Networking Low-Power Energy Harvesting Devices: Measurements and Algorithms

Maria Gorlatova*, Aya Wallwater†, Gil Zussman*
*Electrical Engineering, †Industrial Engineering and Operations Research
Columbia University, New York, NY, 10027
Email: {mag2206, aw2589}@columbia.edu, gil@ee.columbia.edu

***Abstract*—Recent advances in energy harvesting materials and ultra-low-power communications will soon enable the realization of networks composed of energy harvesting devices. These devices will operate using very low ambient energy, such as indoor light energy. We focus on characterizing the energy availability in *indoor environments* and on developing energy allocation algorithms for energy harvesting devices. First, we present results of our *long-term indoor radiant energy measurements*, which provide important inputs required for algorithm and system design (e.g., determining the required battery sizes). Then, we focus on algorithm development, which requires non-traditional approaches, since energy harvesting shifts the nature of energy-aware protocols *from minimizing* energy expenditure *to optimizing* it. Moreover, in many cases, different energy storage types (*rechargeable battery* and a *capacitor*) require different algorithms. We develop algorithms for determining *time fair* energy allocation in systems with *predictable* energy inputs, as well as in systems where energy inputs are *stochastic*.**

***Index Terms*—Energy harvesting, ultra-low-power networking, indoor radiant energy, measurements, energy-aware algorithms.**

I. Introduction

Recent advances in the areas of solar, piezoelectric, and thermal energy harvesting [34], and in ultra-low-power wireless communications [41] will soon enable the realization of perpetual *energy harvesting wireless devices*. When networked together, they can compose rechargeable sensor networks [23], [35], [45], networks of computational RFIDs [19], and Energy Harvesting Active Networked Tags (EnHANTs) [15]–[17]. Such networks will find applications in various areas, and therefore, the wireless industry is already engaged in the design of various devices (e.g., [9]).

In this report we focus on devices that harvest *environmental light energy*. Since there is a *3 orders of magnitude* difference between the light energy available indoor and outdoor [17], [36], significantly different algorithms are required for different environments. However, there is lack of data and analysis regarding the energy availability in such environments. Hence, over the past 16 months we have been conducting a first-of-its-kind measurement campaign that enables *characterizing the energy availability in indoor environments*. We describe the results and show that they provide insights that can be used for the development of energy-harvesting-aware algorithms and systems.

Clearly, there has been an extensive research effort in the area of energy efficient algorithms for sensor networks and for wireless networks in general. However, for devices with renewable energy sources, fundamentally different problems arise. Hence, in the second part of the paper we focus on developing algorithms for determining the energy spending rates and the data rates in various scenarios.

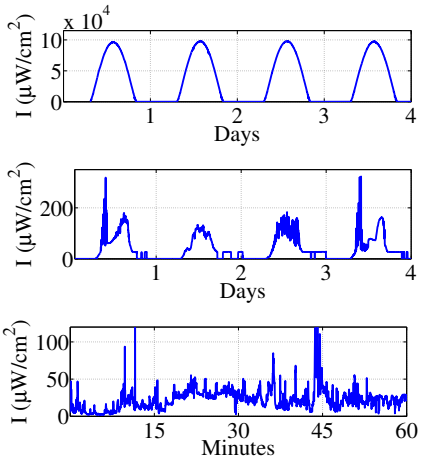

Fig. 1. Examples of different light energy sources: (top) *predictable profile* (Las Vegas, NV [5], outdoors), (middle) *partially-predictable profile* (New York, NY, a static indoor device), and (bottom) *stochastic behavior* (New York, NY, a mobile device in Times Square at *nighttime*).

To describe our contributions, we introduce below several dimensions of the algorithm design space:

- **Environmental energy model**: predictable and partially-predictable energy profile, stochastic process, and model-free.
- **Energy storage type**: battery and capacitor.
- **Ratio of energy storage capacity to energy harvested**: large to small.
- **Time granularity**: sub-seconds to days.
- **Problem size**: stand-alone node, node pair (link), cluster, and multihop network.

The combinations of values along these dimensions induce several "working points", some of which have been studied recently (see Section II).

A. Environmental Energy Models

The model representing harvested energy depends on various parameters such as the energy source (e.g., solar, kinetic), the properties of the environment, and the device's behavior (stationary, semi-stationary, or mobile). Fig. 1 provides examples of radiant (light) energy sources in different settings.

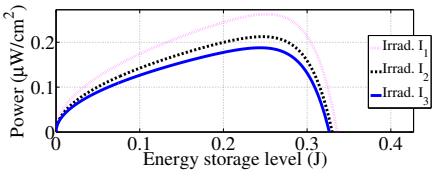

Fig. 2. An example of harvested power vs. storage curves for a capacitor-based light energy harvesting system.

In Fig. 1(a) the energy availability is time-dependent and predictable. On the other hand, in Fig. 1(b) that corresponds to an *indoor* environment, it is time-dependent and periodic, but harder to predict. Time-dependent and somewhat periodic behaviors (along with inputs such as weather forecasts) would allow to develop an *energy profile* [13], [24]. We will refer to ideal energy profiles that accurately represent the future as *predictable profiles*, and to those that are not accurate as *partially-predictable profiles*.

Energy behavior that does not warrant a time-dependent profile appears in Fig. 1(c), which shows the irradiance recorded by a *mobile device* carried around Times Square in New York City at *nighttime*. In this case, the energy can be modeled by a *stochastic process*. Other scenarios where stochastic models are a good fit are 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. Finally, in some settings not relying on an energy model (a *model-free* approach) is most suitable.

B. Energy Storage Types - Linear and Non-Linear

In order to operate when not directly powered by environmental energy, energy harvesting devices need *energy storage*: a *rechargeable battery* or a *capacitor*. *Batteries* can be modeled by an ideal *linear* model, where the changes in the energy stored are linearly related to the amounts of energy harvested or spent, or more realistically by considering their chemical characteristics [37]. Use of *capacitors* for storing harvested energy recently started gaining attention [17], [19], [23], [45]. Specifically, capacitor self-discharge (leakage) is considered in [45]. We consider another important aspect of capacitor-based systems: due to the *highly non-linear output versus voltage characteristics*, in a simple system, the amount of power harvested depends both on the amount of energy provided (*irradiance* in the case of light energy harvesting), and on *the amount of energy stored* [19], [30]. The *non-linear* relations are demonstrated in Fig. 2.

C. Storage Capacity, Decision Timescale, and Problem Size

Storage capacity vs. amount of energy harvested – Energy storage capacity can vary from 0.16J for an EnerChips device [2] to 4700J for an AA battery. The environmental energy availability also varies widely, from thousands of J/cm^2/day in sunny outdoor conditions to under 2J/cm^2/day in indoor environments (see Section IV). Different combinations require different algorithmic approaches. For example, when the storage is small compared to the harvesting rate, the algorithms must continuously keep track of the energy levels, to guarantee that the storage is not depleted or that recharging opportunities are not missed. On the other hand, with relatively large storage, simpler algorithms can be used.

Time granularity – Nodes can characterize the received energy and make decisions on timescales from seconds to days. This timescale is related to the storage-harvesting ratio and the environmental energy model.

Problem/Network Size – Energy harvesting affects nodes' individual decisions, pairwise (link) decisions, and behavior of networked nodes (e.g., routing and rate adaptation).

D. Our Contributions

First, we present the results of a *16 month-long indoor radiant energy measurements campaign* and a *mobile outdoor* light energy study that provide important inputs to the design of algorithms. We discuss the energy available in various indoor environments. We also show that in indoor environments, the energy models are mostly partially-predictable and that simple parameters can significantly improve predictions when the time granularity is at the order of days. To the best of our knowledge, this work is the first to present long-term indoor radiant energy measurements (the traces are available at enhants.ee.columbia.edu and will be made available in CRAWDAD [1]).

Second, we consider *predictable energy profiles* and focus on the simple cases of a *single node and a link*. Our objective is *fair allocation of resources along the time axis*. In particular, we use the *lexicographic maximization* and *utility maximization* frameworks to obtain the energy spending rates for a node and the data rates for a link, both for battery-based systems and for *capacitor-based systems*. To the best of our knowledge our work is the first to consider the *nonlinearity* of the capacitor-based system illustrated in Fig. 2. We provide numerical results that demonstrate its effect.

We also consider a *stochastic model* in which the energy inputs are i.i.d. random variables (e.g., a mobile device outdoor) and show how to treat it as a *Markov Decision Process*. We obtain optimal energy spending policies (both for battery-based and capacitor-based systems) for a single node and a node pair (link) that can be pre-computed in advance.

This report is organized as follows. Section II briefly reviews the related work. Section III presents the model and Section IV describes the measurements. Section V and VI describe algorithms for the predictable profile and stochastic models, respectively. Section VII briefly presents the numerical results. We summarize and discuss future work in Section VIII.

II. Related Work

In this section we briefly review related work using the outlined general settings.

Predictable energy profile: In [22], [24], duty cycle adaptations (mostly for *single nodes*) are considered. For a *network*, various metrics are considered including data collection rates [13], end-to-end packet delivery probability [42], data retrieval

TABLE I
NOMENCLATURE.

I	Irradiance (W/cm^2)
H	Irradiation (J/cm^2)
D	Energy harvested given device physical parameters (J)
K	Number of slots
C	Energy storage capacity (J)
B,B_0,B_K	Energy storage state, initial, and final levels (J)
s	Energy spending rate (J/slot)
Q	Effective energy harvested (J)
$\hat{Q}$	Total energy to be allocated (J)
Δ	Quantization resolution (J)
r	Data rate (bits/s)
c_{tx}, c_{rx}	Energetic costs to transmit and to receive (J/bit)
$U(\cdot)$	Utility function

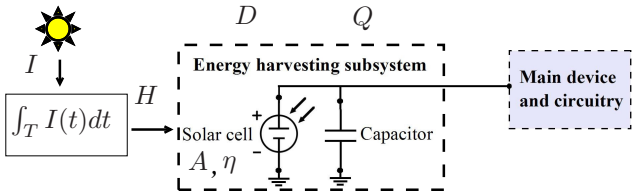

Fig. 3. A schematic diagram of the relationships between energy parameters: irradiance (I), irradiation (H), energy available to a device (D), and energy collected in storage (Q).

rate [43], and routing efficiency [28], [44]. Per-slot *short-term predictions* are assumed in [29].

Partially-predictable energy profile: While considering energy predictable, [24], [29], [32] have provisions for adjustments in cases in which the predictions are not accurate.

Stochastic process: Dynamic activation of energy-harvesting sensors is described in [25]. Admission and power allocation control are developed in [14].

Model-free approach: Duty cycle adjustments for a *single node* (and under the *linear storage* model) are examined in [40]. A *capacitor*-based system is presented and the capacitor leakage is studied in [45].

Additional related work includes a study of the effect of duty cycle adaptations on network-wide parameters [18] and specific considerations for *indoor radiant energy harvesting* [15], [17], [19], [36].

III. Model and Preliminaries

In this work we focus both on *light measurements* and on *resource allocation problems*. The relationships between variables characterizing energy availability are illustrated in Fig. 3. Table I summarizes the notation.

Our measurements record *irradiance*, radiant energy incident onto surface (in W/cm^2), denoted by I. *Irradiation* H_T (in J/cm^2) is the integral of irradiance over a time period T. In characterizing environmental light energy, we are particularly interested in diurnal (daily) environmental energy availability. For $T = 24$ hours, we denote the *daily irradiation* by H_d. The amount of energy (in J) a device with the *given physical characteristics* has access to is denoted by D. For a device with solar cell size A and efficiency η, $D = A\eta H$.

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, respectively. The energy spending rate is denoted by $s(i)$.

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* $D(i)$ *and* $B(i)$: $Q(i) = q(D(i), B(i))$ (see, for example, Fig. 2). We refer to energy storage that is not linear in $D(i)$ (such as a *capacitor*) as *nonlinear storage*. For a *linear* energy storage device (such as a *battery*), $q(D(i), B(i)) = D(i)$ (in general, $Q(i) \leq D(i)$). The 'storage evolution' for the models can be expressed as:

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

We denote the total amount of energy the device is allocating by $\hat{Q}$, where $\hat{Q} = \sum_i Q(i) + (B_0 - B_K)$. For simplicity, some of the developed energy allocation algorithms use *quantized* $B(i)$ and $Q(i)$ values. We denote the quantization resolution by Δ.

We consider the behavior of single nodes and node pairs (links). We denote the endpoints of a link by u and v and their data rates by $r_u(i)$ and $r_v(i)$. For a *single node* we optimize the energy spending rates $s(i)$, which can provide inputs for determining *duty cycle*, *sensing rate*, or *communication rate*. For a *link*, we optimize the communication rates $r_u(i)$ and $r_v(i)$. We denote the costs to transmit and receive a bit by c_{tx} and c_{rx}.

Often the incoming energy varies throughout the day or among different days. We aim to achieve a *time-fair* resource allocation, that is, allocate, as much as possible, the energy in a uniform way with respect to time. We achieve this by using the *lexicographic maximization* and *utility maximization* frameworks. In the former, we *lexicographically maximize* the vector $\{s(0), \ldots, s(K-1)\}$ (for a node), or the vector $\{r_u(0), \ldots, r_u(K-1),\ r_v(0), \ldots, r_v(K-1)\}$ (for a link). Similar approaches have been used to achieve fairness in data generation [13], [26] and in session rate allocations [12]. The network *utility maximization* framework is also well-developed [29], [33], but mostly for fairness among nodes, not across the different time slots. To apply it, $\alpha-$fair functions are used under certain objective functions that will be described in Sections V and VI. $\alpha-$fair functions are the family of *concave* and *non-decreasing* functions parameterized by $\alpha \geq 0$: $U_\alpha(\cdot) = (\cdot)^{1-\alpha}/1-\alpha$, for $\alpha \geq 0$, $\alpha \neq 1$ and $\log(\cdot)$ for $\alpha = 1$. Under our objective function, we use them to achieve $\max - \min$ and proportional fairness [31]. We apply the utility maximization framework to find both the optimal spending rates $s(i)$ and the optimal communication rates $r_u(i)$ and $r_v(i)$.

IV. Characterizing Light Energy

One of the important dimensions of the problem space is environmental energy modeling. Since large-scale *outdoor* solar panels have been used for decades, properties of the Sun's energy were examined in depth [5], [27], [36]. Until recently

TABLE II
LIGHT ENERGY MEASUREMENT SETUPS, AVERAGE DAILY IRRADIATION, AND ACHIEVABLE BIT RATES.

Location index	Location description	Experiment timeline	$\overline{H_d}$ (J/cm^2/day)	$\sigma(H_d)$	r (Kb/s, cont.)
L-1	Students' office. South-facing, 6th floor above ground, windowsill-located setup.	Aug. 15, 2009 – Sept. 13, 2010	1.3	0.72	1.5
L-2	Students' office (same office as setup *L-1*). Setup on a bookshelf far from windows, receiving direct sunlight for a short portion of a day.	Nov. 13, 2009 – Sept. 9, 2010	1.28	0.76	1.5
L-3	Departmental conference room. North-facing, 10th floor above ground, with large unobstructed windows. Windowsill-located setup.	Nov. 7, 2009 - Sept. 13, 2010	63.0	48.0	72.0
L-4	Students' office. South-West-facing, windowsill-located setup.	Nov. 5, 2009 – Sept. 29, 2010	9.2	6.9	7.9
L-5	Students' office (directly under the office of setup *L-1*). Windowsill-located setup.	June 25, 2009 – Oct. 11, 2009	12.3	8.3	13.9
L-6	Students' office. East-facing, often receiving unattenuated reflected outdoor light. Windowsill-located setup.	Feb. 15, 2010 – Sept. 20, 2010	97.3	64.4	112.3
O-1	*Outdoor*: *ECSU* meteostation [5], Elizabeth City, NC.	Jan. 1, 2009 – Dec. 31, 2009	1517	787	1,750
O-2	*Outdoor*: *HSU* meteostation [5], Arcata, CA.	Jan. 1, 2009 – Dec. 31, 2009	1407	773	1,600

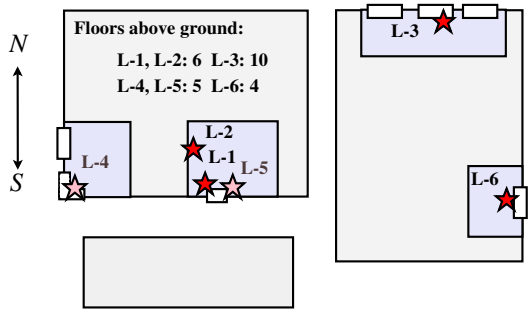

Fig. 4. A schematic diagram of the indoor irradiance measurement locations.

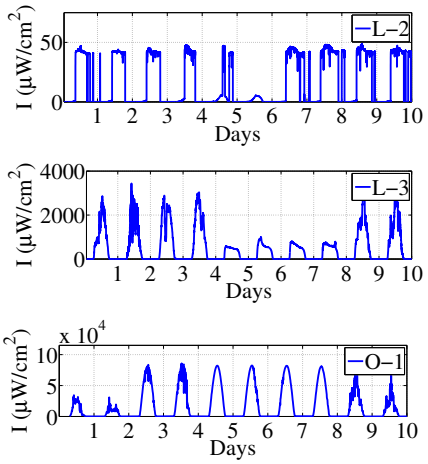

Fig. 5. Sample irradiance measurements in locations *L-2*, *L-3*, and *O-1* (Mar. 2, 2010 - Mar. 12, 2010).

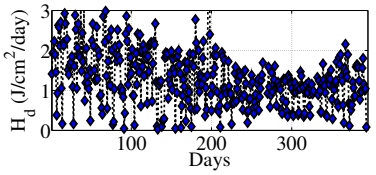

Fig. 6. Long-term daily irradiation (H_d) for setup *L-1* (Aug. 15, 2009 - Sept. 13, 2010).

using *indoor* radiant energy for networking applications was considered impractical, and indoor light was studied mostly in architecture and ergonomics [20], [38]. However, in these domains the important factor is *how humans perceive* the given light (*photometric characterization* – i.e., measurements in Lux) rather than the *energy of the light* (*radiometric characterization*). Photometric measurements by sensor nodes were reported in [3], [19]. Photometric measurements, however, do not provide energetic characterization, and there is a lack of data (e.g., traces) and analysis (e.g., variability, predictability, and correlations) regarding energy availability [36].

To characterize indoor energy availability, since June 2009 we have been conducting a light measurement study in office buildings in New York City. In this study we take long-term measurements of *irradiance* (I, in units W/cm^2) in several indoor locations, and also study a set of shorter-term indoor/outdoor mobile device measurements. Table II provides a summary of the indoor measurement locations. Schematically, they are show in Fig. 4. For the measurements we use TAOS TSL230rd photometric sensors [8] installed on LabJack U3 DAQ devices [4]. These photometric sensors have a high dynamic range, allowing capturing of widely varying irradiance conditions. We verified the accuracy of the sensors with a NIST-traceable [6] Newport 818-UV photodetector [7]. In addition to the indoor measurements, we also analyze a set of outdoor irradiance traces provided by the National Renewable Energy Laboratory (NREL) [5].

The provided measurements and *irradiance traces* can be used to determine the performance achievable by a particular device, for system design (e.g., choosing a suitable energy storage or energy harvesting system component), and for determining which algorithms to use. The traces we have collected can be also used as energy feeds to simulators and emulators. The traces are available at enhants.ee.columbia.edu and will be made available in the CRAWDAD repository [1].

A. Device Energy Budgets and Daily Energy Availability

Sample irradiance measurements (for three setups over the same 10 days) are provided in Fig. 5. One use of the measurements is to determine *energy budgets* for indoor energy harvesting devices. Hence, we calculate the total daily *irradiation* H_d, representing energy incident onto 1cm^2 area over the entire course of a day. Fig. 6 demonstrates the H_d values for setup *L-1*. Table II presents the mean and the standard deviation values, $\overline{H_d}$ and $\sigma(H_d)$, and includes the bit rate r a node would be able to maintain throughout a day when exposed to irradiation H_d. These bit rates are calculated assuming *solar cell efficiency of* $\eta = 1\%$ (i.e., efficiency of

an organic solar cell) and *solar cell size* $A = 10cm^2$. As an energy cost to communicate, 1nJ/bit is used [17].[1] We note that for the different setups, the H_d values vary greatly. The differences are related to office layouts, presence or absence of direct sunlight, as well as the use of shading, windows, and indoor lights.

To predict *daily energy availability* H_d, a node can use a simple *exponential smoothing* approach, calculating a predictor for slot i, $\widehat{H_d}(i)$, as $\widehat{H_d}(1) \leftarrow H_d(0)$, $\widehat{H_d}(i) \leftarrow \alpha \cdot H_d(i-1) + (1-\alpha) \cdot \widehat{H_d}(i-1)$ for α constant, $0 \leq \alpha \leq 1$. The error for such a simple predictor is relatively high. For example, for setup *L-1* the average prediction error is over $0.4\overline{H_d}$, and for setup *L-2* it is over $0.5\overline{H_d}$. For the outdoor datasets the average prediction errors are approximately $0.3\overline{H_d}$.

Improving the energy predictions (for outdoor conditions) using *weather forecasts* has been studied in [27], [39]. We examined whether the H_d values in the *indoor settings* are correlated with the weather data [10]. We determined substantial correlations for some locations. For example, for setup *L-1* the correlation coefficient of the H_d values with the weather data is $r_c = 0.35$ ($p < .001$), and for setup *L-6* it is $r_c = 0.8$ ($p < .001$). This suggests that for some indoor setups the energy predictions may be improved, similar to outdoor environments, by incorporating the weather forecasts into the predictions.

Work week pattern also influences indoor radiant energy in office environments, particularly for setups that do not receive direct sunlight. For setup *L-2*, for example, $\overline{H_d} = 1.63$ J/cm^2 on weekdays, and $\overline{H_d} = 0.37$ J/cm^2 on weekends (it receives, on average, 9.7 hours of office lighting per day on weekdays and under 1 hour on weekends). By keeping separate predictors for weekends and weekdays, the average prediction error for the weekdays is lowered from $0.5\overline{H_d}$ to $0.26\overline{H_d}$.

We also examined correlations between the H_d values of different datasets, and determined statistically significant correlations for a number of setups. For example, for setups *L-1* and *L-2* located in the same room, $r_c = 0.58$ ($p < .001$), and for setups *L-1* and *L-5* facing in the same direction, $r_c = 0.71$ ($p < .001$). This indicates that in a *network* of energy harvesting devices, a device will be able to *infer some information about its peers' energy availability based on its own (locally observed) energy state*.

B. Short Term Energy Profiles

To characterize energy availability at different times of day, we determine the H_T values for different 0.5 hour intervals T, generating *energy profiles* for the setups. Sample energy profiles are shown in Fig. 7, where the left side shows the irradiance curves corresponding to different days overlayed on each other, and the right side shows the $\overline{H_T}$ values, with errorbars representing $\sigma(H_T)$. Due to variations in illumination and occupancy patterns, the energy profiles of different locations can be very different. For example, while setup *L-3* exhibits daylight-dependent variations in irradiance,

[1]The bit rate is calculated as $r = A \cdot \eta \cdot H_d/(3600 \cdot 24)/(10^{-9})$.

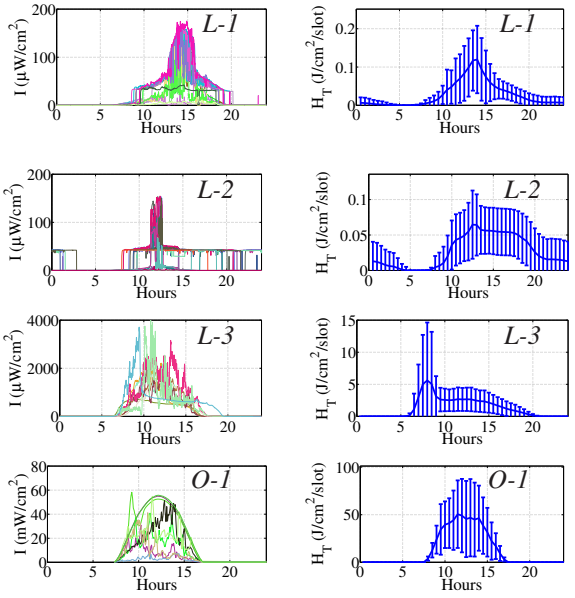

Fig. 7. Sample *energy profiles* for indoor locations *L-1*, *L-2*, *L-3*, and for the outdoor installation *O-1*. Left: irradiance measurements from several different days, overlayed; Right: $\overline{H_T}$ values, with errorbars representing $\sigma(H_T)$.

for setup *L-2* the irradiance is either 0 or 45 μW/cm^2 for most of the day (as this setup receives mostly indoor light). In addition, while for setup *L-2* the lights are often on during late evening hours, for setup *L-3* it is almost never the case. The demonstrated $\sigma(H_T)$ values suggest that these energy inputs generally fall under the *partially predictable profile* energy models.

We have studied whether, similarly to outdoor environments, in the indoor environments the accuracy of the energy profile *for a given day* can be improved when a device has observed some of the incoming energy [11], [27]. We examined correlations between the amount of energy collected in a particular time slot i, $H_T(i)$, and the amount of energy available in some later time slot j, $H_T(j)$ (where $j > i$). The statistically significant r_c values obtained for three different indoor setups are shown on 3D graphs in Fig. 8. We also examined correlations between the amount of energy collected up to a particular time slot j, $\sum_{i \leq j} H_T(i)$, and the energy collected over the subsequent time slots, $\sum_{i>j} H_T(i)$. These correlation coefficients for a number of setups are shown in Fig. 9.

We determined that, in general, are stronger in outdoor settings. For example, for the outdoor setup *O-1* the correlation between the energy received in the 21st time slot (10:30-11:00 AM) and in the 33rd time slot (16:30-17:00 PM) is $r_c = 0.5$ (p<.001), while for the indoor setup *L-1* it is $r_c = 0.2$ (p<.001). For the outdoor setup *O-1*, the correlation between the amount of energy received before 08:00 AM, $\sum_{i \leq 16} H_T(i)$, and the amount of energy received after 08:00 AM, $\sum_{i>16} H_T(i)$, is $r_c = 0.77$ (p<.001), while for the indoor setup *L-3* it is $r_c = 0.31$ (p<.001). These results suggest that profile prediction techniques developed for outdoor systems may be extended to indoor environments, but their performance indoors is likely to be worse.

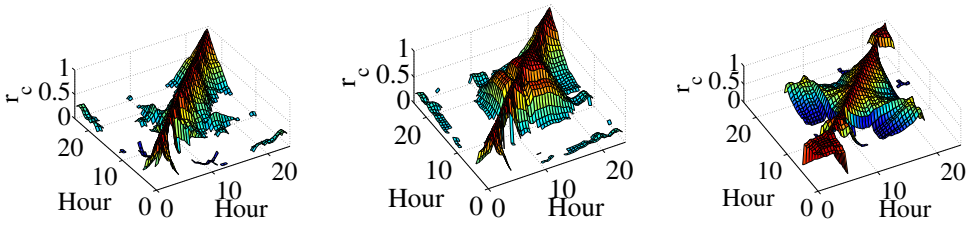

(a) *L-1* (b) *L-2* (c) *L-3*

Fig. 8. Correlations r_c between the amount of energy collected in a particular 0.5 hour time slot, and the amount of energy available in other time slots.

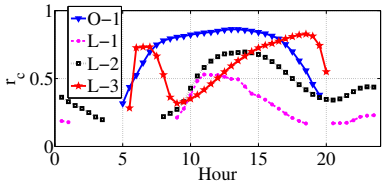

Fig. 9. Correlations between the amounts of energy collected up to a particular time slot, and the amount of energy collected over the subsequent slots.

TABLE III
SUMMARY OF MOBILE MEASUREMENTS.

Meas. index	Measurement description	Date	Time
M-1	Pedestrian walking around university campus (indoor and outdoor environments) carrying a sensor.	4/5/2010	13:06-14:10
M-2	Commuting on public transit, sensor attached to a backpack, measurements outdoors, indoors (subway, train, office).	7/13/2010	15:02–16:42
M-3	Car-based roadtrip, sensor attached to the dashboard.	7/16/2010	12:26–15:23
M-4	Car-based errand running, sensor attached to the dashboard.	7/17/2010	14:48–17:04
M-5	Car-based errand running, sensor attached to the dashboard.	7/18/2010	09:58–12:05
M-6	Pedestrian walking in New York City at nighttime (Theater District and Times Square), sensor attached to a backpack.	7/22/2010	20:02–21:36

C. Mobile Measurements

We have also conducted shorter-term experiments for *mobile devices*. Table III provides a summary of the measurements conducted. A sample irradiance trace for a device carried around Times Square in New York City at *nighttime* was shown in Fig. 1(c). Fig. 10 demonstrates an irradiance trace of a device carried around a set of indoor and outdoor locations (note the log scale of the y-axis) during mid-day on a sunny day (measurements set *M-1*). These measurements highlight the disparity between the light energy available indoors and outdoors. For example, inside a lab, the irradiance was 70μW/cm^2, while in sunny outdoor conditions it was 32mW/cm^2. Namely, the outdoor to indoor energy ratio was *more than 450 times*. In general, we observed that mobile devices' energy levels are poorly predictable and could in some cases be represented by *stochastic* energy models.

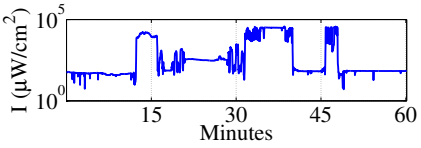

Fig. 10. Irradiance measurements recorded by a mobile device: a mix of indoor and outdoor conditions (note the log scale of the y-axis).

V. PREDICTABLE ENERGY PROFILE

In this section we consider the *predictable profile* energy model (similar to the models studied in [13], [24], [32]). We formulate optimization problems that apply to both *linear* and *nonlinear* energy storage[2] for a *single node* and for pair-wise nodes (*link*), and introduce algorithms for solving them.

A. Single Node: Optimizing Energy Spending

To achieve smooth energy spending for a node, we formulate the following problems using *utility maximization* and *lexicographic maximization* frameworks.

Time Fair Utility Maximization (TFU) Problem:

$$\max_{s(i)} \sum_{i=0}^{K-1} U(s(i)) \tag{2}$$

s.t.:

$$s(i) \leq B(i) \ \forall \ i \tag{3}$$

$$B(i) \leq B(i-1) + Q(i-1) - s(i-1) \ \forall \ i \geq 1 \tag{4}$$

$$B(i) \leq C \ \forall \ i \tag{5}$$

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

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

Recall that $Q(i) = q(D(i), B(i))$. Constraint (3) ensures that a node does not spend more energy than it has stored, (4) and (5) represent the storage evolution dynamics, and (6) sets the initial and final storage levels to B_0 and B_K.

Time Fair Lexicographic Assignment (TFLA) Problem:

$$\text{Lexicographically maximize: } \{s(0), ..., s(K-1)\} \tag{8}$$

s.t.: constraints (3) – (7).

For quantized energy inputs and energy storage, the *TFU* problem can be solved by the following dynamic programming-based algorithm:

In the *TFR* algorithm, for every $\{i, B(i)\}$, we determine $h(i, B(i)) = \max_{s(i) \leq B(i)} [U(s(i)) + h(i+1, \min(B(i) + Q(i) -$

[2]Recall that a *linear* energy storage model applies to a *battery* and that a *non-linear* energy storage model may represent a *capacitor*.

Algorithm 1 Time Fair Rate Assignment (TFR).

$h(i,B) \leftarrow -\infty$, $s(i) \leftarrow 0 \ \forall \ i < K, \forall \ B$;
$h(K, B(K)) \leftarrow -\infty \ \forall \ B(K) < B_K$;
$h(K, B(K)) \leftarrow 0 \ \forall \ B(K) \geq B_K$;
for $i = K-1; i \geq 0; i--$; **do**
  for $B = 0; B \leq C; B \leftarrow B + \Delta$; **do**
    for $s = 0; s \leq B; s \leftarrow s + \Delta$; **do**
      $\hat{s} \leftarrow s$; $\hat{h} \leftarrow U(\hat{s}) + h(i+1, \min(B + q(D(i), B) - \hat{s}, C)$;
      if $\hat{h} > h(i,B)$ **then**
        $h(i,B) \leftarrow \hat{h}$; $s(i) \leftarrow \hat{s}$;
return $h(0, B_0)$, and associated $s(i) \ \forall \ i$

$s(i), C))]$. Going "backwards" from $i = K-1$, we thus determine a vector $\{s(0), ..., s(K-1)\}$ that maximizes $h(0, B_0)$; this vector is the optimal energy allocation. Recall that we denote the energy quantization resolution by Δ. The running time of the *TFR* algorithm is $O(K \cdot [C/\Delta]^2)$.

For *linear storage* ($q(D(i), B(i)) = D(i)$), we refer to the *TFU* and the *TFLA* problems as **TFU-LIN** and **TFLA-LIN**.

Lemma 1: The optimal solutions to the *TFU-LIN* problem and the *TFLA-LIN* problem are equal.

Proof: Given in Appendix I. ■

For solving the *TFLA-LIN* and the *TFU-LIN* problems, we develop the *Progressive Filling (PF)* algorithm (Algorithm 2), inspired by the algorithms for $\max - \min$ fair flow control [12]. The *PF* algorithm starts with $s(i) \leftarrow 0 \ \forall \ i$, and iterates through the slots, increasing the $s(i)$ value of each slot by Δ on every iteration. The algorithm verifies that increasing $s(i)$ does not result in shortage of energy for other slots, or in the lack of final energy B_K. An $s(i)$ value is increased only when it does not interfere with the spending in slots with smaller $s(i)$ values, thus the resulting solution is $\max - \min$ fair. At each step of the *PF* algorithm, the verification subroutine of complexity $O(K)$ is executed. Recall that $\hat{Q} = \sum_i Q(i) + (B_0 - B_K)$. The algorithm takes $\hat{Q}/\Delta$ spending increase steps, and K additional steps to 'fix' the slots. Thus, the *PF* algorithm runs in $O(K \cdot [K + \hat{Q}/\Delta])$ time. Assuming that K is small compared to $\hat{Q}/\Delta$, for C and $\hat{Q}$ that are on the same order, the *PF* algorithm is faster than the *TFR* algorithm.

Finally, when the *energy storage is large compared to the energy harvested*, the *TFLA-LIN and TFU-LIN* problems can be solved easily. Let $\tilde{s}(i) = \hat{Q}/K \ \forall \ i$, and let $\tilde{B}(i) = [\sum_{j=0}^{i-1} Q(j)] - (i-1) \cdot \tilde{s}(i) \ \forall \ i \ 1 \leq i \leq K$. We define the following sets of conditions.

Definition 1: The **LS Conditions** hold, if $B_0 \geq |\min_{1 \leq i \leq K} \tilde{B}(i)|$ and $C - B_0 \geq \max_{1 \leq i \leq K} \tilde{B}(i)$.

Definition 2: The **LS-gen Conditions** hold, if $B_0 \geq [\sum_i Q(i)] \cdot (1 - 1/K)$ and $C - B_0 \geq [\sum_i Q(i)] \cdot (1 - 1/K)$.

Lemma 2: When the *LS conditions* or the *LS-gen conditions* hold, the optimal solution to the *TFLA-LIN* and the *TFU-LIN* problems is $s(i) = \hat{Q}/K \ \forall \ i$.

Proof: Given in Appendix II. ■

Algorithm 2 Progressive Filling (PF).

$A_{fix} \leftarrow \varnothing$; $s(i) \leftarrow 0 \ \forall \ i$;
while $\overline{A_{fix}} \neq \varnothing$ **do**
  for $i = 0; \ i \leq K-1; \ i++$; **do**
    if $i \in \overline{A_{fix}}$ **then**
      $\tilde{s}(j) \leftarrow s(j) \ \forall \ j \in [0, K-1]$; $\tilde{s}(i) \leftarrow \tilde{s}(i) + \Delta$;
      $valid \leftarrow$ **check_validity**$(\tilde{s})$;
      if $valid == TRUE$ **then** $s(i) \leftarrow \tilde{s}(i)$;
      else $A_{fix} := A_{fix} \cup i$;
function check_validity$(\tilde{s})$:
$B(i) \leftarrow 0 \ \forall \ i$; $B(0) \leftarrow B_0$; $valid \leftarrow TRUE$;
for $i = 1; i \leq K; i++$; **do**
  $B(i) \leftarrow \min(C, \ B(i-1) + Q(i-1) - \tilde{s}(i-1))$;
  if $\tilde{s}(i) > B(i)$ **then** $valid \leftarrow FALSE$;
if $B_K < B(K)$ **then** $valid \leftarrow FALSE$;
return $valid$

Verifying that the *LS Conditions* (or the *LS-gen Conditions*) hold and determining the corresponding optimal policy is computationally inexpensive.[3] Thus, in this section we demonstrated a general algorithm (for *linear* and *non-linear* storage) of a relatively high complexity, a faster algorithm for *linear storage*, and a very fast algorithm for *large linear storage*.

B. Link: Optimizing Data Rates

For a *link*, we extend the above optimization problems as follows:

Link Time Fair Utility Maximization (LTFU) Problem:

$$\max_{r_u(i), r_v(i)} \sum_{i=0}^{K-1} [U(r_u(i)) + U(r_v(i))] \tag{9}$$

$$\text{s.t.:} \quad c_{\text{tx}} r_u(i) + c_{\text{rx}} r_v(i) \leq s_u(i) \tag{10}$$

$$c_{\text{tx}} r_v(i) + c_{\text{rx}} r_u(i) \leq s_v(i) \tag{11}$$

$$u, v: \ \text{constraints (3)} - \text{(7)}.$$

Link Time Fair Lexicographic Assignment (LTFL) Problem:

Lexicographically maximize:

$$\{r_u(0), ..., r_u(K-1), r_v(0), ..., r_v(K-1)\} \tag{12}$$

$$\text{s.t.:} \quad (10), (11); \ u, v: \ \text{constraints (3)} - \text{(7)}.$$

Since the optimal solution to the *LTFL* problem is $\max - \min$ fair, it assigns the data rates such that $r_u(i) = r_v(i) \ \forall \ i$. Thus, the *LTFL* problem can be restated as:

$$\text{Lexicographically maximize: } \{r(0), ..., r(K-1)\} \tag{13}$$

$$\text{s.t.:} \quad r(i) \cdot (c_{\text{tx}} + c_{\text{rx}}) \leq \min(s_u(i), s_v(i)) \tag{14}$$

$$u, v: \ \text{constraints (3)} - \text{(7)}$$

where $r(i) = r_u(i) = r_v(i)$.

[3]To determine if the *LS Conditions* hold, a node needs to know $\{Q(1), ..., Q(K-1)\}$, while determining if the *LS-gen Conditions* hold requires only the knowledge of $\sum_i Q(i)$. *LS-gen Conditions* can be used, for example, if light energy harvesting nodes characterize their energy availability by the daily irradiation H_d and do not calculate their *energy profiles* (see Section IV-B).

For quantized energy values, the *LTFU* problem can be solved with an extension of the *TFR* algorithm, referred to as **LTFR**. Over all $\{r_u(i), r_v(i)\}$ such that $c_{\text{tx}}r_u(i) + c_{\text{rx}}r_v(i) = s_u(i) \leq B_u(i), c_{\text{tx}}r_v(i) + c_{\text{rx}}r_u(i) = s_v(i) \leq B_v(i)$, the *LTFR* algorithm determines, for each $\{i, B_u(i), B_v(i)\}$,

$$h(i, B_u(i), B_v(i)) = \max[U(r_u(i)) + U(r_v(i)) + h(i+1, \min(B_u(i) + Q_u(i) - s_u(i), C_u), \min(B_v(i) + Q_v(i) - s_v(i), C_v))].$$

Vectors $\{r_u(0), ..., r_u(K-1)\}$ and $\{r_v(0), ..., r_v(K-1)\}$ that maximize $h(0, B_{0,u}, B_{0,v})$ are the optimal. Since this formulation considers all $\{i,\ B_u(i),\ B_v(i)\}$ combinations and examines all feasible rates $r_u(i)$ and $r_v(i)$ for each combination, the overall complexity of the *LTFR* algorithm is $O(K \cdot [C_u/\Delta]^2 \cdot [C_v/\Delta]^2)$.

For *linear* storage, the *LTFL* problem can be solved by an extension of the *PF* algorithm, referred to as the **LPF** algorithm. Similarly to the *PF* algorithm, the *LPF* algorithm goes through all slots and increases the slots' allocation by Δ when an increase is feasible. Unlike the *PF* algorithm, however, the *LPF* algorithm allocates the energy of both nodes u and v. The running time of the *LPF* algorithm is $O(K \cdot [K + (\hat{Q}_u + \hat{Q}_v)/\Delta])$.

Solving the *LTFU* or the *LTFL* problems directly may be computationally taxing for small devices with limited capabilities. Instead, the nodes may use the following low complexity heuristic algorithms, which do not require extensive exchange of information.

Decoupled Rate Control (DRC) algorithms: Initially, nodes u and v determine *independently from each other* their energy spending rates $s_u(i)$ and $s_v(i)$ for every slot i (i.e., using the *PF* algorithm). Then, for each slot i, under constraints (10) and (11), the nodes obtain a solution to $\max_{r_u(i), r_v(i)} U(r_u(i)) + U(r_v(i))$ if the *LTFU* problem is being solved (**LTFU-DRC** algorithm), and to $\max r(i)$ if the *LTFL* problem is being solved (**LTFL-DRC** algorithm). These subproblems (each considers a single slot i) can be easily solved. For the *LTFL-DRC* algorithm, for example, due to (14), the subproblem solution is $r(i) = \min(s_u(i), s_v(i))/(c_{\text{tx}} + c_{\text{rx}})$.

For linear storage, when *the storage is large compared to the energy harvested* for both u and v (that is, when the *LS conditions* hold for both nodes), solving a *single instance* of the *LTFU-DRC* or *LTFL-DRC* problem obtains the overall solution. Moreover, as shown in the Lemma below, in this case *the DRC solution is optimal*. Thus, in such case the optimal solution can be calculated with little computational complexity.

Lemma 3: If the *LS conditions* hold for nodes u and v, the *LTFU-DRC* and *LTFL-DRC* algorithms obtain the optimal solutions to the *LTFU* and the *LTFL* problems, respectively.

Proof: Given in Appendix III. ∎

In Section VII we provide numerical results demonstrating the rates $\{r_u(i), r_v(i)\}$ obtained by using the *DRC* algorithms to solve the *LTFU* and *LTFL* problems.

VI. Stochastic Energy Models

In this section, we study models in which the energy harvested in a slot is an *i.i.d. random variable* D. For tractability, we assume that D takes one of M discrete values $[d_1, ..., d_M]$ with probability $[p_1, ..., p_M]$. D may represent, for example, the energy harvested by a *mobile* device in a short (seconds or minutes) time slot. For time slots of *days*, it may represent the *daily irradiation* H_d received by a device.[4] We formulate the control problems and determine corresponding policies for a single node and for a node pair (*link*). The formulations apply to *linear* and *nonlinear* energy storage models. For a given distribution of D, the optimal policy needs to be calculated *once*, thus operating according to the optimal policy does not require frequent computations.

Spending Policy Determination (SPD) Problem: For a given distribution of D, determine the energy spending rates $s(i)$ such that:

$$\max_{s(i)} \lim_{K\to\infty} \frac{1}{K} \sum_{i=0}^{K-1} U(s(i)). \tag{15}$$

This *discrete time stochastic control process* is a *Markov Decision Process* (MDP), and can be solved with standard MDP solution techniques. For example, applying *dynamic programming*, we consider a large number of slots K, and going "backwards" from $i = K-1$, for each $\{i, B(i)\}$, determine

$$h(i, B(i)) = \max_{s(i)\leq B(i)} \mathop{\mathbb{E}}_{D}[U(s(i)) + h(i+1, \min[B(i) + q(D(i), B(i)) - s(i), C])] = \max_{s(i)\leq B(i)} [U(s(i)) + \sum_{j=1}^{M} p_{d_j} \cdot h(i+1, \min[B(i) + q(d_j, B(i)) - s(i), C])]. \tag{16}$$

Performing this iterative procedure for a large number of slots K, we obtain, for each energy storage level $B(i)$, a corresponding *stationary* (same for all i) $s(i)$ value that approaches the optimal [21]. Although policy calculations are computationally expensive (the running time of the algorithm is $O([C/\Delta]^2 \cdot M \cdot K)$), such a policy needs to be computed only once for a particular distribution of D.

The MDP formulation can be extended to a *link* as follows.

Link Spending Policy Determination (LSPD) Problem:

$$\max_{r_u(i), r_v(i)} \lim_{K\to\infty} \frac{1}{K} \sum_{i=0}^{K-1} [U(r_u(i)) + U(r_v(i))] \tag{17}$$

Similarly to the *SPD* problem, the *LSPD* problem can be solved with standard approaches to solving MDPs. For example, using dynamic programming, we determine, for each $\{i, B_u(i), B_v(i)\}$,

$$h(i, B_u(i), B_v(i)) = \max \mathop{\mathbb{E}}_{D_u, D_v} [U(r_u(i)) + U(r_v(i)) + h(i+1, min[B_u(i) + q(D_u(i), B_u(i)) - s_u(i), C_u], min[B_v(i) + q(D_v(i), B_v(i)) - s_v(i), C_v])], \tag{18}$$

[4] When the energy storage is relatively large, variations in energy availability *within* a day may be abstracted, and H_d can be used to characterize energy availability.

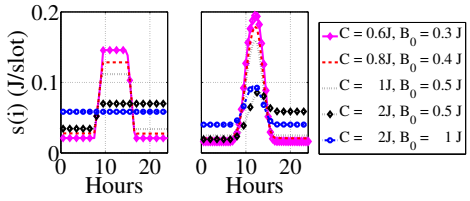

Fig. 11. Energy spending rates $s(i)$: obtained by solving the *TFLA-LIN* and *TFU-LIN* problems (left), and by solving the *TFU* problem for *nonlinear storage* (right).

where the maximization is over all $\{r_u(i), r_v(i)\}$ such that $c_{\text{tx}} r_u(i) + c_{\text{rx}} r_v(i) = s_u(i) \leq B_u(i)$, $c_{\text{tx}} r_v(i) + c_{\text{rx}} r_u(i) = s_v(i) \leq B_v(i)$. This procedure is computationally complex. Similarly to the *SPD* problem, it needs to be solved for a large number of slots K, and has the complexity $O([C_u/\Delta]^2 \cdot [C_v/\Delta]^2 \cdot M_u \cdot M_v \cdot K)$. However, it needs to be computed only once.

The MDP formulations can be easily extended to consider other parameters, such as the cost to change the energy spending rate $s(i)$. In Section VII we will present examples of optimal policies obtained by solving the *LSPD* problem for a linear storage model, and by solving the *SPD* problem for both linear and non-linear storage models.

VII. Numerical Results

This section provides numerical results that demonstrate the use of the algorithms described in Sections V and VI. Measurement traces described in Section IV are used as inputs to the algorithms.

Fig. 11 shows the solutions for the *TFLA* and the *TFU* problems of Section V-A. The energy profile of setup *L-3* (see Fig. 7) was used as an input to the algorithms. The left side of Fig. 11 shows the spending rates $s(i)$ that solve the *TFLA-LIN* and the *TFU-LIN* problems. These spending rates are obtained using the *PF* algorithm. The right side of Fig. 11 shows the solutions of the *TFU* problem with *non-linear energy storage*, where storage state dependency was modeled similarly to the dependency demonstrated in Fig. 2. Such a system (*energy storage with state-dependent inputs*) has not been analyzed before.

Fig. 12 shows the numerical results for the *link* rate determination problems, described in Section V-B. The energy profiles of setups *L-1* and *L-2* (see Fig. 7) were used as inputs to the algorithms. Fig. 12(a) shows the optimal communication rates $\{r_u(i), r_v(i)\}$ obtained by solving the *LTFL* and *LTFU* problems for linear storage. Fig. 12(b) shows the rates $r_u(i)$ and $r_v(i)$ calculated using a simple *LTFU-DRC* algorithm. The *LTFU-DRC* algorithm obtains communication rates $r_u(i)$ and $r_v(i)$ that are similar to those obtained by optimally solving the *LTFU* problem. Fig. 12(c) presents an optimal solution to the *LTFU* problem for *nonlinear energy storage*.

Fig. 13 illustrates the *optimal energy spending policies* obtained by solving the *SPD* problem defined in Section VI. The daily irradiation H_d for setup *L-1* (see Fig. 6) was used as the random variable D. Fig. 13 shows the optimal policies for *linear* and *nonlinear* storage types. Finally, Fig. 14 illustrates

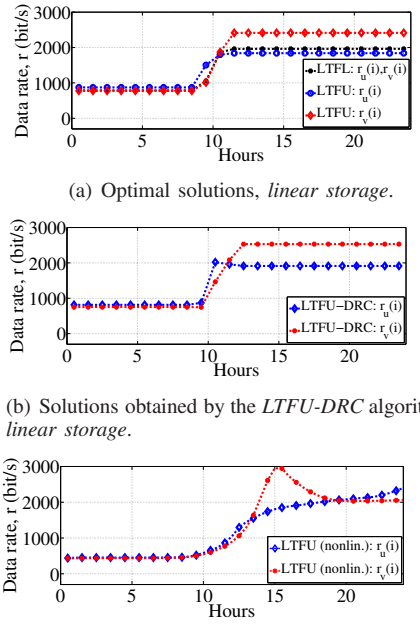

(a) Optimal solutions, *linear storage*.

(b) Solutions obtained by the *LTFU-DRC* algorithm, *linear storage*.

(c) Optimal solution to the *LTFU* problem, *nonlinear storage*.

Fig. 12. Communication rates $r_u(i)$ and $r_v(i)$ obtained by solving the *LTFL* and the *LTFU* problems.

the optimal link rate assignment policy obtained by solving the *LSPD* problem (recall that a policy is computed only once).

VIII. Conclusions and Future Work

Motivated by recent advances in the areas of energy harvesting and ultra-low-power communications, in this work we focus on energy harvesting devices. We describe the first *long-term indoor radiant energy measurements campaign* that provides useful traces, as well as insights into the design of systems and algorithms. We developed algorithms for predictable environment that uniquely determine the spending policies for *linear* and *non-linear* energy storage models. The algorithms for the predictable case also provide insight into the partially-predictable case. We developed algorithms for stochastic environments that can provide nodes with simple pre-computed decisions policies. We used the algorithms to obtain numerical results for various cases.

We covered a few "working points" in the design space described in Section I. Yet, there are still many other working points to study. In particular, although some algorithms have been developed for networks of nodes, most of them are too complex for resource-constrained nodes. We plan to develop simple energy-harvesting-aware algorithms for networks of nodes considering the various other problem dimensions. Moreover, we plan to evaluate these algorithm in an EnHANTs testbed that we are currently building [16].

IX. 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, DHS Task Order #HSHQDC-10-J-00204, and by an NSERC CGS grant. We

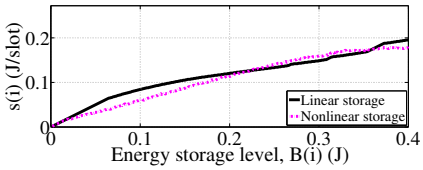

Fig. 13. Optimal energy spending rates $s(i)$ for different energy storage states, obtained by solving the *SPD* problem for *linear* and *nonlinear* storage.

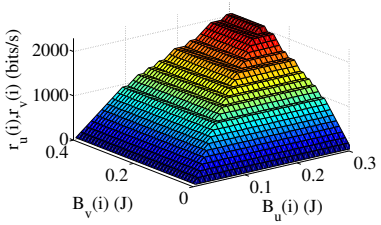

Fig. 14. Optimal communication rates $r_u(i)$ and $r_v(i)$ for different energy storage states, obtained by solving the *LSPD* problem for linear storage.

thank Matthias Bahlke, Enlin Xu and Michael Zapas for their assistance with light energy measurements. We also thank Peter Kinget and John Kymissis for helpful discussions.

REFERENCES

[1] "CRAWDAD – a community resource for archiving wireless data," crawdad.cs.dartmouth.edu/.

[2] "Cymbet EnerChips," www.cymbet.com/content/products.asp.

[3] "Intel Lab data," db.csail.mit.edu/labdata/labdata.html.

[4] "LabJack U3 USB-based multifunction data acquisition and control device," labjack.com/.

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

[6] "National Institute of Standards and Technology (NIST)," www.nist.gov/.

[7] "Newport 818-UV low-power photodetector," www.newport.com/.

[8] "TAOS TSL230rd programmable light-to-frequency converter," www.taosinc.com/.

[9] "Texas Instruments MSP430 Solar Energy Harvesting Development Tool," focus.ti.com/docs/toolsw/folders/print/ez430-rf2500-seh.html.

[10] "Weather Underground Weather History and Data Archive," www.wunderground.com/history/.

[11] M. I. Ali, B. M. Al-Hashimi, J. Recas, and D. Atienza, "Evaluation and design exploration of solar harvested-energy prediction algorithm," in *Design, Automation and Test in Europe (DATE'10)*, Mar. 2010.

[12] D. Bertsekas and R. Gallager, *Data Networks*, 2nd ed. Prentice-Hall, 1992.

[13] 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.

[14] 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.

[15] M. Gorlatova, P. Kinget, I. Kymissis, D. Rubenstein, X. Wang, and G. Zussman, "Energy-harvesting active networked tags (EnHANTs) for ubiquitous object networking," *IEEE Wireless Commun.*, vol. 17, no. 6, pp. 18–25, Dec. 2010.

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

[17] 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.

[18] Y. Gu, T. Zhu, and T. He, "ESC: Energy synchronized communication in sustainable sensor networks," in *Proc. IEEE ICNP'09*, Oct. 2009.

[19] 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.

[20] D. Heil and S. Mathis, "Characterizing free-living light exposure using a wrist-worn light monitor," *Applied Ergonomics*, vol. 33, no. 4, pp. 357–363, 2002.

[21] F. S. Hillier and G. J. Lieberman, *Introduction to Operations Research*, 6th ed. McGraw-Hill, 1995.

[22] J. Hsu, S. Zahedi, A. Kansal, M. Srivastava, and V. Raghunathan, "Adaptive duty cycling for energy harvesting systems," in *Proc. IEEE ISLPED'06*, Oct. 2006.

[23] X. Jiang, J. Polastre, and D. Culler, "Perpetual environmentally powered sensor networks," in *Proc. IEEE IPSN'05*, Apr. 2005.

[24] 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.

[25] 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.

[26] F. P. Kelly, A. Maulloo, and D. Tan, "Rate control in communication networks: shadow prices, proportional fairness and stability," *J. Oper. Res. Soc.*, vol. 49, pp. 237–252, 1998.

[27] M. Kudo, A. Takeuchi, Y. Nozaki, H. Endo, and J. Sumita, "Forecasting electric power generation in a photovoltaic power system for an energy network," *Electrical Eng. in Japan*, vol. 167, no. 4, pp. 16–23, 2009.

[28] 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.

[29] 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.

[30] X. Lu and S. Yang, "Solar energy harvesting for ZigBee electronics," *Sustainability in Energy and Buildings*, pp. 19–27, 2009.

[31] J. Mo and J. Walrand, "Fair end-to-end window-based congestion control," *IEEE/ACM Trans. Netw.*, vol. 8, no. 5, pp. 556–567, 2000.

[32] D. Noh and T. Abdelzaher, "Efficient flow-control algorithm cooperating with energy allocation scheme for solar-powered WSNs," *Wireless Comm. and Mobile Comput.*, 2010.

[33] D. Palomar and M. Chiang, "A tutorial on decomposition methods for network utility maximization," *IEEE J. Sel. Areas Commun.*, vol. 24, no. 8, pp. 1439–1451, 2006.

[34] J. Paradiso and T. Starner, "Energy scavenging for mobile and wireless electronics," *IEEE Pervasive Comput.*, vol. 4, no. 1, pp. 18–27, 2005.

[35] V. Raghunathan, A. Kansal, J. Hsu, J. Friedman, and M. Srivastava, "Design considerations for solar energy harvesting wireless embedded systems," in *Proc. IEEE IPSN'05*, Apr. 2005.

[36] J. Randall, *Designing Indoor Solar Products*, 1st ed. Wiley, 2005.

[37] R. Rao, S. Vrudhula, and D. Rakhmatov, "Battery modeling for energy aware system design," *IEEE Computer*, vol. 36, no. 12, pp. 77–87, 2003.

[38] S. Russel, *The Architecture of Light*. Conceptnine, 2008.

[39] N. Sharma, J. J. Gummeson, D. Irwin, and P. Shenoy, "Cloudy computing: leveraging weather forecasts in energy harvesting sensor systems," in *Proc. IEEE SECON'10*, June 2010.

[40] 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.

[41] D. Wentzloff, F. Lee, D. Daly, M. Bhardwaj, P. Mercier, and A. Chandrakasan, "Energy efficient pulsed-UWB CMOS circuits and systems," in *Proc. IEEE ICUWB'07*, Sept. 2007.

[42] Y. Yang, L. Su, Y. Gao, and T. Abdelzaher, "SolarCode: utilizing erasure codes for reliable data delivery in solar-powered wireless sensor networks," in *Proc. IEEE INFOCOM'10 (mini-conference)*, Mar. 2010.

[43] Y. Yang, L. Wang, D. K. Noh, H. K. Le, and T. F. Abdelzaher, "SolarStore: Enhancing data reliability in solar-powered storage-centric sensor networks," in *Proc. ACM MobiSys'09*, June 2009.

[44] K. Zeng, K. Ren, W. Lou, and P. J. Moran, "Energy aware efficient geographic routing in lossy wireless sensor networks with environmental energy supply," *Wirel. Netw.*, vol. 15, no. 1, pp. 39–51, 2009.

[45] T. Zhu, Z. Zhong, Y. Gu, T. He, and Z.-L. Zhang, "Leakage-Aware Energy Synchronization for Wireless Sensor Networks," in *Proc. ACM MobiSys'09*, June 2009.

I. Appendix: Proof of Lemma 1

Note that $U_\alpha(\cdot)$ functions used as objective functions in *TFU-LIN* problem are concave and twice continuously differentiable for every $\alpha \geq 0$. Here we demonstrate that the claim of Lemma 1 is true for all concave and twice continuously differentiable $U_\alpha(s)$.

First notice the following facts.

Fact 1: The constraint sets of both problems are the same, thus a vector that is a feasible solution to one problem is also a feasible solution to the other.

Fact 2: The nature of constraint set and the utility functions implies that any $\epsilon-$decrease to one of the components of a solution to *TFU-LIN* problem or a solution to *TFLA-LIN* problem yields at most total $\epsilon-$increase to the other components.

Fact 3: Let $y \leq x$ and let f be a twice continuously differentiable concave function. Then for every $\epsilon > 0$, $f(x+\epsilon) + f(y-\epsilon) \leq f(x) + f(y)$.[5]

Let $\bar{s}$ be the optimal solution to the *TFLA-LIN* problem. We want to show that it is also the optimal solution to the *TFU-LIN* problem. Assume not, that is there exists a vector $\tilde{s}$, $\tilde{s} \neq \bar{s}$, which is the optimal solution for the *TFU-LIN* problem. For simplicity, assume that the two vectors differ by only two components (the following arguments also apply if they differ by more than two components). By Fact 2 there exist j, k and $\epsilon > 0$, such that $\tilde{s}(j) = \bar{s}(j) + \epsilon$ and $\tilde{s}(k) = \bar{s}(k) - \epsilon$. Note that $\bar{s}(j) \geq \bar{s}(k)$, otherwise we would get a contradiction for $\bar{s}$ being the solution to the *TFLA-LIN* problem. Therefore, from the uniqueness of the solution,

$$U(\bar{s}(j)) + U(\bar{s}(k)) < U(\tilde{s}(j)) + U(\tilde{s}(k))$$

$$\leq U(\bar{s}(j)) + U(\bar{s}(k))$$

where the second inequality is obtained by applying Fact 3. This contradicts the assumption that $\tilde{s} \neq \bar{s}$ and shows that the unique solution to the *TFLA-LIN* problem is also the solution to the *TFU-LIN* problem. Notice that this also shows the converse, that is, if $\tilde{s}$ is the (unique) solution to the *TFU-LIN* problem, then it also the solution to the *TFLA-LIN* problem. Indeed, any $\epsilon-$increase in one of $\tilde{s}$ components can only come at the expense of "weaker" components, which is exactly the characteristics of the optimal (lexicographically maximal) solution to the *TFLA-LIN* problem.

II. Appendix: Proof of Lemma 2

Let s^* denote the spending policy $s^* = \hat{Q}/K \; \forall \, i$. Below, we first show that the s^* policy is feasible under the *LS Conditions* and the *LS-gen Conditions*. Then we show that, where feasible, this policy is optimal.

We define $\tilde{B}(i) = [\sum_{n=0}^{i-1} Q(n)] - s^* \cdot (i-1)$ for $1 \leq i \leq K$.

[5]Since f is concave and twice differentiable, f' is a decreasing function, and $f'' < 0$. Hence, using Taylor expansion, $f(x+\epsilon) + f(y-\epsilon) = f(x) + f'(x)\epsilon + f(y) + f'(y)(-\epsilon) + o(\epsilon^2) \leq f(x) + f(y) + \epsilon(f'(x) - f'(y)) \leq f(x) + f(y)$.

Lemma 4: When $B_0 \geq |\min_{1\leq i\leq K} \tilde{B}(i)|$ and $C - B_0 \geq \max_{1\leq i\leq K} \tilde{B}(i)$ (the *LS Conditions*), the s^*-policy is feasible.

Proof: Assume the energy storage capacity is sufficiently large. Then the storage state under the s^*-policy at the beginning of the i^{th} slot is $B(i) = B_0 + [\sum_{n=0}^{i-1} Q(n)] - s^* \cdot (i-1) = B_0 + \tilde{B}(i)$. Thus, to avoid running out of energy, $B_0 \geq |\min_{1\leq i\leq K} \tilde{B}(i)|$ is needed. In addition, in each slot i, $C - B(i) \geq 0$ is required. Plugging in the expression for $B(i)$, the condition $C - B_0 \geq \max_{1\leq i\leq K} \tilde{B}(i)$ is obtained. ■

Lemma 5: When $B_0 \geq [\sum_i Q(i)](1 - 1/K)$ and $C - B_0 \geq [\sum_i Q(i)](1 - 1/K)$ (the *LS-gen Conditions*), the s^*-policy is feasible.

Proof: For all energy input vectors Q with a particular $\sum_i Q(i)$, when the *LS-gen Conditions* hold, *LS Conditions* hold, and thus the s^*-policy is feasible. This is easy to see by considering vectors that obtain $\min_{\sum_i Q(i)} \tilde{B}$ and $\max_{\sum_i Q(i)} \tilde{B}$. The minimum is obtained when there is no incoming energy during the whole period, except for the last time slot, that is, $Q(i) = 0 \; \forall \; i < K-1$ and $Q(K-1) = \sum_i Q(i)$. Therefore,

$$\min_{\sum_i Q(i)} \tilde{B} = -\sum_i Q(i) \cdot \frac{K-1}{K} = -\sum_i Q(i) \left(1 - \frac{1}{K}\right).$$

The maximum is obtained when $Q(1) = \sum_i Q(i)$ and $Q(i) = 0$ otherwise. Thus,

$$\max_{\sum_i Q(i)} \tilde{B} = \sum_i Q(i) - \sum_i Q(i) \cdot \frac{1}{K} = \sum_i Q(i) \left(1 - \frac{1}{K}\right).$$

Combining these two bounds with the *LS Conditions*, we obtain the *LS-gen Conditions*. ■

We now show that, when feasible, the s^* policy is the optimal solution to the *TFLA-LIN* and *TFU-LIN* problems.

Optimality of s^* policy for the TLFA-LIN problem: constraints (4) and (6) imply that any other vector s in the feasible domain has at least one i such that $s(i) \leq s^*$. Hence the s^*-policy is the optimal solution to the *TFLA-LIN* problem.

Optimality of s^* policy for the TFU-LIN problem: for the *TFU-LIN* problem, if the s^*-policy is feasible, constraints (3) – (7) reduce to the requirements that $s(i) \geq 0$ and $\sum_{i=0}^{K-1} s(i) \leq \sum_{i=0}^{K-1} Q(i) + (B_0 - B_K) = K \cdot s^*$. We need to prove that in this case the s^*-policy is the optimal solution to

$$\max_{s(i)} \bar{U} = \sum_{i=0}^{K-1} U_\alpha(s(i))$$

for $\alpha \geq 0$, in the feasible domain. We first present the proof for $\alpha = 1$. Recall that $U_1(s) = \log(s)$. Since $\lim_{s\to 0} \log(s) = -\infty$, in order to find the maximum it suffices to search in domains of the form

$$D_\epsilon = \{s = (s(0), \ldots, s(K-1)) :$$
$$\epsilon \leq s(i) \; \forall i, \; \sum_{i=0}^{K-1} s(i) \leq \sum_{i=0}^{K-1} Q(i) + (B_0 - B_K)\Big\}$$

for $\epsilon > 0$ small as we wish. $\bar{U}$ is a smooth function over D_ϵ and D_ϵ is a compact domain, therefore the maximum of the function $\bar{U}$ is attained. $\nabla \bar{U} \neq 0$, hence the maximum lies on the boundary. Analysis of the behavior of $\bar{U}$ on the boundary (using Lagrange multipliers, for example) yields a set of extreme points of the form

$$\left(\underbrace{\epsilon, \ldots, \epsilon}_{i}, \underbrace{\frac{\sum_{i=0}^{K-1} Q(i) + (B_0 - B_K)}{K-i}, \ldots, \frac{\sum_{i=0}^{K-1} Q(i) + (B_0 - B_K)}{K-i}}_{K-i} \right)$$

for $0 \leq i \leq K-1$, including all permutations. Since ϵ is close to zero, the maximum point is the point of the above form, where none of its coordinates equals to ϵ. There is only one point that fits this demand, when $i = 0$, and this is exactly the s^*-policy. The optimality of the s^*-policy for $U_\alpha(\cdot)$ with $\alpha > 1$ can be shown following exactly the same arguments. For $0 \leq \alpha < 1$, the proof is similar, but there is no need for defining D_ϵ domains.

III. Appendix: Proof of Lemma 3

Similarly to a single node case (see Appendix II), when the energy storage is large, constraints (3) – (7) for nodes u and v reduce to simple constraints on the sum of energy spending rates: $\sum_i s_u(i) \leq \sum_i Q_u(i) + (B_{0,u} - B_{K,u})$, $\sum_i s_v(i) \leq \sum_i Q_v(i) + (B_{0,v} - B_{K,v}.)$ Optimal single node energy spending rates $s_u(i)$ and $s_v(i)$ are assigned, by solving the *TFU-LIN* problem (or, equivalently, the *TLFA-LIN* problem), as $s_u^* \ \forall \ i$ and $s_v^* \ \forall \ i$, where $s_u^*(i)$ and $s_v^*(i)$ are according to the s^*-policy definition stated in Lemma 2. Thus, the per-slot sub-problem, either $\max_{r_u(i), r_v(i)} U(r_u(i)) + U(r_v(i))$ or $\max r(i)$, is solved with the same, largest possible, inputs $s_u(i)$, $s_v(i)$ for each slot, and thus the resulting data rates $r_u(i)$ and $r_v(i)$ are also the same for all slots i. The optimality of this solution for the *LTFU* and the *LTFL* problems follows directly from the arguments detailed in Appendix II.