Columbia University, Electrical Engineering
Technical Report #2011-01-05, January 2011

Performance Evaluation of Resource Allocation Policies for Energy Harvesting Devices

Maria Gorlatova*, Andrey Bernstein*†, Gil Zussman*
*Electrical Engineering, Columbia University, New York, NY, 10027
†Electrical Engineering, Technion, Haifa, Israel, 32000
Email: mag2206@columbia.edu, andreyb@techunix.technion.ac.il, gil@ee.columbia.edu

Abstract

This work focuses on resource allocation for energy harvesting devices. We analytically and numerically evaluate the performance of algorithms that determine time fair energy allocation in systems with *predictable* and *stochastic* energy inputs. To gain insight into the performance of networks of devices, we obtain results for the simple cases of a single node and a link. Due to the need for *very low complexity* algorithms, we focus on *very simple policies* (some of which proposed in the past as heuristics) and analytically derive performance guarantees. We also evaluate the performance via simulation, using real-world energy traces that we collected for over a year, and in a testbed of energy harvesting devices developed within the EnHANTs project.

I. Introduction

Recent advances in the areas of solar, piezoelectric, and thermal energy harvesting, and in ultra-low-power wireless communications will soon enable the realization of energy harvesting wireless devices. When networked together, they can compose rechargeable sensor networks [5], [14], networks of computational RFIDs [10], and Energy Harvesting Active Networked Tags (EnHANTs) [9]. Such networks will find applications in various areas, and thus *networking energy harvesting devices* has lately been gaining attention. Work in this area includes design of energy-harvesting aware algorithms [5], [6], [8], [11]–[16], [18], development of energy harvesting devices, and characterizations of different energy sources [8], [10] (for reviews of related work see [5], [8], [9]).

Energy sources may have different characteristics. We consider the *predictable profile energy model* [5], [8], [11], [13] in which ideal energy profiles that accurately represent the future are available, and the *stochastic energy model* [6], [8], [12] in which the energy availability can be modeled by a stochastic process. Examples of the latter include a mobile device harvesting light energy, a floorboard that gathers energy when stepped on, and a solar cell in a room where lights go on and off as people enter and leave. In our model, we also consider linear energy storage device (i.e., a battery) and a non-linear device (i.e., a capacitor).

Under both models, energy availability may have high time-variability [10], [11], [18], and therefore, we aim to, as much as possible, allocate the varying energy *in a uniform way with respect to time*. For that, we use the *lexicographic maximization* and the *network utility maximization* frameworks, which are typically applied to achieving fair resource allocation among different nodes rather than among different time slots. Once the energy spending rates are determined by these frameworks, they can be converted to duty cycle, sensing rate, or communication rate.

Energy harvesting shifts the nature of energy-aware protocols from *minimizing* energy expenditure to *optimizing* it over time. Therefore, the resource allocation problems are highly complex [8]. On the other hand, since the devices are resource constrained, there is a need for *very* low (computation and communication) complexity algorithms. While some attempts have been made to develop algorithms for specific types of networks (e.g., directed graphs [14] and trees [5]), most previous work on implementable algorithms focused on a single node or a link [11], [12], [15], [18]. In order to provide insight into the development of low complexity algorithms for a network, we focus in this work on a single node and a link. We analytically and numerically evaluate the performance of *approximate and heuristic* policies, some of which are proposed in [5], [12], [14], [15]. In particular, for a *single node*, we study the following policies:

- **Optimal (*OPT*)** policies for both the *predictable profile* and the *stochastic* models serve as a benchmark for other policies. For the stochastic case, we use a *Markov Decision Process (MDP)*, prove that *energy state discretization* can be applied, and provide bounds on the performance degradation due to discretization.

TABLE I
NOMENCLATURE.

$D(i)$	Environmental energy (J)
K	Number of slots
C	Energy storage capacity (J)
$B(i)$,B_0,B_K	Energy storage state, initial, and final levels (J)
$s(i)$	Energy spending rate (J/slot)
$Q(i)$	Effective energy harvested (J)
h	Quantization resolution (J)
$r(i)$	Data rate (bits/s)
$c_{\text{tx}}, c_{\text{rx}}$	Energetic costs to transmit and to receive (J/bit)
$U(\cdot)$	Utility function
Z	Objective function value
T	Node downtime

- **Spend-What-You-Get (*SG*)** policy – within a time slot a node spends the expected energy input for that slot, and therefore, the complexity is very low (similar policies are proposed in [14], [15]). For *both* models, we provide performance guarantees.
- **Constant Rate (*CR*)** policy – a node spends energy at a constant rate in all time slots, resulting in very low complexity (it is proposed in [5]). For the *predictable profile* model, we provide a performance guarantee.
- **Energy Storage Threshold-based (*THR*)** policy – a set of energy storage thresholds and corresponding rates are chosen, and the node determines the spending rates based on the current storage level (similar policies are proposed in [12], [15]). We study the parameter settings for the *stochastic* model.
- **Energy Storage-Linear (*SL*)** policy – the spending rate is a *linear function* of the energy storage level. We study the parameter settings for the *stochastic* model.

For *links* (node-pairs) we study the following policies:

- **Optimal (*OPT*)** policies (under which nodes need to exchange their parameters) for both energy models.
- **Decoupled Rate Control (*DRC*)** policies – the nodes first determine *independently* their spending rates, and then jointly calculate the data rates (similar approaches are used in [5], [14]). We examine a few versions:
  - **Node-optimal DRC (*DRC-NOPT*)** – the nodes' spending rates are determined according to the *optimal single-node policy*. We provide a performance guarantee for the *predictable profile* model.
  - ***DRC-SG, DRC-CR,* etc.** – one of the above-described policies is used to solve the two single-node problems. These policies are evaluated numerically for the *predictable profile* model.

Within the *Energy Harvesting Active Networked Tags (EnHANTs)* project [9] we have been developing energy harvesting devices and characterizing the availability of indoor ambient light energy. To evaluate the performance of the algorithms, we use simulations based on traces that we collected for over a year [8] as well as experiments with the EnHANTs prototypes [7]. In many of the considered cases, the simple policies perform very well.

II. MODEL AND PRELIMINARIES

We focus on *discrete-time* models, where the time axis is separated into K slots, and a decision is made at the beginning of a slot i ($i = \{0, 1, ..., K-1\}$). We denote the energy storage capacity by C and the amount of energy stored by $B(i)$ ($0 \leq B(i) \leq C$). We denote the initial and the final energy levels by B_0 and B_K. The energy spending rate is denoted by $s(i)$. The amount of energy a device has access to is denoted by $D(i)$, which can be a given value or a random variable. The *effective* amount of energy a device can harvest from the environment is denoted by $Q(i)$. In general, $Q(i)$ may depend both on the available energy $D(i)$ and on the current energy level: $Q(i) = q(D(i), B(i))$ and hence can be non-linear in $D(i)$ (e.g., the *non-linear energy storage model* applies to a capacitor). For a *linear energy storage model* device (such as a *battery*), $q(D(i), B(i)) = D(i)$ and in general $Q(i) \leq D(i)$ [8]. The 'storage evolution' for the models we consider can be expressed as:

$$B(i) = \min\{B(i-1) + Q(i-1) - s(i-1), C\} \tag{1}$$

Note that for the stochastic energy model, we consider *quantizing* the above energy-related parameters, and denote the quantization resolution by h.

We consider a *single node* and a *node pair (link)*. We denote the endpoints of a link by u and v, the effective amount of energy each node can harvest by $Q_u(i)$ and $Q_v(i)$, and their data rates by $r_u(i)$ and $r_v(i)$. For a *single node* we optimize the energy spending rate vector $s(i)$, which provide inputs for determining *duty cycle*, *sensing rate*, or *communication rate*. For a *link*, we optimize either the spending rates $s_u(i)$ and $s_v(i)$ or the communication rates $r_u(i)$ and $r_v(i)$. We denote the costs to transmit and receive bits by c_{tx} and c_{rx}. The constraints relating node energy spending rates and data rates on a link for slot i are:

$$c_{\text{tx}} r_u(i) + c_{\text{rx}} r_v(i) \leq s_u(i), c_{\text{tx}} r_v(i) + c_{\text{rx}} r_u(i) \leq s_v(i). \tag{2}$$

We focus on time-uniform (time-fair) allocation of resources, and use the *lexicographic maximization* and *network utility maximization* frameworks. In the former, we *lexicographically maximize* an energy spending rate vector (for a stand-alone node), or a data rates vector (for a link). In the latter, we maximize the overall utility, where the utility function for each individual assignment is *concave* and *non-decreasing*. For deriving numerical results we use $U(\cdot) = \log(\cdot)$ or $U(\cdot) = \log(1+(\cdot))$. We denote the total objective function value by Z (i.e., $Z = \sum_i U(\cdot)$), and use subscripts to indicate the policy under which Z was obtained (e.g., Z_{OPT} for the *OPT* policy and Z_{CR} for the *CR* policy). As another performance measure, we consider the *downtime* of a node and a link. Namely, the fraction of slots the node or the link do not spend energy. We denote the downtime of a node by $T = |\{i|s(i) = 0\}|/K$ and of a link by $T^L = |\{i|r_u(i) = 0, r_v(i) = 0\}|/K$.

III. Predictable profile Energy Model

In this section we analyze various policies for a *single node* model and discuss a *link model*. Section V provides numerical results demonstrating the performance of the policies based on real-world energy traces.

A. Single Node

The optimal solution for a single node can be obtained by solving the following problem [8].

Time Fair Utility Maximization (TFU) Problem:

$$\max_{s(i)} \left\{ Z \triangleq \sum_{i=0}^{K-1} U(s(i)) \right\} \tag{3}$$

s.t.:

$$s(i) \leq B(i) \ \forall \ i \tag{4}$$

$$B(i) \leq B(i-1) + Q(i-1) - s(i-1) \ \forall \ i \geq 1 \tag{5}$$

$$B(0) = B_0; \ B(K) \geq B_K; B(i) \leq C \ \forall \ i \tag{6}$$

$$B(i), s(i) \geq 0 \ \forall \ i \tag{7}$$

We now provide bounds on the optimal solution as well as an approximation ratio for the *CR* policy. Observation 1 applies to both *linear* and *non-linear* energy storage models, while Observations 2 and 3 apply to the *linear* energy storage model.

Observation 1: $Z_{\text{OPT}} \leq K \cdot U\left(\left(B_0 - B_K + \sum_{i=0}^{K-1} D(i)\right)/K\right)$.

Proof: Denote the total energy a node has available to it by A, where $A = B_0 - B_K + \sum_i Q(i)$. When the energy storage is sufficiently large, the optimal allocation is $s(i) = A/K \ \forall \ i$ [8]. The corresponding Z value is $\tilde{Z} = \sum_i U(s(i)) = K \cdot U(A/K)$. When the energy storage is smaller, the total energy a node has available to it is also A (or less), and it is allocated less uniformly. Thus, due to concavity of U, the Z value for smaller storage conditions will be smaller, and therefore the above-stated $\tilde{Z}$ is the upper bound. ■

Observation 2: The total energy allocated by the optimal solution is $\sum_i \min(Q(i), C) + B_0 - B_K$. The optimal solution will allocate all available energy if $C > \max(Q(i))$.

Proof: Denote the total energy a node could potentially allocate by A, where $A = B_0 - B_K + \sum_i Q(i)$. Note, however, that when the energy storage is finite, the node will not necessarily be able to capture all available energy. As $s(i) \leq B(i)$ (constraint (4)), $B(i) - s(i) \geq 0$. Hence, due to the storage evolution equation $B(i+1) = \min(B(i)+Q(i)-s(i), C)$, the amount of energy that is "captured" by the system in a slot is at most $\min(Q(i), C)$. The objective function is maximized when as much as possible of the energy is captured and allocated. Hence,

the total amount of energy a node allocates is $A' = \sum_i \min(Q(i), C) + B_0 - B_K$. If $\max(Q(i)) < C$, then $Q(i) \leq C \ \forall \ i$, and hence $A' = \sum_i Q(i) + B_0 - B_K = A$ – all energy available to a node is allocated. ■

Observation 3: Under the *CR* policy, for[1] $B_K = B_0 \leq \sum_i Q(i)$ and $U(s) = \log(1+s)$,

$$Z_{CR} \geq Z_{OPT} \cdot \left(\frac{B_0}{\sum_{i=0}^{K-1} Q(i)} \right).$$

Proof: Denote $\sum_i Q(i)/K$ by A, and use f to denote $f = B_0/\sum_i Q$ (hence $0 < f < 1$, and $B_0/K = f \cdot A$). From Observation 1, $Z_{OPT} \leq K \cdot U(A)$. Since $\sum_i Q(i) > B_0 = B_K$, a node can spend B_0 units of energy and have sufficient B_K in storage at the end of the K slots. Thus the spending in each slot, $s(i)$, will be at least B_0/K. Thus, $Z_{CR} \geq K \cdot U(B_0/K) = K \cdot U(f \cdot A)$, and

$$\frac{Z_{OPT}}{Z_{CR}} \leq \frac{K \cdot U(A)}{K \cdot U(f \cdot A)} = \frac{\log(1+A)}{\log(1+f \cdot A)}$$

For $f = 1$, $\log(1+A)/\log(1+f \cdot A) = 1$. For $0 < f < 1$, this expression is a decreasing function of A, with its maximum achieved for $A \to 0$. Using l'Hôpital's rule,

$$\lim_{A \to 0} \left(\frac{\log(1+A)}{\log(1+f \cdot A)} \right) = \lim_{A \to 0} \left(\frac{1/(1+A)}{f/(1+fA)} \right) = \lim_{A \to 0} \left(\frac{1}{1+A} \cdot \frac{1+fA}{f} \right) = \lim_{A \to 0} \left(\frac{1+fA}{f+fA} \right) = \frac{1}{f}.$$

Hence $Z_{OPT}/Z_{CR} \leq 1/f = \sum_i Q(i)/B_0$, and thus $Z_{CR} \geq Z_{OPT} \cdot [B_0/\sum_i Q(i)]$. ■

The following proposition provides an approximation ratio for the *SG* policy for both *linear* and *non-linear* energy storage models.

Proposition 1: Under the *SG* policy and for $U(s) = \log(s+M)$,

$$Z_{SG} \geq Z_{OPT} \cdot \frac{\log(G(Q'))}{\log(\overline{Q'})},$$

where M is a constant, $\overline{(\cdot)}$ and $G(\cdot)$ denote the *arithmetic* mean and the *geometric* mean of a sequence, and $Q'(i) \triangleq Q(i) + M \ \forall i$.

Proof: For $s(i) = Q(i)$, $Z_{SG} = \sum_i \log(M + Q(i)) = \sum_i \log(Q'(i))$. $G(Q')$, the geometric mean of a sequence, can be calculated as $G(Q') = \sqrt[K]{Q'(0) \cdot Q'(1) \cdot ... \cdot Q'(K-1)}$, and can be transformed as:

$$G(Q') = \sqrt[K]{Q'(0) \cdot Q'(1) \cdot ... \cdot Q'(K-1)} = \exp((1/K) \cdot \sum_i \log(Q'(i)) = \exp((1/K) \cdot Z_{SG}).$$

Thus $Z_{SG} = K \cdot \log(G(Q'))$. From Observation 1 we know that $Z_{OPT} \leq KU(\sum_i Q(i)/K) = K \cdot \log(\overline{Q'})$. Thus, $Z_{SG}/Z_{OPT} \geq [K \cdot \log(G(Q'))]/[K \cdot \log(\overline{Q'})]$, and thus $Z_{SG} \geq Z_{OPT} \cdot \log(G(Q'))/\log(\overline{Q'})$. ■

For example, consider a case of $Q(i)$ such that L samples of $Q(i) = A$, and the rest are equal to zero. Such $Q(i)$ may correspond to the case where the indoor lights are on for a portion of the day. Using Proposition 1, we demonstrate that for $B_K = B_0$ (*energy neutrality* [11]) and for $U(s) = \log(1+s)$, the *SG* policy is a K/L*-approximation algorithm* (for instance, if the indoor lights are on for 8 hours per day, the *SG* policy is a 3-approximation algorithm). Denote $\widehat{Q} = \sum_i Q(i)$. For $U(s) = \log(1+s)$, $U(Q(i) = 0) = 0$, and thus:

$$\frac{Z_{OPT}}{Z_{SG}} \leq \frac{\sum(U(\hat{Q}/K))}{\sum(U(\hat{Q}/L))} = \frac{K \cdot U(\widehat{Q}/K)}{L \cdot U(\widehat{Q}/L)} = \frac{K}{L} \frac{\log(\hat{Q}/K+1)}{\log(\hat{Q}/L+1)} \leq \frac{K}{L},$$

where the last inequality comes from the fact that $K > L$, thus $\hat{Q}/K < \hat{Q}/L$, and hence $\frac{\log(\hat{Q}/K+1)}{\log(\hat{Q}/L+1)} < 1$, and thus K/L is the upper bound. This bound cannot be improved, as for $\hat{Q} \to \infty$ we can demonstrate that

$$\lim_{\hat{Q} \to \infty} \frac{K}{L} \cdot \frac{\log(\hat{Q}+K) - \log(K)}{\log(\hat{Q}+L) - \log(L)} = \lim_{\hat{Q} \to \infty} \frac{K}{L} \cdot \frac{\hat{Q}+K}{\hat{Q}+L} = \frac{K}{L}.$$

[1]Namely, under *energy neutrality* [11], with a relatively small energy storage.

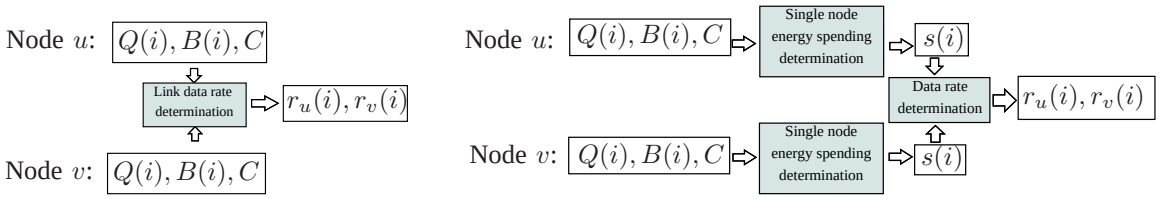

Fig. 1. Overall schematics of an optimal link policy determination (left) and the *DRC* algorithms (right).

B. Node Pair (Link)

The optimal solutions for a link can be obtained by solving the following problems [8].

Link Time Fair Utility Maximization (LTFU) Problem:

$$\max_{r_u(i), r_v(i)} \sum_{i=0}^{K-1} [U(r_u(i)) + U(r_v(i))] \tag{8}$$
$$\text{s.t. : constraints (2) } \forall\ i;\ u, v: \text{ constraints (4)} - \text{(7)}$$

Link Time Fair Lexicographic Assignment (LTFL) Problem:

$$\text{Lexicographically maximize: } \{r_u(0), ..., r_u(K-1), r_v(0), ..., r_v(K-1)\} \tag{9}$$
$$\text{s.t. : constraints (2) } \forall\ i;\ u, v: \text{ constraints (4)} - \text{(7)}$$

The results in this section apply to the *linear* energy storage model. First, we show below that under specific conditions the solutions to both problems are equal.

Proposition 2: When $c_{\text{tx}} = c_{\text{rx}}$, the *LTFL* problem and the *LTFU* problem have the same solution.

Proof: When $c_{\text{tx}} = c_{\text{rx}} = \tilde{c}$, constraints (2) reduce to $\tilde{c} \cdot (r_u(i) + r_v(i)) \leq s_u(i)$ and $\tilde{c} \cdot (r_u(i) + r_v(i)) \leq s_v(i)$, which can be equivalently stated as $r_u(i) + r_v(i) \leq \min(s_u(i), s_v(i))/\tilde{c}$ – that is, in this case the *sum* of the data rates is constrained. Due to the concavity of the objective functions in the *LTFU* problem, the utility of variables constrained like this is maximized when $r_u(i) = r_v(i) = r(i)$. Similarly, the optimal solution to the *LTFL* problem will also have $r_u(i) = r_v(i) = r(i)$, since the optimal solution is max-min fair, and the max-min fairness is achieved when both rates are increased as much as possible. Thus, the *LTFU* problem's objective function reduces to $\max\ \sum_i U(\min(s_u(i), s_v(i))/\tilde{c})$, and the *LTFL* problem's objective function reduces to $lex\ \max[\min(s_u(1), s_v(1))/\tilde{c}, ..., \min(s_u(K-1), s_v(K-1))/\tilde{c}]$, and the constraint set is the same for both *LTFL* and *LTFU*. The equality of optimal solutions to these problems with the same constraint set follows from the proof of Lemma 2 in [8]. ■

We now examine the performance of the following set of algorithms.

Decoupled Rate Control (DRC) Algorithms: For a given link (u, v), the algorithms first determine $s_u(i)$ and $s_v(i)$ for every slot i according to some single-node policy, optimal (*DRC-NOPT*), or approximate (i.e., *DRC-SG*, *DRC-CR*). Then, for each slot i, under constraints (2), the algorithms obtain a solution to

$$\max_{r_u(i), r_v(i)} \{U(r_u(i)) + U(r_v(i))\}. \tag{10}$$

Small per-slot problem (10) can be easily solved. For example if $c_{\text{tx}} = c_{\text{rx}}$, the solution to (10) is

$$r_u(i) = r_v(i) = \min(s_u(i), s_v(i))/(c_{\text{tx}} + c_{\text{rx}}). \tag{11}$$

Fig. 1 shows schematically the difference between solving link problems optimally and applying the *DRC* algorithms.

Proposition 3: $Z_{OPT} \leq K \cdot \tilde{Z}$, where $\tilde{Z}$ is obtained by solving (10) for $s_u(i) \leftarrow [B_{0,u} - B_{K,u} + \sum_i Q_u(i)]/K$, $s_v(i) \leftarrow [B_{0,v} - B_{K,v} + \sum_i Q_v(i)]/K$.

Proof: Consider the case where the energy storage is unlimited. Then constraints (4)-(7) reduce to constraints $\sum_i s_u(i) \leq [B_{0,u} - B_{K,u} + \sum_i Q_u(i)]$, $\sum_i s_v(i) \leq [B_{0,v} - B_{K,v} + \sum_i Q_v(i)]$. Since the objective function is concave and non-decreasing, and since the problem is symmetric for all slots i, the optimal solution will be obtained when each node assigns the same data rate to each slot. The same data rate for each slot will be obtained

when each node assigns the same amount of energy to each slot, that is, $s_u(i) = [B_{0,u} - B_{K,u} + \sum_i Q_u(i)]/K \ \forall \ i$, $s_v(i) = [B_{0,v} - B_{K,v} + \sum_i Q_v(i)]/K \ \forall \ i$, and hence the overall solution will be $K \cdot \tilde{Z}$. This value is an upper-bound since for constrained energy storage conditions, the same amounts of energy will be assigned with more constraints. ■

For example, for the case of $c_{\text{tx}} = c_{\text{rx}}$,

$$\tilde{Z} = 2 \cdot U \left(\frac{\min([B_{0,u} - B_{K,u} + \sum_i Q_u(i)]/K, [B_{0,v} - B_{K,v} + \sum_i Q_v(i)]/K)}{c_{\text{tx}} + c_{\text{rx}}} \right).$$

The Proposition below implies that the *DRC-NOPT* policy obtains the optimal solution to the *LTFL* problem for a link (u, v) in which u *and* v *have the same energy parameters* $Q(i), C, B_0$ and B_K.

Proposition 4: DRC-NOPT solves the *LTFL* problem optimally, if for all slots i, node-optimal $s_u(i) \leq s_v(i)$.

Proof: The optimal solution to the *LTFL* is max-min fair, and thus, for each i, $r_u(i) = r_v(i) = r(i)$, and constraints (2) can be reduced to $r(i) \leq (\min(s_u(i), s_v(i))/[c_{tx} + c_{rx}]$. Thus, the data rate in each slot i is fully determined by the *minimum* of $\{s_u(i), s_v(i)\}$. For $s_u(i) \leq s_v(i) \ \forall i$, the data rates are fully determined by the allocation of $s_u(i)$. The *DRC-NOPT* slot energy spending assignments $s_u(i)$ are lexicographically fair – a spending $s_u(i_1)$ cannot be increased without a decrease in some $s(i_2)$ that is already smaller. Thus, the energy allocation cannot be improved, and thus the *DRC-NOPT* solution is optimal. ■

The following observation discusses the downtime under the *DRC-SG* policy.

Observation 4: Under the *DRC-SG* policy, $\max[T_u, T_v] \leq T^L_{u,v} \leq T_u + T_v$.

Proof: Due to constraints (2), the data rates assigned by a *DRC* policy to a slot i will be zero if $s_u(i) = 0$ or $s_v(i) = 0$, thus $T^L_{u,v}$ is not smaller than $\max[T_u, T_v]$. If both $s_u(i)$ and $s_v(i)$ are non-zero, then the $r(i)$ values maximizing (10) will also be non-zero, thus $T^L_{u,v}$ is not larger than $T_u + T_v$. ■

For example, consider a case where $Q_u(i)$ and $Q_v(i)$ are vectors with L non-zero entries. For a (u, v) where $Q_v(i) = Q_u(i) \ \forall \ i$ ($Q_u(i)$ and $Q_v(i)$ are *synchronized*), $T^L_{u,v} = (K - L)/K$. On the other hand, for a (u, v) where $Q_v(i)$ is *shifted* with respect to $Q_u(i)$, $T^L_{u,v}$ can be as high as $2 \cdot (K - L)/K$.

IV. Stochastic Energy Models

We now study models in which the energy harvested in slot i is a *random process* $\{D(i)\}$. We examine a model of a single node with $\{D(i)\}$ *i.i.d.* random variables. We let D denote the "representative" variable for $D(i)$ and p_D denote its *probability density function* (pdf). In addition, we also briefly discuss the extension of the model to a *link*.

A. Single Node – Optimal Policies and Dicretization Bounds

We formulate the problem as an average cost Markov Decision Process (MDP). Let $\mathcal{B} = [0, C]$ and $\mathcal{S} = [0, C]$ denote the state and action spaces of the MDP, respectively. For any $b \in \mathcal{B}$ and $s \in \mathcal{S}$, the transition density is denoted by $p(\cdot|b, s)$. It determines the next energy storage level $B(i+1)$ given that the current energy storage level is $B(i) = b$ and the spending rate is $s(i) = s$. This transition density is determined by p_D and (1). A policy π is a collection of decision rules $\pi_i : \mathcal{B}^i \times \mathcal{S}^{i-1} \to \Delta(\mathcal{S})$ which at each time i prescribe a probability distribution over the actions ($\Delta(\mathcal{S})$ denotes the probability simplex over the set $\mathcal{S}$). The goal is to find an optimal policy, which maximizes the average utility. In particular, let[2]

$$\lambda_\pi(b) \triangleq \lim_{K \to \infty} (\mathbb{E}_\pi Z(K))/K = \lim_{K \to \infty} \mathbb{E}_\pi \left(\sum_{i=0}^{K-1} U(s(i)) \right) /K$$

denote the asymptotic expected average utility obtained by starting from state $B_0 = b$ and using a given policy π. The optimal average utility is then

$$\lambda^*(b) \triangleq \sup_\pi \lambda_\pi(b).$$

[2] $\mathbb{E}_\pi$ denotes the expectation with respect to the probability law induced by the MDP while using policy π, and $\{s(i)\}$ are the spending rates under this policy.

It is well known (i.e., [17]) that under certain *ergodicity* (or *mixing*) conditions, the optimal average utility does not depend on b. In our case, we use the following mixing condition.

Assumption 4.1 (Mixing): There exists a scalar $\rho \in (0,1]$ and a measure ν with $\nu(\mathcal{B}) \geq \rho$ such that

$$p(A|b,s) \geq \nu(A), \ \forall A \subseteq \mathcal{B}, (b,s) \in \Gamma.$$

For our problem, we prove the following.

Lemma 4.1: If when $B(i) = C$, $s(i) \geq \alpha > 0$ holds for some α, Assumption 4.1 is satisfied. In particular, for the linear case ($q(d,b) = d$) it is satisfied with

$$\nu(y) \triangleq \min_{(b,s)\in\Gamma} p_D(y-b+s), \ y \in \mathcal{B},$$

$$\rho \triangleq \int_{\mathcal{B}} \min_{(b,s)\in\Gamma} p_D(y-b+s)dy > 0.$$

Proof: We prove this Lemma for the case of the linear model. First, trivially

$$\nu(y) \leq p_D(y-b+s) \leq p(y|b,s)$$

for all (y,b,s). Now

$$\rho \triangleq \int_{\mathcal{B}} \nu(y)dy.$$

We show below that if $y > C - \alpha$, where α is defined in (12), we have that

$$\min_{(b,s)\in\Gamma} p_D(y-b+s) > 0.$$

This will imply that $\rho > 0$. Indeed, suppose $y > C - \alpha$. Also, $b - s \leq C - \alpha$. Therefore, $y > b - s$, which implies that $p_D(y-b+s) > 0$. ■

In view of Lemma 4.1, we let

$$\Gamma \triangleq \{(b,s) \in \mathcal{B} \times \mathcal{S} : \ \max(b - C + \alpha, 0) \leq s \leq b\} \tag{12}$$

denote the set of *admissible* state-action pairs.

Under the mixing condition, an optimal policy is *deterministic Markov stationary* policy $\pi^* : \mathcal{B} \to \mathcal{S}$ and can be found by solving the *optimality equation*

$$\lambda + J(b) = \mathcal{T}J(b), \ b \in \mathcal{B},$$

where $\mathcal{T}$ is *Bellman's operator*, defined for any bounded function J as

$$\mathcal{T}J(b) = \max_{s\in\mathcal{S}} \left\{ U(s) + \int_{\mathcal{B}} p(b'|b,s)J(b')db' \right\}.$$

Specifically, a solution (λ^*, J^*) of optimality equation is such that $\lambda^*(b) \equiv \lambda^*$ and an optimal policy is given by

$$\pi^*(b) = \operatorname*{argmax}_{s\in\mathcal{S}} \left\{ U(s) + \int_{\mathcal{B}} p(b'|b,s)J^*(b')db' \right\}.$$

However, since our state and action spaces are infinite, there is no practical algorithm to solve the optimality equation. To address this, we *discretize* the state and action spaces *uniformly*, using a fixed discretization parameter h. We denote thus obtained finite spaces by $\mathcal{B}_h$ and $\mathcal{S}_h$. In particular, if $b \in \mathcal{B}_h$, it is a multiple of h, and similarly for $\mathcal{S}_h$. For any $b \in \mathcal{B}$, we let $x_b \in \mathcal{B}_h$ denote the *representative* point of b in $\mathcal{B}_h$ (which is the closest point to b in $\mathcal{B}_h$).

The discretized set of admissible state-action pairs is then

$$\Gamma_h \triangleq \left\{ (b,s) \in \mathcal{B} \times \mathcal{S}_h : \ |s - s_b| \leq \frac{h}{2} \text{ for some } \max(x_b - C + \alpha, 0) \leq s_b \leq x_b \right\}.$$

Finally, the descretization of the dynamics is as follows:

$$p_h(b'|b,s) \triangleq \frac{p(x_{b'}|b,s)}{\int_{\mathcal{B}} p(x_y|b,s)dy}$$

The corresponding Bellman's operator is then

$$T_h J(b) = \max_{s \ in \mathcal{S}_h} \left\{ U(s) + \int_{b' \in \mathcal{B}} p_h(b'|b,s) J(b') \right\}$$

It is easy to see that this operator returns a *simple* function for any given function J. Moreover, the solution J_h^* of the optimality equation

$$\lambda_h + J_h(b) = T_h J_h(b), \ b \in \mathcal{B} \tag{13}$$

is also a simple function. The solution (λ_h^*, J_h^*) can be found using value/policy iteration algorithms or linear programming (see [17] for details).

We use the results in [2] to provide the performance bounds due to the introduced discretization process. To use these results, in addition to the mixing condition (Lemma 4.1), our MDP model should satisfy the following continuity condition.

Assumption 4.2 (Lipschitz Continuity): There exists a constant $\beta > 0$ such that

$$\left|U(s) - U(s')\right| \leq \beta \left|s - s'\right|, \ \forall s, s' \in \mathcal{S},$$
$$\left\|p(\cdot|b,s) - p(\cdot|b',s')\right\|_v \leq \beta \left\|(b,s) - (b',s')\right\|_\infty, \ \forall (b,s),(b',s') \in \Gamma,$$

where $\|\cdot\|_v$ is the total variation norm.

The first part of Assumption 4.2 can be satisfied by choosing an appropriate utility function. Let us denote its continuity constant by β_U. For the second part, we impose the following on the probability distribution of the energy random variable D.

Assumption 4.3: Suppose that there exists a finite constant $D_{\max}$ such that the variable D takes values in the interval $[0, D_{\max}]$. Let

$$P_{\max} \triangleq \max_{d \in [0, D_{\max}]} p_D(d).$$

Moreover, assume that there exists a finite constant $\beta_D > 0$ such that

$$\left|p_D(d) - p_D(d')\right| \leq \beta_D \left|d - d'\right|, \ \forall 0 \leq d, d' \leq D_{\max}.$$

Lemma 1: Under Assumption 4.3, there exists $\beta = \beta(\beta_U, \beta_D) > 0$ such that Assumption 4.2 is satisfied. In particular, for the linear model ($q(d,b) = d$) we have that

$$\beta = \max\left\{\beta_U, 2(C\beta_D + P_{\max})\right\}.$$

Proof: We prove this Lemma for the linear model case. In this case

$$B(i+1) = \min\left\{B(i) + D - s, C\right\}$$

where $D \sim P_D$. Let $F(y|b,s)$ ($b - s \leq y \leq \min\{b - s + D_{\max}, C\}$) denote the cumulative transition distribution function. We have that

$$\begin{aligned}
F(y|b,s) &= \mathbb{P}\left\{\min\left\{b + D - s, C\right\} \leq y\right\} \\
&= \mathbb{P}\left(\left\{\min\left\{b + D - s, C\right\} \leq y\right\} \bigcap \left\{b + D - s \leq C\right\}\right) \\
&\quad + \mathbb{P}\left(\left\{\min\left\{b + D - s, C\right\} \leq y\right\} \bigcap \left\{b + D - s > C\right\}\right) \\
&= \mathbb{P}\left\{D \leq y - b + s\right\} + \mathbb{I}\left\{C = y\right\} \mathbb{P}\left\{D > C - b + s\right\} \\
&= F_D(y - b + s) + \mathbb{I}\left\{C = y\right\}\left(1 - F_D(C - b + s)\right),
\end{aligned}$$

where F_D is the cdf of D. Thus, the transition density is given by

$$p(y|b,s) = p_D(y-b+s) + \delta(y-C)\left(1 - F_D(C-b+s)\right).$$

Therefore, for any (b,s) and (b',s'), we have that

$$\begin{aligned}
\left\|p(\cdot|b,s) - p(\cdot|b',s')\right\|_v &= \int_0^C \left|p(y|b,s) - p(y|b',s')\right| dy \\
&\leq \int_0^C \left|p_D(y-b+s) - p_D(y-b'+s')\right| \\
&\quad + \int_0^C \left(\delta(y-C)\left|F_D(C-b+s) - F_D(C-b'+s')\right|\right) dy \\
&\leq C\beta_D(\left|b-b'\right| + |s-s'|) + \left|F_D(C-b+s) - F_D(C-b'+s')\right| \\
&\leq 2C\beta_D \left\|(b,s) - (b',s')\right\|_\infty + 2P_{\max} \left\|(b,s) - (b',s')\right\|_\infty,
\end{aligned}$$

where the third inequality follows by Assumption 4.3. Thus, Assumption 4.2 is satisfied with

$$\beta = \max\left\{\beta_U, 2(C\beta_D + P_{\max})\right\}.$$

as required. ■

Hence, the following theorem bounds the distance of the optimal average reward λ_h^* in the discretized model from the optimal average reward λ^*. This theorem is in fact an application of Theorem 4.3.5 in [2] to our case.

Theorem 1: Under Assumption 4.3, there exists $\bar{h} > 0$ and β_λ (depending only on β of Assumption 4.2 and ρ of Assumption 4.1) such that for all $h \in (0, \bar{h}]$

$$|\lambda^* - \lambda_h^*| \leq \beta_\lambda h.$$

In particular, $\bar{h}$ and β_λ can be explicitly written as

$$\bar{h} = \min\left\{\frac{1}{2\beta + 4\beta^2}, \frac{\rho}{\beta + \beta(\beta+0.5) + 4(\beta+\beta^2)}\right\}$$

$$\beta_\lambda = \beta_1\left(1 + \frac{4}{\rho}\right) + \beta_2\left(\frac{\beta}{\rho} + \frac{2\beta}{\rho^2}\right),$$

where

$$\beta_1 \triangleq \beta\left(\beta + \frac{3}{2}\right)$$

and

$$\beta_2 \triangleq \beta\left(5\beta + \frac{9}{2}\right).$$

Proof: By Lemma 1, Assumption 4.2 is satisfied. By Lemma 4.1, Assumption 4.1 is satisfied as well. Thus, the proof follows from Theorem 4.3.5 in [2]. The exact expressions for the parameters β_λ and $\bar{h}$ are obtained from the proofs of Theorems 2.4.1 and 2.4.2 therein. ■

An optimal policy π_h^* (in the discretized model) may be computed *offline*. Therefore, the actual choice of the spending rate by a device can be done by using the precomputed function $\pi_h^* : \mathcal{B}_h \to \mathcal{S}_h$. The quantized policies are used to derive numerical results that appear in Section V.

B. Single Node – Bounds and Heuristic Policies

We now provide some analytical insights into the behavior of the optimal and the *SG* policies for the stochastic energy model. The following observations apply to both energy storage models.

Observation 5: $\mathbb{E}(Z_{OPT}) \leq U(\mathbb{E}(D))$.

Proof: The energy received in a slot i does not exceed $D(i)$, and the overall expected amount does not exceed $K \cdot \mathbb{E}(D)$. The concave objective function U is maximized when the energy is spent uniformly, thus for the total

expected energy $K \cdot \mathbb{E}(D)$, the utility is maximized for energy spending rate $s(i) = [K \cdot \mathbb{E}(D)]/K = \mathbb{E}(D)\ \forall\ i$. Hence $\mathbb{E}(Z_{OPT})$ is bounded as $K \cdot (1/K) \cdot U(\mathbb{E}(D)) = U(\mathbb{E}(D))$. ■

Observation 6: $\mathbb{E}(Z_{SG}) = \mathbb{E}(U(Q))$.

Proof: $\mathbb{E}(Z_{SG}) = \lim_{K\to\infty} \frac{1}{K} \sum_i U(Q(i)) = \mathbb{E}(U(Q))$. ■

We also consider *energy storage state-based* policies, namely the *THR* and the *SL* policies.

- *THR* policy: for a set of storage state thresholds $[B_1, B_2, ...B_T]$ and a set of constants spending rates $[s_1, s_2, ...s_T]$, $s_{THR}(i) \leftarrow 0\ \forall\ B(i) \leq B_1; s_{THR}(i) \leftarrow s_1\ \forall\ B_1 < B(i) \leq B_2; ...; s_{THR}(i) \leftarrow s_T\ \forall\ B(i) > B_T$. That is, for example, for $T = 1$, the *THR* is an ON-OFF policy, and for $T = 2$ is is a bi-level policy.
- *SL* policy: $s_{SL}(i) \leftarrow \alpha_{SL} \cdot [B(i)/C]$ for some parameter α_{SL}.

These policies require choosing parameters, and the policies' performance heavily depends on the choice of the parameters. For policies relying on a small parameter set, simple brute-force algorithms can be used to select the best ones. Consider, for example, a *THR* policy with $T = 1$. A simple algorithms to find the best values for s_1 and B_1 is as follows. For each possible B_1, the algorithms considers all feasible values of s_1, and for each $\{B_1, s_1\}$ combination the algorithm calculates the transition probabilities, determines the stationary probabilities of the states, and calculates Z, choosing the $\{B_1, s_1\}$ combination that maximizes Z. For every state in the state space the algorithm needs to compute transition probabilities, and the resulting stationary storage state probabilities; however, the state space the algorithm considers is relatively small, $O(|C/h|^2)$. In the similar manner, the *SL* policy parameter α_{SL} can be computed by going through at most $O(|C/h|)$ possible α_{SL} values.

Section V demonstrates the performance of different policies using real-world traces.

C. Link Model

The MDP formulation can be extended to a *link* (u, v) as follows. We let $D(i) \triangleq (D_u(i), D_v(i))$ denote the energy harvested in slot i by both devices. We let $D \triangleq (D_u, D_v)$ denote the "representative" variable for $D(i)$ and p_D denote its pdf. In this case, p_D is a *joint* pdf of D_u and D_v. The state space of the MDP is $\mathcal{B} = [0, C]^2$, and the action space at state $b = (b_u, b_v) \in \mathcal{B}$ is given by $\mathcal{S}(b) \triangleq \{(r_u, r_v) :\ c_{\text{tx}} r_u + c_{\text{rx}} r_v = s_u \leq b_u,\ c_{\text{tx}} r_v + c_{\text{rx}} r_u = s_v \leq b_v\}$. The goal is to find an optimal policy that maximizes the average utility $\lim_{K\to\infty} \mathbb{E}_\pi \left(\sum_{i=0}^{K-1} U(r_u(i)) + U(r_v(i)) \right) / K$, which is done using methods similar to those of Section IV-A. Also, corresponding discretization bounds can be obtained.

Similarly to the predictable energy model, the *DRC* algorithms can be used with this model. In this case, the *DRC* policies are calculated using the marginal pdfs of D_u and D_v (rather than the joint pdf), and thus do not account for the dependency between D_u and D_v.

V. Numerical and Experimental Results

A. Trace-based Simulation

TABLE II
EXAMINED LIGHT ENERGY TRACES.

Location	Location description	Experiment timeline
O-1	Outdoor, *ECSU* meteostation [1], Elizabeth City, NC.	Jan. 1, 2009 – Dec. 31, 2009
O-2	Outdoor, *HSU* meteostation [1], Arcata, CA.	Jan. 1, 2009 – Dec. 31, 2009
O-3	Outdoor, Las Vegas meteostation [1], Las Vegas, NV.	Jan. 1, 2009 – Dec. 31, 2009
L-1	Indoor South-facing location [8], New York, NY.	Aug. 15, 2009 – Sept. 13, 2010
L-2	Indoor location receiving mostly indoor light [8], New York, NY.	Nov. 13, 2009 – Sept. 9, 2010
L-3	Indoor North-facing location [8], New York, NY.	Nov. 7, 2009 - Sept. 13, 2010
L-4	Indoor South-West-facing location [8], New York, NY.	Nov. 5, 2009 – Sept. 29, 2010

To evaluate the performance of the various policies, we performed an extensive simulation study using traces from *outdoor* locations [1] and from our measurement campaign, in which we recorded indoor light energy traces at a set of locations at Columbia University for more than a year [8]. The traces are available online at enhants.ee.columbia.edu.

The traces we analyzed are listed in Table II. The traces record *irradiance*, power projected onto a unit surface. To convert the irradiance to the power available to the device, we assumed that the devices had solar cells with 10 cm^2 area and 1% efficiency [9]. For the initial and final storage levels we used $B_0 = B_K = C/2$, where $B_0 = B_K$

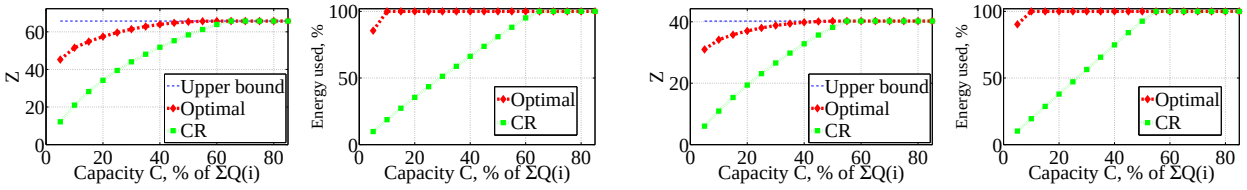

(a) Setup *O-1*. (b) Setup *L-3*.

Fig. 2. Z (left) and % of energy used (right), for a *single node* with a *predictable profile energy*, for the optimal solution and the *CR* policy.

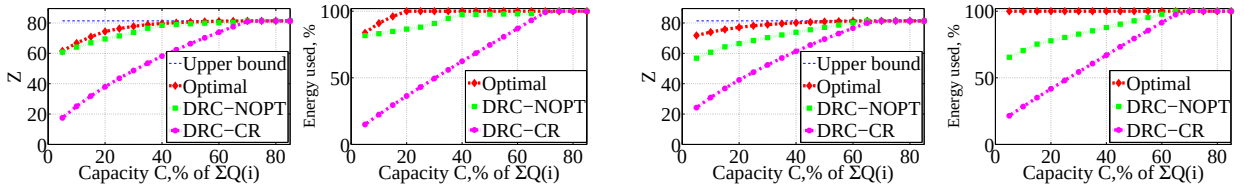

(a) Both C_u and C_v are varied. (b) C_u is kept the same, and C_v is varied.

Fig. 3. Z (left) and energy used (right), for *a link* (u, v) = (*L-1*, *L-2*) with a *predictable profile energy*. The results include the optimal solution, and the *DRC-NOPT* and *DRC-CR* policies.

is chosen to provide for *energy neutrality* [11]. For simplicity, we use $c_{\text{tx}} = c_{\text{rx}}$, and set them to $c_{\text{tx}} = c_{\text{rx}} = 0.5$ nJ/bit [9].

For a *single node* with a *predictable profile energy model*, Fig. 2 illustrates the optimal solution and the performance of the *CR* policy, for energy profiles of two different setups, and shows the upper bound derived in Observation 1. It can be seen that this bound is tight for large C. In our numerical results, the actual ratio between the *CR* solution and the optimal solution is substantially lower than the approximation ratio given in Observation 3. Separately, we evaluate the performance of the *SG* policy. We observe that for *outdoor* setups *O-1* - *O-3*, *SG* policy results in node downtimes T of $0.47 - 0.49$, which corresponds to the overall expected duration of hours of darkness in outdoor environments. For *indoor* setups *L-1* - *L-4* located, the downtimes T are between 0.22 and 0.52.

For a *link* with a *predictable profile model*, we use light energy traces *concurrently recorded in nearby locations*. Fig. 3 illustrates the optimal solution and the performance of the *DRC-CR* policy for a link (u, v) = (*L-1*, *L-2*). Fig. 3(a) shows the case in which both $C_u(i)$ and $C_v(i)$ are varied, while Fig. 3(b) shows the case in which $C_v(i)$ is varied and $C_u(i)$ is kept constant. We note that the *DRC-NOPT* obtains results that are close to the optimal solution in the first case but not in the second case. Separately, we studied the *DRC-SG* policy, and have noticed that for the traces examined , T^L is mostly relatively close to the lower bound derived in Observation 4. For example, for a link (*L-1*, *L-2*), $\max(T_u, T_v) = 0.52$, and $T^L_{u,v} = 0.57$, and for a link (*L-2*, *L-3*), $\max(T_u, T_v) = 0.52$, and $T^L = 0.64$.

For a *single node* and the *stochastic model*, Fig. 4 shows the optimal solution and the solutions obtained by the *SL* and *THR1* (*THR* with one threshold) policies. The policies were evaluated using an empirical pdf of the diurnal energy of setup *L-1*. The calculations of the optimal solutions rely on discretization procedure described in Section IV-A. We can see that for this setup the performance of the *SL* policy is very close to optimal.

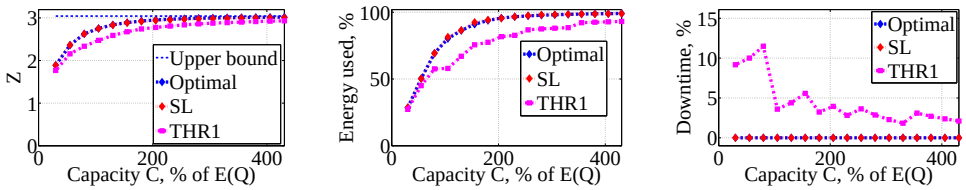

Fig. 4. Z, % of energy spent, and the % of downtime under the optimal solution and the *SL* and *THR1* (ON-OFF) policies, for setup *L-1*.

B. Testbed Experimental Results

To evaluate the performance of the policies in realistic environments, we also used the testbed of energy harvesting devices that we have recently developed [7]. In this testbed, the devices harvest the energy from *indoor light*, and adjust their communication parameters accordingly. We implemented the *CR*, *SG*, and *LS* single-node policies. We also implemented the *DRC* algorithms that can be used with any single-node policy. Testbed implementation allows us to examine the behavior of various policies with *widely varying* and *controlled* energy sources.

For example, under the *DRC-SG* policy, we examined the effect on the performance of the dependence on $Q_u(i)$ and $Q_v(i)$. With strongly correlated $Q_u(i), Q_v(i)$ (i.e., harvesting the energy of the same source), similarly to the light energy traces examined above, $T^L_{u,v}$ is close to the lower bound derived in Observation 4. However, when $Q_u(i)$ and $Q_v(i)$ were independent (i.e., u and v positioned next to two lamps controlled by different people), $T^L_{u,v}$ was closer to the *upper bound* derived in Observation 4. For example, for $T_u = 0.65$ and $T_v = 0.55$, the link downtime $T^L_{u,v}$ was 0.98. Namely, while both u and v had substantial amounts of energy, the data rate on (u, v) was extremely low.

VI. Conclusions

In this work we analyzed and evaluated numerically and experimentally a number of *simple* energy allocation policies for the predictable profile model and the stochastic model. Our analysis applies to linear and non-linear storage models. However, due to the problems' complexity, our analysis applies only to a node or to a link. Most algorithms that were developed for a network are *too complex* for resource-constrained nodes. Therefore, we plan to develop simple algorithms for a network. However, the curse of dimensionality makes it challenging to directly extend the examined stochastic models to larger scenarios, and therefore, approximate solution techniques should be applied (such as Approximate Linear Programming as in [4] and [3]). Moreover, we have so far assumed that the harvested energy is a *stationary* (i.i.d.) process. However, in many environments the energy characteristics changes with time, making *non-stationary models* a better fit. In such cases, the appropriate model is an MDP with *non-homogeneous* or *changing* transition function. Since the changes in the distribution cannot be known in advance, *online learning algorithms* (such as [19]) may be used in such cases.

VII. Acknowledgements

This work was supported in part by the Vodafone Americas Foundation Wireless Innovation Project, NSF grants CNS-0916263, CCF-0964497, and CNS-10-54856, and DHS Task Order #HSHQDC-10-J-00204. We thank Hao Wang and Zainab Noorbhaiwala for their assistance with the experimental work.

References

[1] "Measurement and Instrumentation Data Center, National Renewable Energy Laboratory (NREL), US DOE," www.nrel.gov/midc/.

[2] C.-S. Chow, "Multigrid algorithms and complexity results for discrete-time stochastic control and related fixed-point problems," Ph.D. dissertation, MIT, Cambridge, MA, 1990.

[3] D. De Farias and B. Van Roy, "Approximate linear programming for average-cost dynamic programming," in *Advances in Neural Information Processing Systems 15 (NIPS 03)*, 2003.

[4] ——, "A cost-shaping linear program for average-cost approximate dynamic programming with performance guarantees," *Math. Oper. Res.*, vol. 31, no. 3, pp. 597–620, 2006.

[5] K.-W. Fan, Z. Zheng, and P. Sinha, "Steady and fair rate allocation for rechargeable sensors in perpetual sensor networks," in *Proc. ACM SenSys'08*, Nov. 2008.

[6] M. Gatzianas, L. Georgiadis, and L. Tassiulas, "Control of wireless networks with rechargeable batteries," *IEEE Trans. Wireless. Comm.*, vol. 9, no. 2, pp. 581–593, 2010.

[7] M. Gorlatova, T. Sharma, D. Shrestha, E. Xu, J. Chen, A. Skolnik, D. Piao, P. Kinget, I. Kymissis, D. Rubenstein, and G. Zussman, "Prototyping Energy Harvesting Active Networked Tags (EnHANTs) with MICA2 motes," in *Proc. IEEE SECON'10*, June 2010.

[8] M. Gorlatova, A. Wallwater, and G. Zussman, "Networking low-power energy harvesting devices: Measurements and algorithms," in *Proc. IEEE INFOCOM'11, to appear*, Apr. 2011.

[9] M. Gorlatova, P. Kinget, I. Kymissis, D. Rubenstein, X. Wang, and G. Zussman, "Challenge: ultra-low-power energy-harvesting active networked tags (EnHANTs)," in *Proc. ACM MobiCom'09*, Sept. 2009.

[10] J. Gummeson, S. S. Clark, K. Fu, and D. Ganesan, "On the limits of effective micro-energy harvesting on mobile CRFID sensors," in *Proc. ACM MobiSys'10*, June 2010.

[11] A. Kansal, J. Hsu, S. Zahedi, and M. B. Srivastava, "Power management in energy harvesting sensor networks," *ACM Trans. Embedded Comput. Syst.*, vol. 6, no. 4, 2007.

[12] K. Kar, A. Krishnamurthy, and N. Jaggi, "Dynamic node activation in networks of rechargeable sensors," *IEEE/ACM Trans. Netw.*, vol. 14, no. 1, pp. 15–26, 2006.
[13] L. Lin, N. Shroff, and R. Srikant, "Asymptotically optimal energy-aware routing for multihop wireless networks with renewable energy sources," *IEEE/ACM Trans. Netw.*, vol. 15, no. 5, pp. 1021–1034, 2007.
[14] R.-S. Liu, P. Sinha, and C. E. Koksal, "Joint energy management and resource allocation in rechargeable sensor networks," in *Proc. IEEE INFOCOM'10*, Mar. 2010.
[15] D. Niyato, E. Hossain, and A. Fallahi, "Sleep and wakeup strategies in solar-powered wireless sensor/mesh networks: performance analysis and optimization," *IEEE Trans. Mobile Comput.*, vol. 6, no. 2, pp. 221–236, Feb. 2007.
[16] D. Noh and T. Abdelzaher, "Efficient flow-control algorithm cooperating with energy allocation scheme for solar-powered WSNs," *Wireless Comm. and Mobile Comput.*, 2010.
[17] M. L. Puterman, *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. New York, NY: Wiley, 1994.
[18] C. Vigorito, D. Ganesan, and A. Barto, "Adaptive control of duty cycling in energy-harvesting wireless sensor networks," in *Proc. IEEE SECON'07*, June 2007.
[19] J. Y. Yu and S. Mannor, "Online learning in Markov decision processes with arbitrarily changing rewards and transitions," in *Proc. IEEE Game Theory for Networks (GAMENETS)*, May 2009.