distribution.

Note that - unlike in many types of analysis one may perform on models with stochastic control policies - the sample space and probability measure are both unaffected by the particular choice of energy management policy u. This is accomplished by taking the probability measure $\mathbb{P}$ on the sample space Ω to be the product of the probability measures of three independent processes. In particular, we have $\mathbb{P} = \mathbb{P}_{f_1}\mathbb{P}_{f_2}\mathbb{P}_g$, where $\mathbb{P}_{f_1}$ is the measure induced by the latent-state process of $\{H(t)\}$, $\mathbb{P}_{f_2}$ is the measure of a sequence of i.i.d. uniform random variables on the unit interval, and $\mathbb{P}_g$ is the measure of a sequence of i.i.d. Bernoulli random variables with success probability λ. This ability to decouple the choice of strategy from the choice of probability measure is important in that it allows us to compare the performance of different control policies on a sample-by-sample basis. More precisely, if we let $N_u(t;\omega)$ be the number of successful loop closures attained by a battery management policy u through time t on sample path ω, we can show the following:

Lemma 1 (Pathwise Dominance of Greedy Policy) *For all times* $t \in \mathbb{Z}_{\geq 0}$, *and all samples* $\omega \in \Omega$, *it holds that*

$$N_u(t;\omega) \leq N_{u_g}(t;\omega), \tag{8}$$

i.e. the greedy battery management policy dominates every other battery management policy in terms of successful loop closures at every time along every sample path.

A direct consequence of Lemma 1 is that the greedy battery management policy outperforms all others with respect to the state process $\{x(t)\}$, as stated next.

Corollary 1 (Stochastic Dominance Inequalities) *For all times* $t \in \mathbb{Z}_{\geq 0}$, *and all samples* $\omega \in \Omega$, *it holds that*

$$x_{u_g}^2(t;\omega) \leq x_u^2(t;\omega), \tag{9}$$

and hence it also holds that

$$\mathbb{E}[x_{u_g}^2(t)] \leq \mathbb{E}[x_u^2(t)], \tag{10}$$

where the expectation is taken with respect to the measure $\mathbb{P} = \mathbb{P}_{f_1}\mathbb{P}_{f_2}\mathbb{P}_g$.

As a direct consequence of (10), it also holds that $\{x_u(t)\}$ is mean-square stable only if $\{x_{u_g}(t)\}$ is mean-square stable. Since the mean-square stability of $\{x_{u_g}(t)\}$ implies stabilizability of the EHCS, Theorem 1 is proven as well.

B. A Stability Test for Greedy Battery Management

We now develop an efficient test for determining the stability of an EHCS under the greedy battery management policy, by embedding the process into the dynamics of a MJLS. The essence of the embedding we develop is brought to light by noting that the process $\{S(t) \triangleq (B(t), \ell(t), \gamma(t))\}$ is Markovian, and contains everything necessary to model the dynamics of $\{x_{u_g}(t)\}$ as a MJLS.

The process $\{S(t)\}$ evolves on the state space $\mathcal{S} \triangleq [B_{\text{cap}}]_0 \times \mathcal{L} \times \{0,1\}$. We demonstrate that $\{S(t)\}$ is Markovian by verifying that the Markov property holds, i.e. that $S(t+1)$ is independent of $S(t-1)$, when conditioned on $S(t)$. To do so, we demonstrate that transition probabilities $\mathbb{P}(S(t+1) = s'|S(t) = s) \triangleq \psi(s'|s)$ are constants determined by the EHCS's specification. Define b, ℓ, γ, and ε, to be the battery level, latent harvesting process state, loop closure state, and energy usage of the system at state s; define b', ℓ', γ', and ε' likewise for s'. Note that ε is implicitly determined as a function of b and ℓ by (7), and is not explicitly part of the process $\{S(t)\}$. We now show that the probability of transitioning to some state s' from a particular state s depends only on whether or not the evolution of the battery adheres to the battery dynamics, the probability of the required transition occuring in the latent space of the harvesting process, and the probability of the plant successfully receiving feedback in state s'. Since we have assumed $\{L(t)\}$ to be Markovian and independent of the packet drop process, one may check that for all states s' such that the battery dynamics

$$b' = [\![b + h(\ell) - \varepsilon]\!]_0^{B_{\text{cap}}} \tag{11}$$

hold, we have the decomposition

$$\psi_{(s'|s)} = \begin{cases} \lambda L(\ell', \ell) & b' + h(\ell') \geq \bar{\varepsilon}, \gamma' = 1; \\ (1-\lambda) L(\ell', \ell) & b' + h(\ell') \geq \bar{\varepsilon}, \gamma' = 0; \\ 0 & b' + h(\ell') < \bar{\varepsilon}, \gamma' = 1; \\ L(\ell', \ell) & b' + h(\ell') < \bar{\varepsilon}, \gamma' = 0; \end{cases} \tag{12}$$

where $L(\ell', \ell)$ denotes the probability of transition from latent harvesting state ℓ to state ℓ' (see Section II-C). For all other states s', we have $\psi(s'|s) = 0$. Intuitively, the right hand side of (12) partitions the states of $\mathcal{S}$ into events corresponding to whether or not the system has enough energy for the greedy policy to attempt transmitting feedback, and whether or not a feedback signal is received by the plant. The first two events correspond to the cases in which the sensor attempts feedback under the greedy policy; the last two correspond to the opposite.

By noting that the value of $\gamma(t)$ is embedded in $S(t)$, we may define $\gamma(S(t))$ to be the state of the loop closure variable at $S(t)$, and write the plant state dynamics (1) as

$$x(t+1) = a_s x(t), \quad a_s \triangleq \begin{cases} a_c, & \gamma(s) = 1; \\ a_o, & \gamma(s) = 0; \end{cases} \tag{13}$$

where $s = S(t)$. From this, it follows that $\{x(t)\}$ is a MJLS with mode process $\{S(t)\}$. We may now state the stability test we have established in the following theorem.

Theorem 2 (Energy Harvesting Control System Stabilizability Test) *Let $\mathcal{S}$ be the state space of the mode transition process $S(t) = \{B(t), H(t), \gamma(t)\}$ generated by an EHCS. The EHCS is stabilizable if and only if the linear program*

$$\begin{aligned} \underset{(v,r)}{\text{minimize}} \quad & v \\ & a_s^2 \sum_{s' \in S} \psi_{(s'|s)} r_{s'} - r_s \leq v, \ \forall s \in \mathcal{S}, \\ & r_s \geq 1, \ \forall s \in \mathcal{S}, \end{aligned} \tag{14}$$

is unbounded below, where $\psi(s'|s)$ is defined by (12) for all pairs (s, s') such that (11) holds, and $\psi(s'|s) = 0$ otherwise.

Theorem 2 provides a simple, efficient test for determining the stability of an EHCS $(a_c, a_o, h, L, \lambda, \bar{\varepsilon}, B_{\text{cap}})$. In particular, we note that the number of variables and constraints of (14) both grow linearly in the cardinality of $\mathcal{S}$. As a consequence, the complexity of the stability test is at worst polynomial in $|\mathcal{L}|B_{\text{cap}}$, which follows from noting that $|\mathcal{S}| = 2|\mathcal{L}|(B_{\text{cap}}+1)$, and that linear programs can be solved in polynomial time with respect to the number of variables and constraints (see, e.g., [16]).

C. Existence and Computation of Critical Battery Capacity

We now show that stabilizability is a monotone property of EHCSs, with respect to the system's battery capacity. Hence, we cannot reduce the set of stabilizing battery management policies by increasing the storage capacity of the system. Formally, this is stated as follows:

Theorem 3 (Monotonicity of Stabilizability for Energy Harvesting Control Systems) *Stabilizability is a monotone increasing property of EHCSs, i.e. if the system $(a_c, a_o, h, L, \lambda, \bar{\varepsilon}, B_{\text{cap}})$ is stabilizable, then the system $(a_c, a_o, h, L, \lambda, \bar{\varepsilon}, B_{\text{cap}} + 1)$ is stabilizable.*

As a direct consequence of Theorem 3, we can infer that a critical capacity threshold B_{crit} exists for every EHCS. If the battery capacity of the system is chosen to be above B_{crit}, then the system is stabilizable; otherwise, it is not. Note that we may formally prove this by a simple contradiction argument: if no such threshold existed, then the monotonicity property established by Theorem 3 would not hold. This fact helps us in establishing an efficient search method for identifying the critical battery capacity.

Algorithm 1 is essentially a binary search on the battery capacity parameter of the EHCS. At every iteration, it queries a test (see Theorem 2) which determines the stability of the system under the greedy battery management policy with the battery capacity set to the midpoint of known upper and lower bounds, the result of which determines how the bounds are modified. When the lower bound agrees with or surpasses the upper bound, the algorithm terminates. We now formally establish its correctness and time complexity.

Theorem 4 (Critical Battery Capacity Search Analysis) *Algorithm 1 computes $\min\{B_{\text{crit}}, B_{\text{max}}\}$ correctly and in $O(\ln(B_{\text{max}}))$ queries to the stability test (14). Moreover, if no $B \in [B_{\text{max}}]_0$ stabilizes the EHCS, the algorithm returns with a certificate of non-stabilizability.*

The importance of Theorem 4 is in that it reduces the search complexity of identifying the critical battery capacity threshold. This allows the use of large values for our initial guess B_{max} without adding significant computational demands. Algorithm 1 dramatically outperforms a brute force search by limiting the number of queries made to stability tests with a large battery capacity, which are more demanding (see Section III-B).

IV. An Example Energy Harvesting System

We now consider a detailed numerical simulation of an example EHCS. We consider the EHCS with $a_c = 0.95$, $a_o = 1.02$, $\lambda = 0.98$, $\bar{\varepsilon} = 4$, and $\{H(t)\}$ as a process with latent Markov process $\{L(t)\}$ with transition matrix

$$L = \begin{pmatrix} 0.70 & 0.15 & 0.00 & 0.00 \\ 0.30 & 0.70 & 0.15 & 0.00 \\ 0.00 & 0.15 & 0.70 & 0.30 \\ 0.00 & 0.00 & 0.15 & 0.70 \end{pmatrix},$$

and energy function $h(\ell) = \ell - 1$. As constituted, $\{H(t)\}$ is a skip-free random walk on $[H_{\text{max}}]_0$, and can be interpreted as a stochastic model for wind speed. We determine the critical battery capacity threshold to be 3 by using Algorithm 1.

Algorithm 1 Critical Battery Capacity Search

Initialization:

$\bar{B}_{\text{crit}}(0) \leftarrow B_{\max}$;
$\underline{B}_{\text{crit}}(0) \leftarrow 0$;
$\operatorname{stab}(B) \triangleq$ Stability test of Theorem 2;

Main Program:

1: **if** $\bar{B}_{\text{crit}}(k) = \underline{B}_{\text{crit}}(k)$ **then**
2: Run $\operatorname{stab}(\bar{B}_{\text{crit}}(k))$.
3: **if** System is stable with battery capacity $\bar{B}_{\text{crit}}(k)$ **then**
4: Return $B_{\text{crit}} = \bar{B}_{\text{crit}}(k)$;
5: **else**
6: System is not stabilizable for any $B_{\text{cap}} \in [B_{\max}]_0$;
7: **end if**
8: **else if** $\underline{B}_{\text{crit}}(k) > \bar{B}_{\text{crit}}(k)$ **then**
9: System is not stabilizable for any $B_{\text{cap}} \in [B_{\max}]_0$;
10: **else**
11: $B_{\text{cap}}(k) \triangleq \lfloor \frac{\bar{B}_{\text{crit}}(k)+\underline{B}_{\text{crit}}(k)}{2} \rfloor$.
12: Run $\operatorname{stab}(B_{\text{cap}}(k))$
13: **if** System is stable with battery capacity $B_{\text{cap}}(k)$ **then**
14: $\bar{B}_{\text{crit}}(k+1) \leftarrow B_{\text{cap}}(k)$;
15: **else**
16: $\underline{B}_{\text{crit}}(k+1) \leftarrow B_{\text{cap}}(k) + 1$;
17: **end if**
18: Go to 1;
19: **end if**

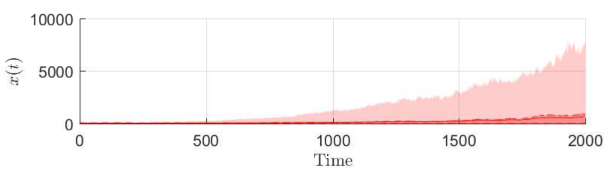

(a) Battery Capacity $B_{\text{cap}} = B_{\text{crit}} - 1 = 2$.

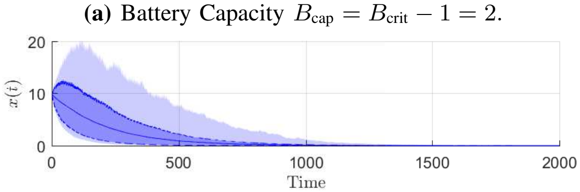

(b) Battery Capacity $B_{\text{cap}} = B_{\text{crit}} = 3$.

Fig. 3: A plot of a 10k sample Monte Carlo simulation of the example EHCS, with the mean trajectory given by a solid line, the 80% confidence interval in dark shade, with the 98% confidence interval given in lighter shade.

The evolution of $\{x(t)\}$ with initial condition $x(0) = 10$, and two levels of B_{cap} is given in Figure 3. By inspection, one can see that for $B_{\text{cap}} = 2$, $\{x(t)\}$ is unstable, as the sample expectation grows without bound, whereas for $B_{\text{cap}} = 3$ the system appears to be stable, with the expectation decaying to 0, and the confidence interval also vanishing in the limit.

V. Conclusions and Future Work

In this paper, we established that the problem of computing the minimal storage capacity required to stabilize a scalar energy harvesting control system can be solved in a computationally efficient manner. By demonstrating that one only needs to check the stability of the system under a particular greedy battery management policy, and that stabilizability is a monotone system property with respect to storage capacity, we reduce the problem of identifying the critical battery storage capacity to solving a sequence of linear programs, the length of which grows logarithmically with respect to the maximum allowable battery capacity. Future work can come in many directions, including considering a situation in which multiple sensors communicate to multiple plants, and generalizing the plant's dynamical model.

Acknowledgments

This is supported by the TerraSwarm Research Center, one of six centers supported by the STARnet phase of the Focus Center Research Program (FCRP), a Semiconductor Research Corporation program sponsored by MARCO and DARPA.

References

[1] S. Ulukus, A. Yener, E. Erkip, O. Simeone, M. Zorzi, P. Grover, and K. Huang, "Energy harvesting wireless communications: A review of recent advances," *IEEE Journal on Selected Areas in Communications*, vol. 33, no. 3, pp. 360–381, 2015.

[2] S. Sudevalayam and P. Kulkarni, "Energy harvesting sensor nodes: Survey and implications," *IEEE Communications Surveys & Tutorials*, vol. 13, no. 3, pp. 443–461, 2011.

[3] S. Basagni, M. Y. Naderi, C. Petrioli, and D. Spenza, "Wireless sensor networks with energy harvesting," *Mobile Ad Hoc Networking: The Cutting Edge Directions*, pp. 701–736, 2013.

[4] A. Nayyar, T. BaÅ§ar, D. Teneketzis, and V. V. Veeravalli, "Optimal strategies for communication and remote estimation with an energy harvesting sensor," *IEEE Transactions on Automatic Control*, vol. 58, pp. 2246–2260, Sept 2013.

[5] M. Nourian, A. S. Leong, and S. Dey, "Optimal energy allocation for kalman filtering over packet dropping links with imperfect acknowledgments and energy harvesting constraints," *IEEE Transactions on Automatic Control*, vol. 59, no. 8, pp. 2128–2143, 2014.

[6] M. Nourian, A. S. Leong, S. Dey, and D. E. Quevedo, "An optimal transmission strategy for kalman filtering over packet dropping links with imperfect acknowledgements," *IEEE Transactions on Control of Network Systems*, vol. 1, no. 3, pp. 259–271, 2014.

[7] S. Knorn and S. Dey, "Optimal sensor transmission energy allocation for linear control over a packet dropping link with energy harvesting," in *2015 54th IEEE Conference on Decision and Control (CDC)*, pp. 1199–1204, IEEE, 2015.

[8] Y. Li, F. Zhang, D. E. Quevedo, V. K. N. Lau, S. Dey, and L. Shi, "Power control of an energy harvesting sensor for remote state estimation," *IEEE Transactions on Automatic Control*, vol. PP, no. 99, pp. 1–1, 2016.

[9] O. Ozel and V. Anantharam, "State estimation in energy harvesting systems," in *Information Theory and Applications Workshop*, 2016.

[10] S. Priya, "Advances in energy harvesting using low profile piezoelectric transducers," *Journal of electroceramics*, vol. 19, no. 1, pp. 167–184, 2007.

[11] K. Gatsis, A. Ribeiro, and G. J. Pappas, "Optimal power management in wireless control systems," *IEEE Transactions on Automatic Control*, vol. 59, pp. 1495–1510, June 2014.

[12] J. Zhang, K. Tan, J. Zhao, H. Wu, and Y. Zhang, "A practical SNR-guided rate adaptation," in *INFOCOM 2008. The 27th Conference on Computer Communications.*, pp. 2083–2091, 2008.

[13] N. Ploplys, P. Kawka, and A. Alleyne, "Closed-loop control over wireless networks," *IEEE Control Systems Magazine*, vol. 24, no. 3, pp. 58–71, 2004.

[14] F. Zhang, X. Liu, S. A. Hackworth, R. J. Sclabassi, and M. Sun, "In vitro and in vivo studies on wireless powering of medical sensors and implantable devices," in *2009 IEEE/NIH Life Science Systems and Applications Workshop*, pp. 84–87, April 2009.

[15] O. L. V. Costa, M. D. Fragoso, and R. P. Marques, *Discrete-time Markov jump linear systems*. Springer Science & Business Media, 2006.

[16] N. Megiddo, *On the complexity of linear programming*. IBM Thomas J. Watson Research Division, 1986.

[17] O. L. Costa and M. D. Fragoso, "Stability results for discrete-time linear systems with markovian jumping parameters," *Journal of mathematical analysis and applications*, vol. 179, no. 1, pp. 154–178, 1993.