Networking Low-Power Energy Harvesting Devices: Measurements and Algorithms

Maria Gorlatova, *Student Member*, *IEEE*, Aya Wallwater, and Gil Zussman, *Senior Member*, *IEEE*

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 energy harvested from indoor lights. We focus on characterizing the light 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 nontraditional 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 calculating *time fair* energy allocation in systems with *deterministic* energy inputs, as well as in systems where energy inputs are *stochastic*.

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

1 Introduction

Recent advances in the areas of solar, piezoelectric, and thermal energy harvesting [40], and in ultra-low-power wireless communications [49] will soon enable the realization of perpetual *energy harvesting wireless devices*. When networked together, they can compose rechargeable sensor networks [26], [41], [54], networks of computational RFIDs [20], and Energy Harvesting Active Networked Tags (EnHANTs) [15], [18]. Such networks will find applications in various areas, and therefore, the wireless industry is already engaged in the design of various devices (e.g., [5]).

In this paper, we focus on devices that harvest *environmental light energy*. Since there is a *three orders of magnitude* difference between the light energy available indoor and outdoor [18], [42], 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 two years, 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 this paper, we focus on developing algorithms for determining the energy spending rates and the data rates in various scenarios.

To describe our contributions, we introduce below several dimensions of the vast algorithm design space for energy-harvesting devices:

- *Environmental energy model:* deterministic 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:* subseconds 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 2).

1.1 Environmental Energy Models

The model representing harvested energy depends on various parameters such as the energy source (e.g., solar or kinetic), the properties of the environment, and the device's behavior (stationary, semistationary, or mobile). Fig. 1 provides examples of radiant (light) energy sources in different settings. In Fig. 1a, the energy availability is time-dependent and predictable. On the other hand, in Fig. 1b 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* [12], [27]. We will refer to ideal energy profiles that accurately represent the future as *deterministic*

- *M. Gorlatova and G. Zussman are with the Department of Electrical Engineering, Columbia University, 1300 S.W. Mudd, 500 West 120th Street, New York, NY 10027. E-mail: mag2206@columbia.edu, gil@ee.columbia.edu.*
- *A. Wallwater is with the Department Industrial Engineering and Operations Research, Columbia University, 1300 S.W. Mudd, 500 West 120th Street, New York, NY 10027. E-mail: aw2589@columbia.edu.*

Manuscript received 19 Oct. 2011; revised 7 June 2012; accepted 21 June 2012; published online 10 July 2012.
For information on obtaining reprints of this article, please send e-mail to: tmc@computer.org, and reference IEEECS Log Number TMC-2011-10-0580.
Digital Object Identifier no. 10.1109/TMC.2012.154.

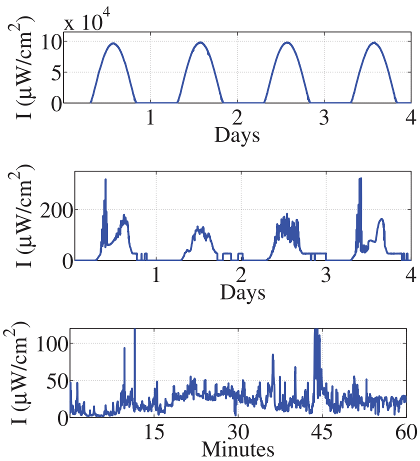

Fig. 1. Examples of different light energy sources: (top) *deterministic* profile (Las Vegas, NV [3], outdoors), (middle) *partially predictable profile* (New York, NY, a static indoor node), and (bottom) *stochastic process* (New York, NY, a mobile node in Times Square at *nighttime*).

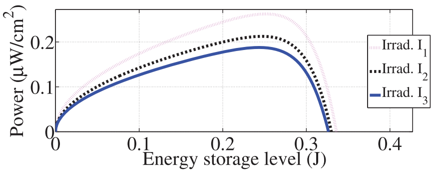

Fig. 2. An example of harvested power versus energy storage curves for a capacitor-based light energy harvesting system.

profiles, and to those that are inaccurate as *partially predictable profiles*.

Energy behavior that does not warrant a time-dependent profile appears in Fig. 1c, 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.

1.2 Energy Storage Types—Linear and Nonlinear

To operate when not directly powered by environmental energy, energy harvesting devices need *energy storage*: a *rechargeable battery* or a *capacitor*. *Rechargeable 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 [43]. Use of *capacitors* for storing harvested energy recently started gaining attention [18], [20], [26], [54]. In this paper, we consider *nonlinearity* of capacitor-based devices: in a simple capacitor-based device, the amount of power harvested depends both on the amount of energy provided (*irradiance*), and on *the amount of energy stored* [20], [35]. The nonlinear relations are demonstrated in Fig. 2.[1]

1.3 Storage Capacity, Decision Timescale, and Problem Size

Storage capacity versus amount of energy harvested. Energy storage capacity can vary from 0.16 J for an EnerChips device [1] to 4,700 J for an AA battery. The environmental energy availability also varies widely, from thousands of $\mathrm{J/cm^2/day}$ in sunny outdoor conditions to under $2\,\mathrm{J/cm^2/day}$ in indoor environments (see Section 4). 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).

1.4 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. The indoor light energy traces that we have collected are available at enhants.ee.columbia.edu and in the CRAWDAD repository [19]. To the best of our knowledge, this work is the first to present long-term indoor radiant energy measurements.

Second, we formulate *resource allocation problems* for energy-harvesting devices. The energy available to such devices often *varies in time* (e.g., throughout the day or among different days). Hence, in this paper, we aim to achieve *"smooth" allocation of resources along the time axis in the presence of varying environmental energy*.

We consider *deterministic energy profile* and *stochastic* environmental energy models, for battery-based systems and for *capacitor-based systems*, and focus on the cases of a *single node* and a node pair (*link*). For the deterministic profile environmental energy model, we use the lexicographic maximization and utility maximization frameworks to obtain the energy spending rate allocations for a node and the data rate allocations for a link. For the stochastic environmental energy model, we consider the case in which the energy inputs are i.i.d. random variables (e.g., a mobile device outdoors), and show how to treat it as an *average-cost Markov decision process* (MDP). 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 precomputed in advance. To the best of our

1. Solar cells have highly nonlinear output-versus-voltage characteristics, as demonstrated in Fig. 2. In simple systems, the voltage of the solar cell is determined by the voltage of the energy storage device. Within the battery operating range, the battery voltage is nearly constant. Capacitor voltage, on the other hand, is directly related to the energy stored on a capacitor, and changes substantially as energy is harvested or spent.

knowledge, our work is the first that models the *nonlinearity* of the capacitor-based system illustrated in Fig. 2. We provide numerical results that demonstrate its effect.

This paper is organized as follows: Section 2 reviews the related work. Section 3 presents the model and Section 4 describes the measurements. Sections 5 and 6 describe algorithms for the deterministic profile and stochastic process energy models, respectively. Section 7 presents the numerical results. We summarize and discuss future work in Section 8.

2 Related Work

Energy efficiency in wireless networks has long been a subject of research. In comparison, *energy harvesting* in wireless networks has only recently started gaining attention. The developments in this area include wireless energy-harvesting device design and development [26], [41], [46], [47], [53], [54], and exploration of theoretic and algorithmic approaches.

In this paper, we characterize *indoor light energy* for low-power energy-harvesting devices. Since large-scale *outdoor* solar panels have been used for decades, properties of the Sun's energy were examined in depth [3], [31], [42]. Practical outdoor solar energy considerations for energy-harvesting in sensor networks (e.g., light obstructions and scattering) were discussed in [47]. Until recently using *indoor* light energy for networking applications was considered impractical, and indoor light was studied mostly in the areas of architecture and ergonomics [21], [44]. 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 [2] and [20]. Photometric measurements, however, do not provide energetic characterization, and there is currently lack of data (e.g., traces) and analysis (e.g., variability, predictability, and correlations) regarding energy availability [42].

This paper also deals with *resource allocation for energy-harvesting devices*. The related work in this area can be classified according to the environmental energy model employed and related to the general settings described in the previous section:

- *Deterministic profile.* In [23] and [27], duty cycle adaptations (mostly for *single nodes*) are considered. Transmission power adaptation and transmission scheduling for a scenario with an energy-harvesting transmitter and two receivers are examined in [50] and [8], respectively. For a *network*, various metrics are considered including data collection rates [12], data retrieval rates [51], throughput maximization [11], and routing efficiency [33]. *Per-slot short-term* predictions are used to obtain (per-slot) data rates in [34].
- *Partially predictable profile.* While considering energy predictable, [11], [27], [34], [38] have provisions for adjustments in cases in which the predictions are inaccurate.
- *Stochastic process.* Dynamic activation of energy-harvesting sensors is described in [25] for a *single node*, and for a *cluster* in [28]. Admission and power allocation control policies are developed in [13]. Routing and scheduling policies are developed in [30]. Maximizing the utility of the *average* data rates via joint power allocation and energy management is examined in [24]. Energy allocation policies for source-channel coding are developed in [10].
- *Model-free approach.* Duty cycle adjustments for a single node (and under the linear storage model) are examined in [48]. A *capacitor*-based system is presented and the capacitor leakage is studied in [54].

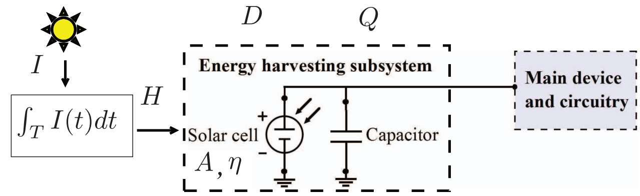

Fig. 3. A schematic diagram of the relationships between energy parameters: irradiance (I), irradiation (H), energy available to a device (D), and energy harvested by the device (Q).

We aim to allocate nodes' resources in a "smooth" way with respect to time. The need for policies that enable such behavior in energy-harvesting devices has been previously noted [12], [27], [37], [48]. Smoothing node duty cycles using a control theory approach is examined in [48]. Energy allocation vectors with minimal variance are sought in [37]. Both [37] and [48] consider linear energy storage models and focus mainly on single node scenarios. We note that the approach introduced in [52] for throughput optimization in QoS-constrained single node scenarios (for non-energy-harvesting devices) can also be used to achieve smooth energy allocation in energy-harvesting devices (where finite energy storage constraint can be related to the QoS buffer constraint [50]). A throughput optimization framework for energy-harvesting nodes [11], developed in parallel with our work, can also be extended to achieve smooth resource allocation. However, applications of these frameworks to energy-harvesting scenarios result in implicit assumptions of *linear* energy storage. The model developed in this paper allows incorporating general (linear and nonlinear) energy storage models. Furthermore, we formulate problems and present practical algorithms for both *single node* and *link* scenarios.

We note that resource allocation in energy harvesting devices has some similarities with power consumption scheduling in power networks (e.g., [32] and references therein). However, these works consider scenarios where energy sources are centralized and infinite. In contrast, in our settings energy availability is restricted, and is specific to each node and each time slot.

3 Model and Preliminaries

In this paper, we focus both on *light measurements* and on *resource allocation problems*. The relationships between variables characterizing energy availability are illustrated in Fig. 3. Table 1 summarizes the notation.

We focus on *discrete-time* models, where the time axis is separated into K slots, and a decision is made at the

TABLE 1
Nomenclature

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

beginning of a slot i ($i = \{0, 1, \ldots, K-1\}$). 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.

Our measurements record *irradiance*, radiant energy incident onto surface (in $\mathrm{W/cm^2}$), denoted by I. *Irradiation* H_T (in $\mathrm{J/cm^2}$) 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 *solar cell with given physical properties* (size, efficiency) can harvest in a time slot i is denoted by $D(i)$. For a solar cell with area A and efficiency η, $D(i) = A \cdot \eta \cdot H(i)$. For the numerical results presented in this paper we use $A = 10\ \mathrm{cm}^2$ and $\eta = 1\%$ (i.e., efficiency of an organic solar cell) [18].

The energy a node harvests from the environment in a time slot i is denoted by $Q(i)$. $Q(i)$ is a function of $D(i)$, and *may also depend on* $B(i)$. Specifically, for a *battery*-based device, $Q(i) = D(i)$. For a *capacitor*-based device, $Q(i) = q(D(i), B(i))$, where $q(D(i), B(i))$ is a *nonlinear* function of $D(i)$ (see Section 1.2). We refer to energy storage where $Q(i)$ is linear in $D(i)$ as *linear energy storage*, and to energy storage where $Q(i)$ is nonlinear in $D(i)$ as *nonlinear energy storage*. Functions $q(D(i), B(i))$ for a capacitor, derived from capacitors' electric properties, are shown in Fig. 2. To derive numerical results for nonlinear energy storage, we use $q(D(i), B(i)) = D(i) - D(i) \cdot (B(i) - C/2)^2 / (\beta_{\text{nonlin}} \cdot (C/2)^2)$, where β_{nonlin} is the energy storage nonlinearity parameter.[2] These functions have properties similar to the functions shown in Fig. 2.

The energy spending rate is denoted by $s(i)$. The "storage evolution" of energy harvesting devices can be expressed as

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

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 use these as subscripts for link enpoints' energy variables (e.g., C_u and $B_{0,u}$ correspond, respectively, to node u's energy storage capacity and initial storage state). We denote the data rates of u and v by $r_u(i)$ and $r_v(i)$, respectively. For a *single node*, we optimize the energy spending rates $s(i)$, which can provide inputs for determining *transmission power*, *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 allocate the energy or the data rates as much as possible in a *uniform way with respect to time*. We achieve this objective by using the *lexicographic maximization* and *utility maximization* frameworks. These frameworks are typically applied to achieve fairness *among nodes* [9], [12], [29], [34], [39]. In this paper, we apply them to achieve *time-fair resource allocation*. In the lexicographic maximization framework, 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). In utility maximization framework, we maximize $\sum_{i=0}^{K-1} U(s(i))$ (for a single node) or $\sum_{i=0}^{K-1} [U(r_u(i)) + U(r_v(i))]$ (for a link), where $U(\cdot)$ are *concave nondecreasing twice-differentiable continuous* functions (e.g., $U(\cdot) = \log(\cdot)$, $U(\cdot) = \sqrt{(\cdot)}$, $U(\cdot) = (\cdot)^{1-\alpha}/(1-\alpha), \alpha > 1$).[3] To derive numerical results, we use $U(\cdot) = \log(\cdot)$ or $U(\cdot) = \log(1 + (\cdot))$. In general, the solutions obtained by applying the two frameworks are not the same. The solutions are identical in certain cases, such as those examined in Lemma 1 and in Observation 1.

4 Characterizing Light Energy

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 $\mathrm{W/cm^2}$) in several indoor locations, and also study a set of shorter-term indoor and outdoor *mobile device* measurements. Table 2 provides a summary of the indoor measurement locations. The locations are shown schematically in Fig. 4. For the measurements, we use TAOS TSL230rd photometric sensors [4] installed on LabJack U3 DAQ devices. These photometric sensors have a high dynamic range, allowing to capture widely varying irradiance conditions. We verified the accuracy of the sensors with a NIST-traceable Newport 818-UV photodetector. In addition to the indoor measurements, we also analyze a set of outdoor irradiance traces provided by the US Department of Energy National Renewable Energy Laboratory (NREL) [3].

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 also be used as energy feeds to simulators and

2. We note that $Q(i) = D(i)$ for $\beta_{nonlin} \to \infty$.

3. We note that the utility maximization framework achieves *proportional fairness* for $U(\cdot) = \log(\cdot)$ and *max-min fairness* for $U(\cdot) = (\cdot)^{1-\alpha}/(1-\alpha)$ with $\alpha \to \infty$ [36].

TABLE 2
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, 13th 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 [3], Elizabeth City, NC.	Jan. 1, 2009 – Dec. 31, 2009	1517	787	1,750
O-2	*Outdoor*: *HSU* meteostation [3], Arcata, CA.	Jan. 1, 2009 – Dec. 31, 2009	1407	773	1,600

emulators. The traces are available at enhants.ee.columbia.edu and in the CRAWDAD repository [19].

4.1 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 $1\ \text{cm}^2$ area over the entire course of a day. Fig. 6 demonstrates the H_d values for setup *L-1*. Table 2 presents the average and the standard deviation values, $\overline{H_d}$ and $\sigma(H_d)$. These bit rates are calculated assuming *solar cell efficiency of* $\eta = 1\%$ (i.e., efficiency of an organic solar cell) and *solar cell size* $A = 10\ \text{cm}^2$. 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. Table 2 also shows the bit rate r a node would be able to maintain throughout a day when exposed to irradiation H_d. As an energy cost to communicate, 1 nJ/bit is used [18].[4] These bit rates can be seen as "communication budgets" for light energy harvesting devices (such as EnHANTs [18], [46], [53]) deployed in indoor environments.

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, $0 \le \alpha \le 1$. 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 data sets, 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 [31] and [45]. We examined whether the H_d values in the *indoor settings* are correlated with the weather data [6]. We determined statistically significant correlations for all setups except *L-2*.[5] 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, $\overline{H_d} = 1.63\ \text{J/cm}^2$ on weekdays, and $\overline{H_d} = 0.37\ \text{J/cm}^2$ 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 data sets, and determined statistically significant correlations for a number of setups. For example, for setups *L-1* and *L-2* located in the same room, $\rho = 0.56$ ($p < .001$), and for setups *L-1* and *L-5* facing in the same direction, $\rho = 0.71$ ($p < .001$).[6] 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*.

4.2 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, for setup *L-2* the irradiance is either 0 or 45 $\mu\text{W/cm}^2$ 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.

4. The bit rate is calculated as $r = A \cdot \eta \cdot H_d/(3{,}600 \cdot 24)/(10^{-9})$.

5. The correlation coefficients of the H_d values with the weather data are as follows: *L-1*: $\rho = 0.35$ ($p < .001$), *L-2*: no statistically significant ($p < .05$) correlation, *L-3*: $\rho = 0.137$ ($p < .05$), *L-4*: $\rho = 0.29$ ($p < .001$), *L-5*: $\rho = 0.24$ ($p < .05$), and *L-6*: $\rho = 0.71$ ($p < .001$).

6. We also detected the following statistically significant correlations: $\{L\text{-}1, L\text{-}3\}$: $\rho = -0.19$ ($p < .05$), $\{L\text{-}3, L\text{-}4\}$: $\rho = 0.52$ ($p < .001$), $\{L\text{-}3, L\text{-}6\}$: $\rho = 0.25$ ($p < .05$), and $\{L\text{-}4, L\text{-}6\}$: $\rho = 0.47$ ($p < .001$).

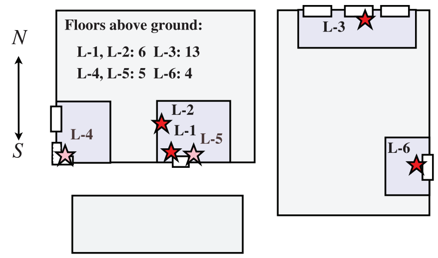

Fig. 4. A schematic diagram of the indoor irradiance measurement locations *L-1*-*L-6*.

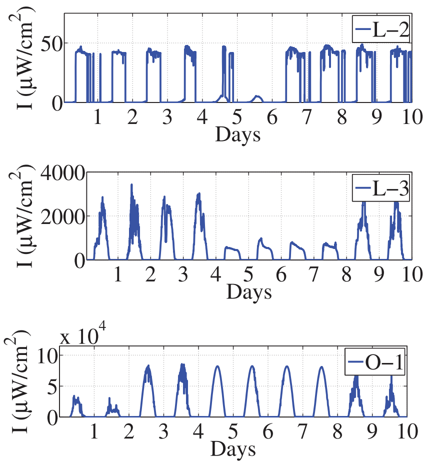

Fig. 5. Sample irradiance measurements in locations *L-2*, *L-3*, and *O-1* (2 March 2010-12 March 2010).

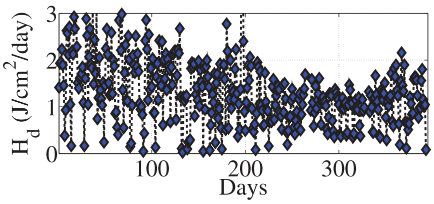

Fig. 6. Long-term daily irradiation (H_d) for setup *L-1* (15 August 2009-13 September 2010).

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 [7], [31]. 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 $j > i$). 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, $\sum_{i>j} H_T(i)$.[7] We determined that such correlations are present in indoor environments, but are generally 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 $\rho = 0.5$ ($p < .001$), while for the indoor setup *L-1* it is $\rho = 0.2$ ($p < .001$). 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 $\rho = 0.77$ ($p < .001$), while for the indoor setup *L-3* it is $\rho = 0.31$ ($p < .001$). 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.

4.3 Mobile Measurements

We have also conducted shorter term experiments for *mobile devices*. Table 3 provides a summary of the measurements conducted, demonstrating average irradiance $\overline{I}$, standard deviation of the irradiance $\sigma(I)$, and the corresponding sustainable bit rate r. It can be observed that energy availability differs drastically for different experimental conditions.

A sample irradiance trace for a measurement setup carried around Times Square in New York City at *nighttime* (measurement set *M-6*) was shown in Fig. 1c. Fig. 8 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 70 $\mu\mathrm{W/cm}^2$, while in sunny outdoor conditions it was 32 $\mathrm{mW/cm}^2$. Namely, the outdoor to indoor energy ratio was *more than 450 times*. In general, we observed that mobile devices' energy levels are diverse, poorly predictable, and could in some cases be represented by *stochastic* energy models.

5 Deterministic Energy Profile

In this section, we consider the *deterministic profile* energy model (similar to the models studied in [12], [27], and [38]). We formulate optimization problems that apply to both *linear* and *nonlinear* energy storage[8] for a *single node* and for a *link*, and introduce algorithms for solving the formulated problems.

5.1 Single Node: Optimizing Energy Spending

To achieve smooth energy spending for a node, we formulate the following problems where we optimize the node energy allocation vector $\{s(i)\}$ using the *utility maximization* and *lexicographic maximization* frameworks.

Time Fair Utility Maximization (TFU) Problem:

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

$$\text{s.t.:} \qquad s(i) \leq B(i)\ \forall\ i, \tag{3}$$

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

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

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

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

7. Additional correlation results are available in [16].

8. Recall that a *linear* energy storage model applies to a *battery* and that a *nonlinear* energy storage model may represent a *capacitor*.

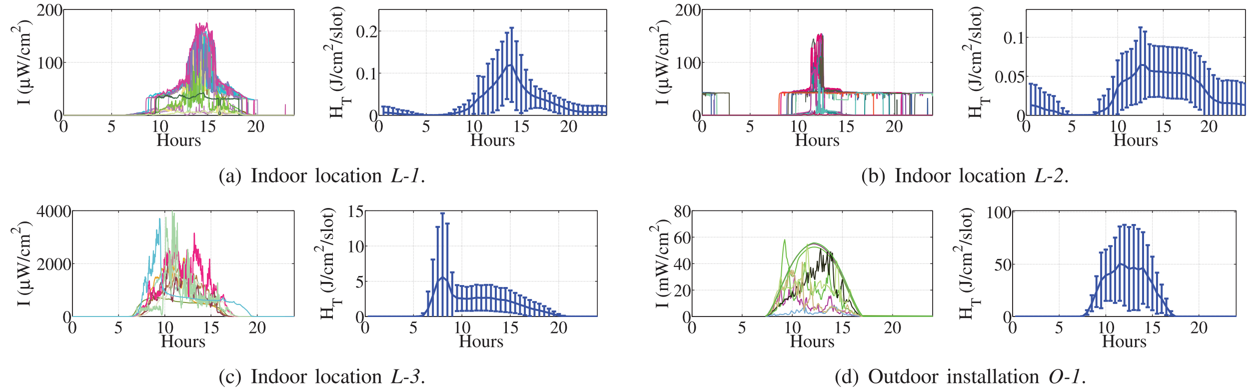

(a) Indoor location *L-1*. (b) Indoor location *L-2*.

(c) Indoor location *L-3*. (d) Outdoor installation *O-1*.

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 $\sigma(H_T)$.

TABLE 3
Mobile Light Energy Measurements—Average Irradiance and Achievable Bit Rates

Meas. index	Measurement description	Experiment start time	Experiment duration	$\overline{I}$, (μW/cm^2)	$\sigma(I)$	r (Kb/s)
M-1	Pedestrian walking around university campus (indoor and outdoor environments) carrying a measurement setup.	4/5/2010, 13:06	1:00:09	4,700	9,760	470
M-2	Commuting on public transit, measurement setup attached to a backpack, measurements outdoors, indoors (subway, train, office).	7/13/2010, 15:02	1:40:30	134	448	0.45
M-3	Car-based roadtrip, measurement setup attached to the dashboard.	7/16/2010, 12:26	2:57:00	11,416	4,370	38.1
M-4	Car-based errand running, measurement setup attached to the dashboard.	7/17/2010, 14:48	2:15:24	7,475	4,741	747.5
M-5	Car-based errand running, measurement setup attached to the dashboard.	7/18/2010, 09:58	2:06:00	16,472	5,563	1,647.4
M-6	Pedestrian walking in Times Square, New York City at nighttime, measurement setup attached to a backpack.	7/22/2010, 20:02	1:31:59	22.9	17	1.49

Recall that $U(s(i))$ is a concave nondecreasing function. Recall, additionally, that for *linear energy storage*, $Q(i) = D(i)$, and for *nonlinear energy storage*, $Q(i) = q(D(i), B(i))$ (see Section 3). Constraint (3) ensures that a node does not spend more energy than it has stored, (4) and (5) represent the energy storage evolution dynamics, and (6) sets the initial and final energy storage levels to B_0 and B_K.

Time Fair Lexicographic Assignment (TFLA) Problem:

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

$$\text{s.t.: constraints (3)-(7).}$$

Fig. 9 shows an example of node energy allocation vectors $\{s(i)\}$ obtained by solving the *TFU* and the *TFLA* problems. Fig. 9a shows the energy profile $\{D(i)\}$ used as an input to these problems. This energy profile corresponds to the light energy available in an indoor location *L-1* (see Table 2). Fig. 9b shows the energy allocation vectors $\{s(i)\}$ obtained by solving the *TFLA* problem under the *linear energy storage model* and by solving the *TFU* problem under the *nonlinear energy storage model*.[9]

Next, we provide a general algorithm (for *linear* and *nonlinear* energy storage) of a relatively high complexity, a faster algorithm for *linear energy storage*, and a very fast algorithm for *large linear energy storage*.

Assuming energy inputs and energy storage to be quantized, the *TFU* problem can be solved by the dynamic programming-based *Time Fair Rate Assignment* (*TFR*) algorithm (Algorithm 1).[10]

9. The solutions were obtained for the following parameters: $C = 0.5 \cdot \hat{Q}$, $B_0 = B_K = 0.4 \cdot C$, $U(s(i)) = \log(s(i))$, and $\beta_{\text{nonlin}} = 1.05$.

10. While other ways of solving the *TFU* problem can be considered, dynamic programming offers a natural solution approach.

Algorithm 1 Time Fair Rate Assignment (TFR).

$h(i, B) \leftarrow -\infty$, $s(i) \leftarrow 0 \ \forall \ i < K, \forall \ B$;
$h(K, B) \leftarrow -\infty \ \forall \ B < B_K$;
$h(K, B) \leftarrow 0 \ \forall \ B \geq B_K$;
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 $s(i) \ \forall \ i$

In the *TFR* algorithm, for each $\{i, B(i)\}$ we determine

$$h(i, B(i)) = \max_{s(i) \leq B(i)} [U(s(i)) + h(i+1, \min(B(i) + Q(i) - s(i), C))].$$

Going "backwards" from $i = K - 1$, we thus obtain a vector $\{s(0), \ldots, s(K-1)\}$ that maximizes $h(0, B_0)$; this is the optimal energy allocation vector. Recall that we denote the energy quantization resolution by Δ. In the TFR algorithm, we calculate $h(i, B(i))$ for each of the $K \cdot (C/\Delta)$ tuples

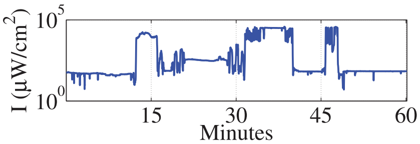

Fig. 8. Irradiance measurements recorded by a mobile device: a mix of indoor and outdoor conditions (note the $\log$ scale of the y-axis).

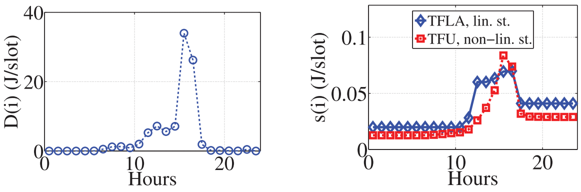

(a) (b)

Fig. 9. (a) Node energy profile $\{D(i)\}$, and (b) the corresponding energy allocation vectors $\{s(i)\}$ obtained by solving the *TFLA* problem (for linear energy storage model) and the *TFU* problem (for nonlinear energy storage model).

$\{i, B(i)\}$. Maximizing an instance of $h(i, B(i))$ requires considering all $s(i)$ such that $s(i) \le B(i) \le C$. Thus, for each tuple $\{i, B(i)\}$, the TFR algorithm performs at most C/Δ operations. The running time of the *TFR* algorithm is, therefore, $O(K \cdot [C/\Delta]^2)$.

For *linear energy storage* ($q(D(i), B(i)) = D(i)$, i.e., a *battery*), we refer to the *TFU* and the *TFLA* problems as *TFU-LIN* and *TFLA-LIN*. For these problems, we obtain the following Lemma, whose proof is given in Appendix I, which can be found on the Computer Society Digital Library at http://doi.ieeecomputersociety.org/10.1109/TMC.2012.154.

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

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 [9]. 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.

Algorithm 2 Progressive Filling (PF).

```
A_fix ← ∅; s(i) ← 0 ∀ i;
while ¬A_fix ≠ ∅ do
  for i = 0; i ≤ K − 1; i++; do
    if i ∈ ¬A_fix then
      s̃(j) ← s(j) ∀ j ∈ [0, K − 1]; s̃(i) ← s̃(i) + Δ;
      valid ← check_validity(s̃);
      if valid == TRUE then s(i) ← s̃(i);
      else A_fix := A_fix ∪ i;
function check_validity(s̃):
B(i) ← 0 ∀ i; B(0) ← B_0; valid ← TRUE;
for i = 1; i ≤ K; i++; do
  B(i) ← min(C, B(i − 1) + Q(i − 1) − s̃(i − 1));
  if s̃(i) > B(i) then valid ← FALSE;
if B(K) < B_K then valid ← FALSE;
return valid
```

Finally, when the *energy storage is large compared to the energy harvested*, the *TFLA-LIN and TFU-LIN* problems can be solved easily. Below we define *Large Storage (LS)* and *generalized Large Storage (LS-gen) Conditions*, and demonstrate that when they hold, the optimal policy is a simple one.[11] Let $s(i) = \hat{Q}/K \; \forall \, i$, and let $\tilde{B}(i) = [\sum_{j=0}^{i-1} Q(j)] - (i-1) \cdot s(i) \; \forall \, i \; 1 \le i \le K$.

Definition 1. *The LS Conditions hold, if*

$$B_0 \ge |\min_{1 \le i \le K} \tilde{B}(i)|$$

and

$$C - B_0 \ge \max_{1 \le i \le K} \tilde{B}(i).$$

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

The proofs of the following lemmas are given in Appendices II and III, which are available in the online supplemental material, respectively.

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

Lemma 3. *When the LS Conditions or the LS-gen Conditions hold, the optimal solution to the* TFU-LIN, *for* $U(s(i))$ *that are twice differential strictly concave on* $(0, \hat{Q}]$, *and that satisfy*

$$\begin{aligned} & (I) && U'(\cdot) > 0 \text{ on } (0, \hat{Q}] \text{ and } U'(0) = 0, \\ \text{or} \;\; & (II) && U'(\cdot) > 0 \text{ on } [0, \hat{Q}], \\ \text{or} \;\; & (III) && U'(\cdot) > 0 \text{ on } (0, \hat{Q}] \text{ and } \lim_{x \to 0} U(x) = -\infty, \end{aligned}$$

is $s(i) = \hat{Q}/K \; \forall \, i$.

Examples of $U(\cdot)$ that satisfy (*I*), (*II*), and (*III*) include: (*I*): $U(\cdot) = (\cdot)^{1-\alpha}/(1-\alpha)$ for $0 < \alpha < 1$ [36], (*II*): $U(\cdot) = \log(\alpha + (\cdot))$ for $\alpha > 0$, used, for $\alpha = 1$, in, for example, [11], [14], and (*III*): $U(\cdot) = \log(\cdot)$, used in, for example, [34]. Verifying that the *LS Conditions* (or the *LS-gen Conditions*) hold and determining the corresponding optimal policy is computationally inexpensive.

5.2 Link: Optimizing Data Rates

For a *link*, we formulate the following problems where we optimize data rate allocation vectors $\{r_u(i)\}$, $\{r_v(i)\}$.

Link Time Fair Utility Maximization (LTFU) problem:

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

$$\text{s.t.}: \quad c_{\mathrm{tx}} r_u(i) + c_{\mathrm{rx}} r_v(i) \le s_u(i) \; \forall \, i, \tag{10}$$

$$\begin{aligned} & c_{\mathrm{tx}} r_v(i) + c_{\mathrm{rx}} r_u(i) \le s_v(i) \; \forall \, i, \\ & u, v\text{: constraints (3)-(7).} \end{aligned} \tag{11}$$

11. To determine if the *LS Conditions* hold, a node needs to know $\{Q(0), \ldots, Q(K-1)\}$, 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 4.2).

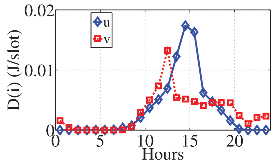

(a)

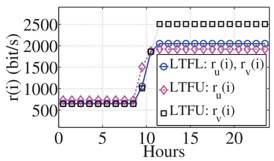

(b)

Fig. 10. (a) Energy profiles of link endpoints u and v, and (b) the corresponding data rate allocation vectors $\{r_u(i)\}$ and $\{r_v(i)\}$ obtained by solving the *LTFL* and the *LTFU* problems.

Link Time Fair Lexicographic Assignment (LTFL) problem:

Lexicographically maximize :

$$\{r_u(0),\dots,r_u(K-1),r_v(0),\dots,r_v(K-1)\} \quad (12)$$

s.t.: (10), (11); u, v: constraints (3)-(7).

Since the optimal solution to the *LTFL* problem is $\max-\min$ fair, it assigns the data rates such that $r_u(i)=r_v(i)\forall\ i$ (since for the $\max-\min$ fairness objective no increase in one of the rates can "outweigh" the decrease in the other). Thus, the *LTFL* problem can be restated as

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

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

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

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

Examples of solutions to the *LTFU* and *LTFL* problems are shown in Fig. 10. Fig. 10a shows the energy profiles of nodes u and v. These energy profiles correspond to the light energy available in indoor locations *L-1* and *L-2* (see Table 2) on the same day. Fig. 10b shows the data rate allocation vectors $\{r_u(i)\}$ and $\{r_v(i)\}$ obtained by solving the *LTFU* and the *LTFL* problems.[12]

In general, the solutions to the *LTFL* and *LTFU* problems are not the same. The following observation, whose proof appears in Appendix IV, which is available in the online supplemental material, identifies a case where the solutions are identical.

Observation 1. When $c_{\rm tx}=c_{\rm rx}$, the *LTFL* problem and the *LTFU* problem have the same solution.

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 $c_{\rm tx}r_u(i)+\ c_{\rm rx}r_v(i)=s_u(i)\le B_u(i), c_{\rm tx}r_v(i)+c_{\rm rx}r_u(i)=s_v(i)\le B_v(i)$, the *LTFR* algorithm determines, for each $\{i,B_u(i),B_v(i)\}$,

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

Vectors $\{r_u(0),\dots,r_u(K-1)\}$ and $\{r_v(0),\dots,r_v(K-1)\}$ 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* energy 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, due to (14), the subproblem solution is $r(i)=\min(s_u(i),s_v(i))/(c_{\rm tx}+c_{\rm rx})$. For the *LTFU-DRC* algorithm, a closed-form $O(1)$ solution to the subproblem can be obtained for each particular function $U(s(i))$. For example, for $U(s(i))=\log(s(i))$ with $c_{\rm tx}=\rho c_{\rm rx}$, $\rho>1$ [18], for the case of $s_v(i)=\gamma s_u(i)$, $0\le\gamma\le 1$, the optimal solution is either $\{r_u(i),r_v(i)\}=(s_u(i)/(c_{\rm rx}\cdot(\rho^2-1)))\{\rho-\gamma,\gamma\cdot\rho-1\}$ or $\{r_u(i),\ r_v(i)\}=\{s_v(i)/(2\cdot c_{\rm rx}),s_v(i)/(2\cdot c_{\rm tx})\}$.

For linear energy storage, when *the storage is large compared to the energy harvested* for both u and v, 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. The proofs of the following lemmas are given in Appendices V and VI, which are available in the online supplemental material, respectively.

Lemma 4. *If the* LS Conditions *or the* LS-gen Conditions *hold for nodes u and v, the* LTFL-DRC *algorithm obtains the optimal solution to the* LTFL *problem.*

Lemma 5. *If the* LS Conditions *or the* LS-gen Conditions *hold for nodes u and v, for $U(\cdot)$ that are twice differential strictly concave on $(0,R]$ where $R=\max\{\hat{Q}_u,\hat{Q}_v\}/\min\{c_{\rm tx},c_{\rm rx}\}$, and that satisfy*

$$\begin{aligned}&(I) && U'(\cdot)>0 \text{ on } (0,R] \text{ and } U'(0)=0,\\ \text{or}\ &(II) && U'(\cdot)>0 \text{ on } [0,R],\\ \text{or}\ &(III) && U'(\cdot)>0 \text{ on } (0,R] \text{ and } \lim_{x\to 0}U(x)=-\infty,\end{aligned}$$

12. The solutions were obtained for the following parameters: $C_u=C_v=0.5\cdot\hat{Q}_u$, $B_{0,u}=B_{0,v}=B_{K,u}=B_{K,v}=0.25\cdot C_u$, $c_{\rm tx}=0.1$ nJ/bit, $c_{\rm rx}=1$ nJ/bit [18], $U(r(i))=\log(r(i))$, and $Q(i)=D(i)$ (linear energy storage).

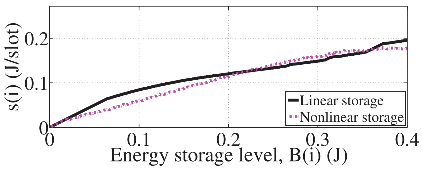

Fig. 11. Optimal energy spending rates $s(i)$ corresponding to different energy storage levels, obtained by solving the *SPD* problem for *linear* and *nonlinear* energy storage models.

the *LTFU-DRC* algorithm obtains the optimal solutions to the *LTFU* problem.

In Section 7, we provide numerical results demonstrating the rate allocation vectors $\{r_u(i), r_v(i)\}$ obtained by using the *DRC* algorithms to solve the *LTFU* and *LTFL* problems.

6 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 $[d_1, \ldots, d_M]$ with probability $[p_1, \ldots, p_M]$. 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 (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). We formulate the control problems and determine corresponding policies for a single node and for a *link*. The formulations apply to *linear* and *nonlinear* (e.g., a *capacitor*) 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 an *average cost MDP*, and can be solved with standard MDP solution techniques. For example, using value iteration approach and applying dynamic programming, we consider a large number of slots K, and going "backwards" from $i = K-1$, for each $\{i, B(i)\}$, determine

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

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 values of i) $s(i)$ value that approaches the optimal [22]. Although policy calculations are computationally expensive (the running time of this 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. Fig. 11 presents example optimal energy spending policies obtained by solving the *SPD* problem for *linear* and *nonlinear* energy storage models. The daily irradiation H_d for setup *L-1* (see Fig. 6) is used as the random variable D.[13]

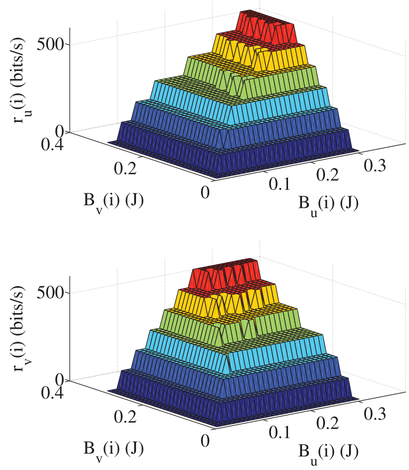

Fig. 12. Optimal communication rates $r_u(i)$ (top) and $r_v(i)$ (bottom) corresponding to different energy storage states, obtained by solving the *LSPD* problem.

For a *link*, we define the following problem.

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 value iteration approach, we determine, for each $\{i, B_u(i), B_v(i)\}$,

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

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. Fig. 12 demonstrates example optimal link rate assignment policy $\{r_u(i), r_v(i)\}$ as a function of $\{B_u(i), B_v(i)\}$ obtained by solving the *LSPD* problem. The daily irradiation H_d for

13. The solutions were obtained for the following parameters: $C = 2.7 \cdot \mathbb{E}(D(i))$, $U(s(i)) = \log(1 + s(i))$, and $\beta_{\text{nonlin}} = 1.3$.

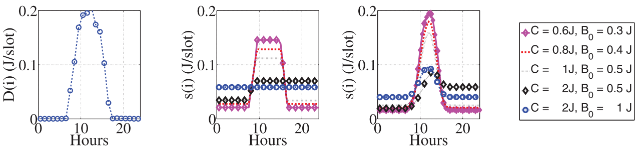

(a) (b) (c)

Fig. 13. (a) Energy profile $\{D(i)\}$, and energy spending rate assignments $\{s(i)\}$, obtained by (b) solving the *TFLA-LIN* and *TFU-LIN* problems, and by (c) solving the *TFU* problem for *nonlinear* energy storage.

setup *L-1* (see Fig. 6) is used as the random variable D_u and the random variable D_v.[14]

The MDP formulations can be easily extended to consider other parameters, such as the cost to change the energy spending rates $s(i)$, or the cost to change data rates $r_u(i)$ and $r_v(i)$.

7 Numerical Results

This section provides numerical results demonstrating the use of the algorithms described in Section 5. Measurement traces described in Section 4 are used as inputs to the algorithms.

Fig. 13 shows the optimal energy spending allocation vectors $\{s(i)\}$ for the *TFU* and the *TFLA* problems presented in Section 5.1, for different values of energy storage capacity C and initial energy storage state B_0.[15] The energy profile $\{D(i)\}$ used as an input to these algorithms is shown in Fig. 13a. It corresponds to the average daily energy profile for the indoor location *L-3* (see Table 2). Fig. 13b demonstrates the energy spending rate allocations $\{s(i)\}$ that solve the *TFLA-LIN* and the *TFU-LIN* problems (that is, *linear* energy storage model). These spending rates were obtained using the *PF* algorithm. It can be observed that larger energy storage allows for "smoother" energy allocation. For this energy profile $\{D(i)\}$, the *LS Conditions* described in Section 5 are matched when $C = 2$ J and $B_0 = 1$ J. It can be observed that in this case the energy spending rate allocation vector $\{s(i)\}$ corresponds to the optimal policy given by Lemmas 2 and 3. Fig. 13c shows the optimal solutions of the *TFU* problem with *nonlinear energy storage* (for $\beta_{\text{nonlin}} = 1.1$) obtained using the *TFR* algorithm. Such a system has not been analyzed before.

Fig. 14 shows the numerical results for the link data rate determination problems presented in Section 5.2. The energy profiles of indoor setups *L-1* and *L-2* (see Fig. 10a) were used as inputs to the algorithms. The optimal solutions to the *LTFL* and the *LTFU* problems for linear energy storage model have been shown in Fig. 10. Fig. 14a shows the optimal solution to the *LTFU* problem for *nonlinear energy storage* (for $\beta_{\text{nonlin}} = 1.1$) obtained using the *LTFR* algorithm. Fig. 14b shows the communication rate assignment vectors $\{r_u(i)\}$ and $\{r_v(i)\}$ calculated using a simple *LTFU-DRC* algorithm for linear energy storage. In this example, the *LTFU-DRC* algorithm obtains data rate assignments $\{r_u(i), r_v(i)\}$ that are similar to those obtained by optimally solving the *LTFU-LIN* problem.

14. The solutions were obtained for the following parameters: $C = 2.1 \cdot \mathbb{E}(D(i))$, $c_{\text{tx}} = 1$ nJ/bit, $c_{\text{rx}} = 2$ nJ/bit, $U(s(i)) = \log(1 + s(i))$, and $Q(i) = D(i)$ (linear energy storage).

15. The solutions were obtained for the following parameters: $B_K = B_0$ and $U(s(i)) = log(s(i))$.

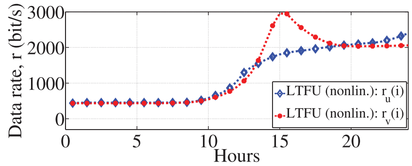

(a) Optimal solution to the *LTFU* problem, *nonlinear storage*.

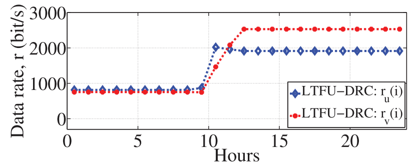

(b) Solutions obtained by the *LTFU-DRC* heuristic algorithm, *linear storage*.

Fig. 14. Communication rates $r_u(i)$ and $r_v(i)$ obtained by solving the *LTFL* and the *LTFU* problems.

8 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 described the first *long-term indoor radiant energy measurements campaign* that provides useful energy traces, as well as insights into the design of systems and algorithms. We developed algorithms for deterministic environments that uniquely determine the energy management and data rate control policies for *linear* and *nonlinear* energy storage models, for *single node* and *node pair* (*link*) scenarios. 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 precomputed decision 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 1. 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. In our ongoing work, we are analyzing the performance of simple policies for energy harvesting devices for single node and link cases [14]. We plan to develop simple energy-harvesting-aware algorithms for networks of nodes, additionally considering various other problem dimensions. Moreover, we plan to evaluate these algorithms in an energy harvesting active networked tags (EnHANTs) testbed that we are currently building [46], [53].

Acknowledgments

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 by an NSERC CGS grant. The authors thank Matthias Bahlke, Enlin Xu and Michael Zapas for their assistance with light energy measurements. They thank Peter Kinget and John Kymissis for helpful discussions.

REFERENCES

[1] "Cymbet EnerChips," www.cymbet.com/content/products.asp, 2011.

[2] "Intel Lab Data," db.csail.mit.edu/labdata/labdata.html, 2013.

[3] "Measurement and Instrumentation Data Center, National Renewable Energy Laboratory (NREL), US DOE," www.nrel.gov/midc, 2013.

[4] "TAOS TSL230rd Programmable Light-to-Frequency Converter," www.taosinc.com, 2013.

[5] "Texas Instruments MSP430 Solar Energy Harvesting Development Tool," focus.ti.com/docs/toolsw/folders/print/ez430-rf2500-seh.html, 2013.

[6] "Weather Underground Weather History and Data Archive," www.wunderground.com/history, 2013.

[7] M.I. Ali, B.M. Al-Hashimi, J. Recas, and D. Atienza, "Evaluation and Design Exploration of Solar Harvested-Energy Prediction Algorithm," *Proc. Conf. Design, Automation and Test in Europe (DATE '10),* Mar. 2010.

[8] M.A. Antepli, E. Uysal-Biyikoglu, and H. Erkal, "Optimal Scheduling on an Energy Harvesting Broadcast Channel," *Proc. IEEE Int'l Symp. Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks (WiOpt '11),* May 2011.

[9] D. Bertsekas and R. Gallager, *Data Networks,* second ed. Prentice-Hall, 1992.

[10] P. Castiglione, O. Simeone, E. Erkip, and T. Zemen, "Energy-Neutral Source-Channel Coding in Energy-Harvesting Wireless Sensors," *Proc. IEEE Int'l Symp. Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks (WiOpt '11),* May 2011.

[11] S. Chen, P. Sinha, N. Shroff, and C. Joo, "Finite-Horizon Energy Allocation and Routing Scheme in Rechargeable Sensor Networks," *Proc. IEEE INFOCOM,* Apr. 2011.

[12] K.-W. Fan, Z. Zheng, and P. Sinha, "Steady and Fair Rate Allocation for Rechargeable Sensors in Perpetual Sensor Networks," *Proc. Sixth ACM Conf. Embedded Network Sensor Systems (SenSys '08),* Nov. 2008.

[13] 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, Feb. 2010.

[14] M. Gorlatova, A. Bernstein, and G. Zussman, "Performance Evaluation of Resource Allocation Policies for Energy Harvesting Devices," *Proc. IEEE Int'l Symp. Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks (WiOpt '11),* May 2011.

[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 Comm.,* vol. 17, no. 6, pp. 18-25, Dec. 2010.

[16] M. Gorlatova, A. Wallwater, and G. Zussman, "Networking Low-Power Energy Harvesting Devices: Measurements and Algorithms," Technical Report 2010-12-15, Columbia Univ., http://enhants.ee.columbia.edu/images/papers/CU_EE_2010_12_15.pdf, Dec. 2010.

[17] M. Gorlatova, A. Wallwater, and G. Zussman, "Networking Low-Power Energy Harvesting Devices: Measurements and Algorithms," *Proc. IEEE INFOCOM,* Apr. 2011.

[18] M. Gorlatova, P. Kinget, I. Kymissis, D. Rubenstein, X. Wang, and G. Zussman, "Challenge: Ultra-Low-Power Energy-Harvesting Active Networked Tags (EnHANTs)," *Proc. ACM MobiCom,* Sept. 2009.

[19] M. Gorlatova, M. Zapas, E. Xu, M. Bahlke, I. Kymissis, and G. Zussman, "CRAWDAD Data Set Columbia/Enhants (v. 2011-04-07)," http://crawdad.cs.dartmouth.edu/columbia/enhants, Apr. 2011.

[20] J. Gummeson, S.S. Clark, K. Fu, and D. Ganesan, "On the Limits of Effective Micro-Energy Harvesting on Mobile CRFID Sensors," *Proc. ACM MobiSys,* June 2010.

[21] 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.

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

[23] J. Hsu, S. Zahedi, A. Kansal, M. Srivastava, and V. Raghunathan, "Adaptive Duty Cycling for Energy Harvesting Systems," *Proc. IEEE Int'l Symp. Low Power Electronics and Design (ISLPED '06),* Oct. 2006.

[24] L. Huang and M. Neely, "Utility Optimal Scheduling in Energy Harvesting Networks," *Proc. ACM MobiHoc,* May 2011.

[25] N. Jaggi, K. Kar, and A. Krishnamurthy, "Rechargeable Sensor Activation under Temporally Correlated Events," *Proc. IEEE Int'l Symp. Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks (WiOpt '07),* Apr. 2007.

[26] X. Jiang, J. Polastre, and D. Culler, "Perpetual Environmentally Powered Sensor Networks," *Proc. IEEE Fourth Int'l Symp. Information Processing in Sensor Networks (IPSN '05),* Apr. 2005.

[27] A. Kansal, J. Hsu, S. Zahedi, and M.B. Srivastava, "Power Management in Energy Harvesting Sensor Networks," *ACM Trans. Embedded Computing Systems,* vol. 6, no. 4, 2007.

[28] K. Kar, A. Krishnamurthy, and N. Jaggi, "Dynamic Node Activation in Networks of Rechargeable Sensors," *IEEE/ACM Trans. Networking,* vol. 14, no. 1, pp. 15-26, Feb. 2006.

[29] F.P. Kelly, A. Maulloo, and D. Tan, "Rate Control in Communication Networks: Shadow Prices, Proportional Fairness and Stability," *J. Operational Research Soc.,* vol. 49, pp. 237-252, 1998.

[30] M. Khouzani, S. Sarkar, and K. Kar, "Optimal Routing and Scheduling in Multihop Wireless Renewable Energy Networks," *Proc. Information Theory and Application Workshop (ITA),* Feb. 2011.

[31] 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.,* vol. 167, no. 4, pp. 16-23, 2009.

[32] N. Li, L. Chen, and S. Low, "Optimal Demand Response Based on Utility Maximization in Power Networks," *Proc. IEEE Power and Energy Soc. General Meeting,* July 2011.

[33] L. Lin, N. Shroff, and R. Srikant, "Asymptotically Optimal Energy-Aware Routing for Multihop Wireless Networks with Renewable Energy Sources," *IEEE/ACM Trans. Networking,* vol. 15, no. 5, pp. 1021-1034, Oct. 2007.

[34] R.-S. Liu, P. Sinha, and C.E. Koksal, "Joint Energy Management and Resource Allocation in Rechargeable Sensor Networks," *Proc. IEEE INFOCOM,* Mar. 2010.

[35] X. Lu and S. Yang, "Solar Energy Harvesting for ZigBee Electronics," *Proc. Sustainability in Energy and Buildings Conf. (SEB '09),* Apr. 2009.

[36] J. Mo and J. Walrand, "Fair End-to-End Window-Based Congestion Control," *IEEE/ACM Trans. Networking,* vol. 8, no. 5, pp. 556-567, Oct. 2000.

[37] D. Noh, L. Wang, Y. Yang, H. Le, and T. Abdelzaher, "Minimum Variance Energy Allocation for a Solar-Powered Sensor System," *Proc. IEEE Fifth Int'l Conf. Distributed Computing in Sensor Systems (DCOSS '09),* June 2009.

[38] D. Noh and T. Abdelzaher, "Efficient Flow-Control Algorithm Cooperating with Energy Allocation Scheme for Solar-Powered WSNs," *Wireless Comm. and Mobile Computing,* 2010.

[39] D. Palomar and M. Chiang, "A Tutorial on Decomposition Methods for Network Utility Maximization," *IEEE J. Selected Areas Comm.,* vol. 24, no. 8, pp. 1439-1451, Aug. 2006.

[40] J. Paradiso and T. Starner, "Energy Scavenging for Mobile and Wireless Electronics," *IEEE Pervasive Computing,* vol. 4, no. 1, pp. 18-27, Jan. 2005.

[41] V. Raghunathan, A. Kansal, J. Hsu, J. Friedman, and M. Srivastava, "Design Considerations for Solar Energy Harvesting Wireless Embedded Systems," *Proc. IEEE Fourth Int'l Symp. Information Processing in Sensor Networks (IPSN '05),* Apr. 2005.

[42] J. Randall, *Designing Indoor Solar Products,* first ed. Wiley, 2005.

[43] R. Rao, S. Vrudhula, and D. Rakhmatov, "Battery Modeling for Energy Aware System Design," *IEEE Computer,* vol. 36, no. 12, pp. 77-87, Dec. 2003.

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

[45] N. Sharma, J.J. Gummeson, D. Irwin, and P. Shenoy, "Cloudy Computing: Leveraging Weather Forecasts in Energy Harvesting Sensor Systems," *Proc. IEEE Seventh Int'l Comm. Soc. Conf. Sensor Mesh and Ad Hoc Comm. and Networks (SECON '10),* June 2010.

[46] G. Stanje, P. Miller, J. Zhu, A. Smith, O. Winn, R. Margolies, M. Gorlatova, J. Sarik, M. Szczodrak, B. Vigraham, L. Carloni, P. Kinget, I. Kymissis, and G. Zussman, "Demo: Organic Solar Cell-Based EnHANT Prototypes," *Proc. Ninth ACM Conf. Embedded Networked Sensor Systems (SenSys '11),* Nov. 2011.

[47] J. Taneja, J. Jeong, and D. Culler, "Design, Modeling, and Capacity Planning for Micro-Solar Power Sensor Networks," *Proc. IEEE Int'l Conf. Information Processing in Sensor Networks (IPSN '08),* Apr. 2008.

[48] C. Vigorito, D. Ganesan, and A. Barto, "Adaptive Control of Duty Cycling in Energy-Harvesting Wireless Sensor Networks," *Proc. IEEE Int'l Comm. Soc. Conf. Sensor Mesh and Ad Hoc Comm. and Networks (SECON '07),* June 2007.

[49] D. Wentzloff, F. Lee, D. Daly, M. Bhardwaj, P. Mercier, and A. Chandrakasan, "Energy Efficient Pulsed-UWB CMOS Circuits and Systems," *Proc. IEEE Int'l Conf. Ultra-Wideband (ICUWB '07),* Sept. 2007.
[50] J. Yang, O. Ozel, and S. Ulukus, "Broadcasting with a Battery Limited Energy Harvesting Rechargeable Transmitter," *Proc. IEEE Int'l Symp. Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks (WiOpt '11),* May 2011.
[51] 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," *Proc. ACM MobiSys,* June 2009.
[52] M. Zafer and E. Modiano, "A Calculus Approach to Energy-Efficient Data Transmission with Quality-of-Service Constraints," *IEEE/ACM Trans. Networking,* vol. 17, no. 3, pp. 898-911, June 2009.
[53] J. Zhu, G. Stanje, R. Margolies, M. Gorlatova, J. Sarik, Z. Noorbhaiwala, P. Miller, M. Szczodrak, B. Vigraham, L. Carloni, P. Kinget, I. Kymissis, and G. Zussman, "Demo: Prototyping UWB-Enabled EnHANTs," *Proc. ACM MobiSys,* June 2011.
[54] T. Zhu, Z. Zhong, Y. Gu, T. He, and Z.-L. Zhang, "Leakage-Aware Energy Synchronization for Wireless Sensor Networks," *Proc. ACM MobiSys,* June 2009.

Maria Gorlatova received the BSc (summa cum laude) and MSc degrees in electrical engineering from the University of Ottawa, Ontario, Canada, in 2004 and 2006, respectively. Since 2008, she has been working toward the PhD degree in electrical engineering at Columbia University, New York. Her research interests include networking low-power devices. Her industry experience includes positions at Defence R&D Canada and Telcordia Technologies. She is a recipient of the Columbia University Presidential Fellowship, the 2012 Google Anita Borg Fellowship, the Canadian Graduate Scholar (CGS) NSERC Fellowships, and a corecipient of the 2011 ACM SenSys Best Student Demo Award and the 2011 IEEE Communications Society Award for Outstanding Paper on New Communication Topics. She is a student member of the IEEE.

Aya Wallwater received the BA degree in mathematics in 2007 (cum laude) and the MSc degree in mathematics in 2010, both from the Technion, Israel. She is currently working toward the PhD degree in industrial engineering and operation research at Columbia University. Her research interests include the area of stochastic systems and efficient rare events simulation.

Gil Zussman received the BSc degree in industrial engineering and management and the BA degree in economics from the Technion, Israel, in 1995 (both summa cum laude), the MSc degree (summa cum laude) in operations research from Tel-Aviv University, Israel, in 1999, and the PhD degree in electrical engineering from the Technion, Israel, in 2004. Between 2004 and 2007, he was a postdoctoral associate at the Massachusetts Instutite of Technology. He is currently an associate professor of electrical engineering at Columbia University. His research interests include the area of wireless networks. He is a corecipient of the IFIP Networking 2002 Best Student Paper Award, the OPNETWORK 2002 and ACM SIGMETRICS 2006 Best Paper Awards, and the 2011 IEEE Communications Society Award for Outstanding Paper on New Communication Topics. He received the Marie Curie Outgoing International Fellowship, the Fulbright Fellowship, the DTRA Young Investigator Award, and the US National Science Foundation CAREER Award, and was a member of a team that won first place in the 2009 Vodafone Foundation Wireless Innovation Project competition. He is a senior member of the IEEE.